

IBM System i5 V5R4 Technical Overview

More power to the System i family and
i5/OS

Expanding innovative solutions

More management tools for
your IT solutions and business



Jim Cook
Donald Otis
Roshan Tarrosa
Louis Cuypers
Joonwoo Kim
Jos Vermaere
Will Witten

Redbooks



International Technical Support Organization

IBM System i5 V5R4 Technical Overview

February 2007

Note: Before using this information and the product it supports, read the information in “Notices” on page xiii.

First Edition (February 2007)

This edition applies to i5/OS Version V5R4 and other i5/OS-based licensed programs at the corresponding release. WebSphere-based application servers and application development tools include up to Version 6.1.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xiii
Trademarks	xiv
Preface	xvii
The team that wrote this redbook	xviii
Become a published author	xx
Comments welcome	xxi
Chapter 1. Introducing the IBM System i5 hardware and i5/OS V5R4 overview	1
1.1 Overview of System i5 2006 announcements	2
Chapter 2. IBM System i5 hardware	5
2.1 System i5 models: more power and flexibility	7
2.2 System i5 hardware changes	8
2.2.1 System i5 520 hardware changes	8
2.2.2 System i5 550 hardware changes	10
2.2.3 System i5 570 hardware changes	12
2.2.4 System i5 595 hardware changes	14
2.3 System i5 memory features	15
2.4 IBM System i5 maximums summary	17
2.5 System i5 models: the IOP-less approach and strategy	18
2.5.1 Media Devices that can attach to adapters with no controlling IOP	22
2.5.2 RAID5 adapters with write cache	22
2.5.3 Internet Small Computer System Interface controller	23
2.5.4 Small Computer System Interface controller for disks with 90 MB cache	25
2.5.5 High performance disk controllers with 1.5 GB write cache, 1.6 GB read cache ..	27
2.5.6 IBM TotalStorage EXP24 disk enclosures	28
2.5.7 Disk Controllers for EXP24 disk enclosures	29
2.5.8 Four Gbps fibre channel adapters	30
2.5.9 Existing communication adapters without input/output processor	32
2.6 AIX and Linux	33
2.7 iSCSI support	35
2.8 V5R4 POWER5-based systems and i5/OS console support	35
2.8.1 Hardware Management Console	36
2.8.2 Thin console	36
2.9 VXA-32 Tape Drive support	37
2.10 IBM System Storage TS3400 Tape Library	37
2.11 CPU power redundancy	39
2.12 Rack-related changes during 2006	39
2.13 ytty	39
2.14 IBM System i5 code administration and terminology	40
2.15 V5R4 RAID protection	42
2.15.1 RAID level 6	46
2.16 SAN and TotalStorage DS6000 and DS8000	48
2.16.1 Flashcopy	49
2.16.2 Peer-to-Peer Remote Copy	51
2.17 10 GB Ethernet LAN IOA	52
2.18 IBM Configurator for ebusiness	53
2.18.1 IOA-level mirroring #0308	54

2.18.2 Custom Data Protection Specify #0269	55
2.18.3 5372-IS5 Partition Preload Specifies	55
2.19 Customer Specified Placement and LPAR preload	55
2.20 System i5 editions	56
2.20.1 Editions at a glance	57
2.20.2 IBM System i5 520 Collaboration Edition	59
2.20.3 IBM System i5 Solution editions	60
2.20.4 Capacity Backup (CBU) Editions	61
2.21 Upgrade paths from one edition to another edition	65
Chapter 3. i5/OS work management	69
3.1 Work management enhancements	70
3.1.1 Job log management	70
3.1.2 Subsystem enhancements	73
3.1.3 Call stack display enhancements	74
3.1.4 New job interrupt exit program support	76
3.1.5 WRKACTJOB display now shows job's current user profile	76
3.2 Daylight saving time support for March 2007	78
3.3 Memory affinity	79
Chapter 4. DB2 Universal Database for i5/OS	83
4.1 Strategic initiatives	84
4.2 Application flexibility and portability	84
4.2.1 Integration of free-format RPG and SQL	84
4.2.2 Standards compliance	85
4.2.3 Improved DB2 family compatibility	86
4.3 Database application development enhancements	86
4.3.1 iSeries .NET provider	86
4.3.2 DB2 Connect Unlimited Edition for iSeries	90
4.3.3 IBM EWLM support added to CLI, DRDA, .NET, ODBC, JDBC	91
4.3.4 Stored procedures	95
4.4 On Demand and availability	96
4.5 Performance and performance management enhancements	97
4.5.1 Performance enhancements	97
4.5.2 Enhanced performance tooling	100
4.5.3 DB2 performance analysis with iSeries Navigator	101
4.5.4 Tools to monitor and analyze SQL performance	109
4.5.5 Tools to show current activity	113
4.5.6 Tools to show current database limits and sizes	113
4.5.7 Content Manager	114
4.6 Additional information	116
Chapter 5. i5/OS integration of IBM System x-base servers	117
5.1 IBM System x integration enhancements	118
5.1.1 Support for integrated Intel servers	120
5.1.2 Licensed program product changes	120
5.1.3 Support withdrawal	120
5.1.4 Support for Windows 2003 Volume Shadow Copy	120
5.1.5 Expand disk drive support	121
5.2 System x and BladeCenter support via iSCSI	121
5.2.1 iSCSI introduction	123
5.2.2 iSCSI installation requirements	124
5.2.3 iSCSI-attached xSeries or BladeCenter installation	127
5.2.4 iSeries Navigator support for iSCSI connections	132

Chapter 6. i5/OS-based security	135
6.1 V5R4 i5/OS security overview	136
6.2 System integrity	136
6.3 Intrusion detection	139
6.4 Cryptographic support	141
6.4.1 Cryptographic support for AS/400 (5722-CR1)	141
6.5 Enterprise identity mapping	142
6.6 Improved antivirus scanning	144
6.7 Portable Utilities for i5/OS	144
Chapter 7. i5/OS-based backup and recovery	147
7.1 i5/OS: saving and restoring data	148
7.2 Spooled file support	148
7.3 Dynamic device allocation	149
7.4 IFS parallel save/restore support	149
7.4.1 Automate parent directory creation	150
7.4.2 IFS restore performance improvements	150
7.5 New SAVSYSINF and RSTSYSINF commands	150
7.5.1 SAVSYSINF and RSTSYSINF considerations	151
7.5.2 SAVSYSINF backup strategy	152
7.6 Changes to SAV commands	152
7.6.1 SAVLIB, SAVOBJ, and SAVCHGOBJ	152
7.6.2 Save data queue contents	152
7.6.3 Parallel library save enhancement	153
7.7 Backup and Recovery Media Services	153
7.7.1 BRMS iSeries Navigator GUI client	155
7.7.2 Backup enhancements	159
7.7.3 Recovery and retrieval enhancements	161
7.7.4 Devices/media management enhancements	162
7.7.5 Miscellaneous enhancements	168
Chapter 8. i5/OS-based virtual tape	169
8.1 Virtual tape support	170
8.2 Key advantages	170
8.3 Considerations	170
8.4 Overview and flow	171
8.4.1 Implementation of virtual tape	172
8.4.2 Virtual volumes stored as IFS stream files	172
8.4.3 Image catalog	173
8.4.4 Virtual media	173
8.4.5 Virtual tape devices with guest partitions (Linux, AIX)	174
8.5 Implementing virtual tape	174
8.5.1 Create virtual tape devices	175
8.5.2 Create the image catalog	175
8.5.3 Create a virtual tape image	177
8.5.4 Working with virtual devices	179
8.5.5 Working with image catalogs	180
8.5.6 Working with virtual volumes	181
8.6 Dependent catalog overview	181
8.6.1 Using dependent catalogs	182
8.6.2 Reference catalog	182
8.6.3 Dependent catalog	183
8.6.4 Creating a dependent catalog	184
8.6.5 Shadow catalog views	185

8.6.6	Catalog security changes	186
8.6.7	New image catalog APIs	186
8.7	V5R4 BRMS virtual tape support	187
Chapter 9.	Accessing 5/OS via the Web	189
9.1	iSeries Navigator tasks on the Web	190
9.1.1	Product packaging and software requirements	190
9.1.2	Activation and setup	190
9.1.3	i5/OS V5R4 iSeries Navigator Tasks on the Web - basic functions	196
9.2	iSeries Access for Web enhancements	207
9.2.1	Summary of V5R4 enhancements	208
9.2.2	Servlet version enhancements	209
9.2.3	Portlet version enhancements	215
9.2.4	Extract server data	217
Chapter 10.	i5/OS-based communications and general connectivity	221
10.1	SNA and the Enterprise Extenders	222
10.1.1	SNA function support under EE	223
10.2	IPv6 support	225
10.2.1	More IPv6 summary information	226
10.3	TCP/IP configuration enhancements	228
10.4	Virtual IP and Virtual Ethernet	229
10.4.1	Virtual IP enhancements	229
10.4.2	Virtual Ethernet and the Multiple Proxy ARP support	231
10.4.3	Proxy ARP support for a Virtual IP address on a different subnet	232
10.5	Virtual Private Networking and the i5/OS V5R4	233
10.5.1	Traffic Flow Confidentiality	234
10.5.2	Extended Sequence Number (ESN)	235
10.6	Point-to-Point Protocol and Layer 2 Tunnel Protocol	236
10.7	HTTP server enhancements	237
10.8	iSeries NetServer changes	237
10.9	iSeries Access family in V5R4	240
10.9.1	iSeries Access for Windows	240
10.9.2	iSeries Access for Linux	243
Chapter 11.	i5/OS-based consoles and logical partitioning	245
11.1	V5R4 console support	246
11.1.1	Operations Console support	246
11.1.2	Hardware Management Console	249
11.1.3	Thin Console	251
11.1.4	Service support level console function changes with V5R4	254
11.2	Enhancements to logical partitioning	255
11.2.1	Advanced Power Virtualization	255
11.2.2	Virtual Partition Manager	256
11.2.3	LPAR simplification update	257
11.2.4	IBM System Planning Tool	258
11.3	General LPAR updates	264
11.3.1	Reserve Capacity on Demand update	265
Chapter 12.	i5/OS-based application development languages and tools	269
12.1	Application developer roadmap	270
12.1.1	JavaServer Pages (JSPs)	272
12.1.2	eXtensible Markup Language	272
12.1.3	Servlet	272

12.1.4	Portlet	273
12.1.5	Web services	273
12.2	IBM System i5 Initiative for Innovation program	274
12.2.1	Application innovation	274
12.2.2	Tools innovation	274
12.2.3	System i5 Innovation	275
12.3	Control Language	275
12.3.1	Subroutine support	275
12.3.2	Variables	276
12.3.3	Built-in functions	277
12.4	ILE RPG programming	278
12.4.1	Control specifications	278
12.4.2	File specifications	278
12.4.3	Definition specifications	279
12.4.4	Operation codes	279
12.4.5	Built-in functions	281
12.4.6	Free format SQL	282
12.4.7	Rules for naming in embedded SQL are updated	282
12.5	ILE COBOL programming	283
12.6	Java	283
12.6.1	IBM technology for Java Virtual Machine	283
12.6.2	What is new for IBM Developer Kit for Java	286
12.6.3	What is new for the IBM Toolbox for Java	288
12.7	Structured Query Language (SQL)	290
12.7.1	ISO timestamp format	290
12.7.2	Special register	290
12.7.3	Recursive common table expressions and recursive views	291
12.7.4	OLAP expressions	297
12.7.5	SELECT enhancements	299
12.7.6	Richer toolbox of SQL functions	301
12.7.7	Instead of Triggers (IOTs)	302
12.7.8	Miscellaneous enhancements	302
12.8	XML Toolkit	303
12.9	WebSphere Development Studio and Client	304
12.9.1	WDSC V6.0.1	306
12.10	WebFacing and HATS	307
Chapter 13.	i5/OS-based WebSphere application support	313
13.1	WebSphere Application Server V6.0	314
13.1.1	WebSphere Application Server programming model	314
13.1.2	WebSphere Application Server V6.0 system management features	318
13.1.3	WebSphere Application Server V6.0 security	322
13.1.4	WebSphere Application Server V6.0 supplied applications	322
13.1.5	WebSphere Application Server V6.0 on i5/OS packaging	324
13.2	WebSphere Application Server Version 6.1 on i5/OS	327
13.2.1	Multiple WebSphere Application Server versions coexistence	328
13.2.2	JDK 5.0 support	328
13.2.3	WebSphere Application Server installation changes	328
13.2.4	WebSphere Application Server administration	329
13.2.5	WebSphere Application Server security enhancements	330
13.2.6	Runtime i5/OS subsystem	331
13.2.7	Application development enhancements	331
13.3	IBM Express Runtime	334

13.3.1	Middleware included in Express Runtime V2.1.1	334
13.3.2	Express Runtime V2.1.1 development and deployment tools.	335
Chapter 14.	i5/OS-based Web enablement enhancements	337
14.1	iSeries Access for Web, iSeries Navigator tasks on Web, and new integrated Web application server	338
14.2	IBM Express Runtime Web Environments for i5/OS	339
14.2.1	Hardware and software requirements for 5733-SO1	342
14.2.2	Installing the Web Enablement Environment	343
14.2.3	Starting to use the Web Enablement Environment	344
14.2.4	Using IBM Express Runtime Web Environments for i5/OS with applications	350
14.3	Web performance management tools	350
14.3.1	Web Performance Monitor	352
14.3.2	Web Performance Advisor	358
14.4	WebSphere Application Server performance data	370
14.5	Simplified use of SSL with i5/OS HTTP Administration server	372
Chapter 15.	i5/OS printing and printer output enhancements	377
15.1	Introducing printing and printer output enhancements	378
15.2	i5/OS printing and printer output management	378
15.3	Native save/restore spooled files	379
15.3.1	History and changes	379
15.3.2	Saving and restoring spool file examples	381
15.3.3	OUTQ and spooled files in iASP considerations	383
15.3.4	Restrictions and limitations	384
15.4	Spool file, policy-based output queue management	386
15.4.1	V5R4 introduces spooled file expiration	386
15.4.2	Action plan: gaining control over your output queue	387
15.5	Changes Work With Spooled Files (WRKSPLF)	388
15.6	Print output formatting and routing	389
15.6.1	Infoprint Server for iSeries	389
15.6.2	Infoprint Designer for iSeries	390
15.6.3	Mapping Suite	390
15.6.4	Host Print Transform	391
15.6.5	DDS printer file enhancements	395
15.6.6	CHGIPLA modification aids spooled file retention	398
15.7	Printing hardware update (selected devices)	398
15.7.1	RFID: Infoprint 6700	398
15.7.2	MFP devices	399
15.7.3	Infoprint 15xx and 16xx workgroup printers	400
Chapter 16.	i5/OS-based performance update	403
16.1	i5/OS V5R4 performance updates	404
16.2	CPW and MCU ratings	404
16.2.1	Database performance update	408
16.2.2	Journaling	410
16.2.3	Reorganize table (physical file member)	411
16.2.4	Disk hardware performance - smaller cache disk controllers	417
16.2.5	Fastest disk controllers (adapters) performance	419
16.2.6	Virtual I/O restructuring	421
16.2.7	Improved NetServer performance	421
16.2.8	Java and WebSphere performance	421
16.2.9	Object conversion	422
16.2.10	LPAR considerations	423

16.2.11 Performance updated on integrated xSeries servers	424
16.3 Performance tools updates	427
16.3.1 iSeries Navigator Monitor	427
16.3.2 Performance Tools for i5/OS 5722-PT1	428
16.3.3 iSeries Navigator performance tools plug-in	430
16.4 Changes to job trace and job analysis reports	430
16.4.1 The Trace Analysis Summary report (QPPTTRC1)	431
16.4.2 The Trace Analysis I/O Summary report (QPPTTRC2)	432
16.5 Changes in the System report.	432
16.5.1 Changes to the disk utilization section of the System report.	433
16.5.2 Changes to the workload section of the system report.	434
16.5.3 HTTP Server summary	435
16.6 Changes in the Component Report.	435
16.6.1 HTTP Server activity	436
16.6.2 Domino server statistics	437
16.7 WRKSYSACT	437
16.8 Display performance data	438
16.9 Job Watcher	439
16.10 IBM Systems Workload Estimator.	441
16.11 IBM Disk Magic	444
16.12 Performance Management for System i5	444
16.12.1 PMiSeriesIS tool	449
16.12.2 Integrating PM for System i5 and the IBM System Workload Estimator.	451
16.13 Web Performance tools.	451
16.14 Non-IBM sizing tools	451
16.14.1 Midrange Performance Group (MPG) Performance Navigator	451
16.14.2 BMC PATROL for iSeries - Predict.	452
Chapter 17. i5/OS and IBM Director.	455
17.1 IBM Director 5.10	456
17.2 IBM Director overview.	456
17.2.1 IBM Director topology	457
17.3 Reducing complexity, delivering value	460
17.4 IBM Director V5.10 Console	461
17.4.1 Starting and stopping	463
17.4.2 Discovery of managed objects	464
17.5 Controlling your system.	466
17.5.1 Event action plans.	466
17.5.2 Job Scheduler	468
17.5.3 Process Management.	470
17.5.4 Resource monitors	471
17.6 IBM Director 5.10 extensions	471
17.6.1 IBM Director 5.10 base extensions	472
17.6.2 Free download extensions	472
17.6.3 Fee-based extensions.	474
17.6.4 Software Distribution (Premium Edition).	474
17.7 IBM Director for i5/OS Version 5.20	479
17.8 IBM Director and Tivoli	481
17.9 References	481
Chapter 18. i5/OS and Lotus products	483
18.1 Lotus products.	484
18.2 Lotus Domino for i5/OS.	484

18.2.1 Domino releases on i5/OS	484
18.2.2 End of marketing	486
18.3 Lotus Workplace on i5/OS	486
18.3.1 Workplace Services Express (WSE)	486
18.3.2 Workplace Collaboration Services (WCS)	487
18.3.3 IBM Workplace Designer	489
18.4 References	490
Appendix A. Hardware Management Console microcode and system firmware compatibility levels	491
HMC machine code upgrade	492
Firmware history: releases and service pack levels	493
2007 Daylight saving time changes support	493
Viewing HMC firmware and system firmware levels	494
Appendix B. IBM System i5 memory details	499
System i5 model DDR1, DDR2 memory and plugging rules	500
System i5 520 and 550 model memory technology considerations	500
System i5 570 model memory technology considerations	501
System i5 595 model memory technology considerations	504
Appendix C. System i5 Capacity Backup edition software-based details	509
System i CBU edition software details	510
Appendix D. IBM Support Assistant	515
The IBM Support Assistant	516
Installation and setup	517
Search	519
Product information	521
Tools	523
Support	525
Help	526
IBM Education Assistant	527
Appendix E. System i IP Telephony powered by 3Com	529
System i IP Telephony, Voice over IP, and 3Com	530
3Com offerings on System i configurations	532
3Com IP Telephony Suite for System i	532
IBM System i5 520 model Telephony Express packages	535
Appendix F. IBM-Zend PHP on i5/OS	541
IBM-Zend PHP Web deployment environment	542
Appendix G. i5/OS V5R4 IPv6 restrictions	543
IPv6 support restrictions through V5R4	544
Appendix H. IBM future planning statements	545
Planning statements	546
Marketing withdrawals - hardware	546
Software withdrawals	547
General software planning statements	547
Appendix I. IBM System iSociety	549
iSociety overview and facilities	550

Related publications	551
IBM Redbooks	551
Online resources	551
How to get IBM Redbooks	553
Help from IBM	553
Index	555

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the products and/or the programs described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.


This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo)  ™	DB2®	POWER™
developerWorks®	DRDA®	POWER Hypervisor™
ibm.com®	DS4000™	POWER4™
iSeries™	DS6000™	POWER5™
i5/OS®	DS8000™	POWER5+™
pSeries®	Electronic Service Agent™	QuickPlace®
xSeries®	FlashCopy®	Rational®
z/OS®	Footprint®	Redbooks™
z/VM®	Infoprint®	Sametime®
zSeries®	Informix®	ServeRAID™
z9™	Integrated Language Environment®	System i™
Advanced Function Presentation™	IBM®	System i5™
AnyNet®	IPDS™	System p™
Asset ID™	Language Environment®	System p5™
AFP™	Lotus Enterprise Integrator®	System x™
AIX 5L™	Lotus Notes®	System z™
AIX®	Lotus®	System z9™
AS/400e™	Lotusphere®	System Storage™
AS/400®	Micro-Partitioning™	System/36™
Bar Code Object Content Architecture™	MQSeries®	System/38™
BladeCenter®	Netfinity®	Tivoli®
BCOCA™	NetServer™	TotalStorage®
Cloudscape™	Notes®	Virtualization Engine™
Distributed Relational Database Architecture™	OpenPower™	VisualAge®
Domino®	OS/400®	WebSphere®
DB2 Connect™	Passport Advantage®	Workplace™
DB2 Universal Database™	Power PC®	Workplace Collaborative Learning™
	PowerPC®	Workplace Managed Client™
	Print Services Facility™	1-2-3®

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates.

mySAP, SAP R/3, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Snapshot, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Enterprise JavaBeans, EJB, Java, Javadoc, JavaBeans, JavaMail, JavaScript, JavaServer, JavaServer Pages, JDBC, JDK, JMX, JRE, JSP, JVM, J2EE, J2SE, Solaris, Streamline, Sun, Sun Microsystems, XGL, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Excel, Expression, IntelliSense, Internet Explorer, JScript, Microsoft, PowerPoint, Visual Basic, Visual C++, Visual Studio, Windows Server, Windows Vista, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Itanium, Pentium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

BMC Patrol ©

Midrange Performance Group ® (MPG) Performance Navigator ®

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbook provides a technical overview of IBM System i™ family hardware and i5/OS®-based software at Version 5 Release 4 announced January 2006 through February 2007. It is intended for those already familiar with System i and i5/OS who need a broad perspective of what is new with V5R4, either to consider upgrading to V5R4 or to take full advantage of V5R4 capabilities.

Several, but not all, V5R4-based licensed program capabilities are also included. These include Lotus® products, WebSphere® Application Server versions, WebSphere Development Studio for iSeries™ capabilities, and IBM Director for i5/OS capabilities.

Selected i5/OS V5R3 level content is included where the developers of this book determined the function should be included. Support introduced with V5R3 is clearly identified.

Information in this book is generally of an intermediate level of detail. Providing this information in one place, this is an ideal document by which to obtain a good understanding of System i and i5/OS V5R4 capabilities. This book cannot provide all of the hardware and software details contained in documentation located at various Web sites. For example, it does not contain all processor models and I/O hardware supported details or how to attach cables to the system or I/O enclosures.

It does contain System i hardware and software information on capabilities announced January 2006 through February 2007.

Throughout this book we refer you to Web sites for more information. Information about 2007 hardware and software should also be available at these Web sites:

- ▶ The IBM Systems - Servers Web site
<http://www.ibm.com/servers/index.html>
- ▶ The IBM System i Web site
<http://www.ibm.com/systems/i/>
- ▶ The IBM Systems Information Center Web site
<http://publib.boulder.ibm.com/eserver>

This site includes the following links as of December 2006:

- IBM Hardware (POWER5™ or later technologies) Information Center. This includes information about IBM System i5™ and IBM System p5™ models, capacity on demand, and other hardware such as newer disk controllers and cabling information:
<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>
- The IBM iSeries (System i) Information Center Web site:
<http://publib.boulder.ibm.com/series/>
- IBM eServer™ pSeries® (System p™) and AIX® Information Center:
<http://publib16.boulder.ibm.com/pseries/index.htm>
- IBM Systems Software Information Center. This site includes information about products such as IBM Director, that run on multiple operating systems.
<http://publib.boulder.ibm.com/eserver/swic.html>

- ▶ IBM ITSO Redbooks™ Web site (<http://www.ibm.com/redbooks>) for System i and i5/OS-based documents, such as:
 - *High Availability Considerations: SAP R/3 on DB2 for OS*, SG24-2003 (Refer to the Web site for the latest update level.)
 - *IBM System i5 Handbook IBM i5/OS Version 5 Release 4 January 2006*, SG24-7486 (Refer to the Web site for the latest update level).
 - *PCI and PCI-X Placement Rules for IBM System i5, eServer i5, and iSeries servers with i5/OS V5R4 and V5R3*, REDP-40111 (Refer to the Web site for the latest update level.)
 - Additional hardware-based redpapers planned for 2007 announcements
- ▶ IBM System i planning, including upgrades information Web site
<http://www.ibm.com/systems/support/i/planning/migrationupgrade.html>

The team that wrote this redbook

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Rochester Center. The residents created the work during three separate time periods in order to include announced content through August 2006.

The original group of residents created a set of moderately detailed Microsoft® PowerPoint® V5R4 Technical Overview presentations. These presentations are at an March 2006 level of content and can be accessed from the IBM iSeries (System i) technical support Web site in the Technical library category available for viewing at:

<http://www.ibm.com/eserver/support/iseriess>

This group of residents included:

Louis Cuypers is an IBM IT Specialist in Belgium. He has 28 years of experience in Technical Support. He has contributed to several ITSO publications, including ITSO technical overview and iSeries System Handbook and Builder. He has also authored the iSeries LPAR publications.

Joonwoo Kim is an IT Specialist in Korea, working on the IBM System i family of systems since 1997. He has provided technical support to customers in the areas of hardware and multiple software areas including i5/OS performance and work management, availability, database, IBM WebSphere, LPAR, and server consolidation. He currently works for Global Technology Services as a Technical Availability representative.

Jos Vermaere is a Systems Architect and Consulting IT Specialist in IBM Belgium. He joined IBM Belgium in 1979 and has held several positions, including a three-year assignment in the Dallas ITSO (1988–1991). He has contributed to many ITSO publications, including ITSO technical overviews, IBM Redbooks on globalization, and performance-related material.

Will Witten is a Senior Information Technology Specialist with the IBM Advanced Technical Sales Support organization in Rochester, MN. He has had extensive experience with the Rochester, MN based systems (S/3x, AS/400®, and System i technologies). For the last several years Will has assisted with large and complex customer opportunities involving Lotus DOMINO on the iSeries. He also assists with Critical Situations and presents at user group conferences (Lotusphere®, IBM Technical Conference, COMMON, and so forth). Prior to working for IBM, he was the Data Processing Manager for an independent oil and gas exploration company.

The following residents and project leader produced this book by summarizing the technical overview presentation set content and integrating it with new or updated content that became available March 2006 through December 2006:

Jim Cook is a Senior Software Engineer at the ITSO, Rochester Center. He leads teams that produce IBM System i announcement presentation sets that are maintained on the System i technical support Web sites and presents at ITSO iSeries Forums internationally. Jim also produces IBM Redbooks about various System i and i5/OS-related topics.

Donald Otis is a Senior IT Specialist working in Techline, System i support, in the United States. A professional hire, he has over 20 years of retail management, inventory control, warehousing and transportation, and office administration experience. He has worked at IBM for nine years, initially in the Printing Systems Division. His current focus is on System i technology and configuration, printing, virtualization, and open source and collaboration tools.

Roshan Tarrosa is a System i5 Advisory IT Specialist for the IBM Global Technology Services of IBM Philippines. He has over 10 years of iSeries experience, providing technical support to diverse sets of clients across industries. He also works for the IBM ASEAN ITS Regional Support Center supporting customers across the region. As an IT Specialist he specializes in iSeries Access and iSeries connectivity (SNA and TCP/IP), Linux®, and Windows® integration on iSeries; backup and recovery management; and security for iSeries. He has worked at IBM for 15 years.

Hundreds of Rochester, Guadalajara, Austin, Toronto, and Raleigh development, marketing, and technical support people created the products and source material used for this book. At the risk of omitting some of the major contributors, we acknowledge special thanks to the following:

Ian Jarman
Mark Olson
Doug Fulmer
Hoovey Halverson
Kyle Wurgler
Mark Anderson
Kent Milligan
Amartey Pearson
Tracy Smith
Allyn Walsh
Gene Rentz
Steve Dickes
Tim Kramer
Mike Aho
Brian Nordyke
George Romano
Bob Dick
Carole Miner
Doug Beauchene
Deb Saugen
Dave Bhaskaran
Sanjay Patel
Scott Maxson
Bob Gintowt
Dave Novey
Al Smith
Tracy Smith
John Kochan

Craig Wilcoz
Eduardo Aguila
John McGinn
Hector Fierros
Ginny McCright
Guy Vig
Jeff Uehling
Jim Naylor
Charles Graham
Bruce Vining
Kathy Tri
John Kasperski
Larry Youngren
Peg Levering
James Thompson
Randy Grimm
Scott Urness
Jeff Thelen
Dave Legler
Pat Fleming
Paul A Wolf
Kevin Kirkeby
Bill Shaffer
Ken Dittrich
Terry Luebbe
Pat Glenski
Randy Erickson
Josep Cors
Randy Nelson
Jeff Meaden
Brian Podrow
Sue Baker
Don Pischke

The following are members of the Rochester International Support Center:

Hernando Bedoya
Aleksandr Nartovich
Nick Harris
Gary Mullen-Schulz

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review book form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400



Introducing the IBM System i5 hardware and i5/OS V5R4 overview

This chapter highlights IBM System i announcements through February 2007, including hardware and software. It provides a preview of the specific topics in this book.

See “Preface” on page xvii for summary information about where to go for documentation more detailed than the contents of this book. See also the additional resources listed in “Related publications” on page 551.

From a hardware viewpoint, the key documentation to be used with this book are *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155; *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition)*, REDP-4011; and the System i Planning Web site. Details on cable connection and setup on POWER5-based hardware are contained in the IBM Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>

1.1 Overview of System i5 2006 announcements

In January 2006 IBM announced new System i5 processor models and new I/O capabilities, including I/O adapters (IOAs) that do not require attachment to an I/O Processor (IOP), and the new i5/OS V5R4 release.

During the May through February 2007 time period there were additional System i5-based hardware and i5/OS-based software announcements. This chapter provides a brief overview of the content of these announcements. The chapters that follow provide a moderate level of detail on these capabilities.

With these announcements IBM continues on its roadmap committed to System i innovation and leadership within the IBM POWER5 processor technology by refreshing the System i5 product line with the latest advanced POWER5 processors and some enhanced I/O device capabilities. Especially important are the new processors within the popular IBM System i5 520 model system configurations and the integration of IBM's fastest POWER5+™ processors within the System i5 570 models.

The new V5R4 release of i5/OS contains even more integrated capabilities with selected IBM middleware enhancements that promotes innovation, while enhancing the value of your current investments in solutions in the areas of:

- ▶ Integrated middleware
- ▶ Simplified Information Technology (IT) infrastructure
- ▶ Security and compliance
- ▶ Business resiliency

i5/OS V5R4 contains a wide range of medium to small enhancements and new functions. Taken together, this set of capabilities provides significant enhancements to assist you in providing solutions innovation, simplifying your information technology, and supporting your business resiliency and compliance requirements.

There are also enhancements to many System i licensed programs with i5/OS V5R4.

This chapter lists many of the key V5R4-based enhancements covered in the following chapters of this book. Selected V5R3 capabilities that became available during 2005 are also included in some of the topics covered in this technical overview. This includes simplifying deployment of new POWER5+ systems with support for both V5R3 and V5R4.

Topics covered here include:

- ▶ New System i5 models, including the POWER5+ based 520, 550, 570, and 590 models and an update on Solution and Capacity Backup edition offerings
- ▶ New I/O capabilities, including new disk controllers and an expanded Operations Console over the LAN and a Thin Console alternative for the i5/OS partition management console
- ▶ New Java™ and enhancements to CL, RPG, and COBOL
- ▶ Database management and enhanced SQL
- ▶ WebSphere Development Studio, including enhanced WebFacing and HATS integration
- ▶ Enhanced integration with Microsoft Windows applications
- ▶ Print and output management
- ▶ DB2® Content Manager
- ▶ iSeries Access for Web and iSeries Navigator on the Web

- ▶ Communications support including iSeries Access family product enhancements and iP network enhancements, including Enterprise Extenders (SNA over IP)
- ▶ Updates on console and LPAR support
- ▶ Streamlined i5/OS work management
- ▶ Enhanced hardware storage protection, network intrusion detection, i5/OS security
- ▶ Expanded save and restore options and new virtual tape support
- ▶ Journaling and clustering enhancements
- ▶ Addition of RAID6 to RAID5 disk hardware protection
- ▶ Extended Windows operating system integration, including support of System i Small Computer Standardized Interface (iSCSI) attachment to BladeCenter® and xSeries® servers
- ▶ Update on automated detection and cleaning of viruses stored in i5/OS, AIX 5L™, or Linux partition
- ▶ IBM Express Runtime Web Environments for i5/OS (This product was introduced October 2006 and contains a wide range of capabilities intended to get you operational in a Web application serving environment as quickly as possible.)
- ▶ General update on performance management and performance tools under i5/OS
- ▶ Overview of IBM-Zend PHP on i5/OS and 3Com Telephony offerings on System i5 models
- ▶ Overview information about iSociety and IBM Support Assistant

We also cover additional V5R3 and V5R4 topics that we determined should be included in an overview of System i support. These topics include:

- ▶ Unique LPAR topics such as the Advanced POWER™ Virtualization feature normally associated with IBM System p5 and eServer p5 systems, Virtual Partition Manager, and on-going *LPAR simplification*
- ▶ Additional system console hardware choices available with V5R4
- ▶ Overview of WebSphere Application Server V6.0 and V 6.1 on i5/OS
- ▶ Overview of Virtualization Engine™ Console and IBM Director V5.10 and later
- ▶ Overview of iSeries Navigator tasks on the Web
- ▶ Lotus Domino® and Workplace™ products update

In addition to the specific topic chapters, several appendixes supply additional selected topic details. Appendixes include one summarizing IBM's *planning* statements for future enhancements or withdrawal from marketing and support, and one identifies corresponding levels of HMC software and IBM System i5 system firmware.

Attention: In this book we use the term *i5/OS partition* to mean both:

- ▶ A single partition within an IBM System i configuration running i5/OS when more than one partition may be configured.
- ▶ The entire IBM System i configuration running i5/OS when there are no specifically configured partitions; that is, the system is one i5/OS partition that owns all hardware.

For more details on the topics covered in this book, consider accessing the following Web sites and associated documentation. Some of these document resources are specifically mentioned in this book under the appropriate topics:

- ▶ iSeries Information Center (V5R2, V5R3, V5R4) at:
<http://publib.boulder.ibm.com/iseriess>
- ▶ IBM Hardware Information Center (POWER5-based hardware) at:
<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>
- ▶ IBM Software Information Center at (Virtualization Engine, IBM Director, and so on) at:
<http://publib.boulder.ibm.com/eserver/swic.html>
- ▶ ITSO Technical Overview set of presentations (V5R2, V5R3, V5R4) under the Technical Library link at the technical support Web site:
<http://www-03.ibm.com/servers/eserver/support/iseriess/index.html>
- ▶ IBM Redbooks and redpapers Web site at:
<http://www.ibm.com/redbooks>
- ▶ iSeries performance topics at:
<http://www.ibm.com/eserver/iseriess/perfmgmt>



IBM System i5 hardware

Information in this chapter is generally of an intermediate level of detail. We refer you to sources for more detailed information about specific topics, including the Web sites listed in the Preface and “Related publications” on page 551. This chapter discusses the new and enhanced IBM System i5 models, memory, and I/O features announced from January 2006 through February 2007.

Note that this chapter does not contain all supported System i hardware announced or withdrawn through December 2006. It also does not contain hardware information announced during 2007. From a hardware viewpoint you can find more details at the Web sites listed below. Key hardware documentation to be used with this book includes *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155, and the System i Planning Web site. How to access this and additional documentation is included below in the listed Web sites:

- ▶ The IBM Systems Information Center Web site

<http://publib.boulder.ibm.com/eserver>

This site includes a link to the IBM Hardware (POWER5 or later technologies) Information Center. This includes information about IBM System i5 and IBM System p5 models, capacity on demand, and other hardware such as newer disk controllers and cabling information.

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>

- ▶ The IBM iSeries (System i) Information Center Web site

<http://publib.boulder.ibm.com/iserries/>

- ▶ IBM Redbooks Web site for System i and i5/OS-based documents

<http://www.ibm.com/redbooks>

- *High Availability Considerations: SAP R/3 on DB2 for OS*, SG24-20032155 (Refer to the Web site for the latest update level.)
- *IBM System i5 Handbook IBM i5/OS Version 5 Release 4 January 2006*, SG24-7486 (Refer to the Web site for the latest update level.)
- *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition)*, REDP-4011 (Refer to the Web site for the latest update level.)

- Additional hardware-based redpapers planned for 2007 announcements
- ▶ IBM System i planning, including upgrades information Web site
<http://www.ibm.com/systems/support/i/planning/migrationupgrade.html>

2.1 System i5 models: more power and flexibility

The IBM system i5 family continues to grow in processing power and added flexibility with special *edition* or *solution* configurations. The 520, 550, and 570 have all been enhanced with POWER5+ processors, and the 595 is enhanced with a 1.9 GHz POWER5 processor. With the introduction of processor capacity on demand (permanent CoD, On/Off, Trial and Reserve) in the 1/2-way 520, and the IBM Accelerator for System i5 with Value and Express editions, capacity on demand is now available on each member of the System i5 server family.

Additional flexibility has been added to the Capacity Backup edition offerings and there are more applications available as special Solution edition offering and packages.

The POWER5-based systems also get faster memory cards. Some of the faster memory cards also have a corresponding size card that is lower priced and slightly slower than its higher speed counterpart.

The System i5 server family structure is simplified in several areas. For example, the 570 starts now at a 2-way to 4-way, and an 8-way to 16-way is introduced simplifying the available offerings while maintaining flexibility to host different workloads and adapting to different operating environments.

We also discuss the new options for using several existing and new adapters without a requirement for a Input Output Processor controlling these adapters, including a new Internet Small Computer System Interface (iSCSI) adapter enabling expanded integration with Windows operating systems running on IBM BladeCenter and a wider range of IBM xSeries systems.

Figure 2-1 shows the System i5 family of servers.



Figure 2-1 System i5 family of servers

2.2 System i5 hardware changes

This section is devoted to the hardware changes in the System i5 family of processor models. Each model of the i5 systems is available in several editions.

2.2.1 System i5 520 hardware changes

The new System i5 i520 now comes in a 1-way or 1-way to 2-way processors. The models announced during January 2006 come with a 1.9 GHz POWER 5 processor, increasing the processing power rating up to 7100 CPW. The new 520 model has a maximum memory of 32 GB of DDR2 memory and a maximum capacity of 39 TB of disk space.

The model 520 can run on an i5/OS V5R3 with a V5R3M5 Licensed Internal Code (LIC) or later with a minimum required cumulative PTF level of C6045530.

This new 520 model processor enclosure comes with the following:

- ▶ 1.9 GHz POWER5+ processor
- ▶ Eight DDR2 memory slots
- ▶ Five PCI-X slots
- ▶ One PCI-X DDR slot (only for dual mode IOAs (smart IOAs), which can run with or without an IOP)
- ▶ Two imbedded 10/100/1000 Mb Ethernet LANs
- ▶ Embedded I/O controller (IOP-less capable)
- ▶ Optional write cache/RAID-5 (IOP-less capable)
- ▶ No Base IOP
- ▶ Base WAN IOA for ECS (uses PCI-X slot)
- ▶ Two HMC ports
- ▶ Optional HSL-2 loop
- ▶ Eight disk slots (four optionally activated)
- ▶ Two DVD drive slots
- ▶ One tape drive slot

In contrast to 520 models available before the new January 2006 520 models, these new models come with no base IOP and no base HSL loop. These are optional on the new models. Like the 2004/2005 model 520, every 520 has eight disk slots or bays. With the base order, four of the slots are enabled. Use of the other four slots requires a chargeable enablement feature.

In addition, there is one tape slot for Quarter Inch Cartridge (QIC), Linear Technology Option (LTO), or VXA tape drives. Also available are two slim line slots for DVD drives, of which OS/400® can drive only one slot. Both slim line slots can be driven only by AIX 5L or Linux if these operating systems are running in a partition that owns the processor enclosure (also referred to as the CEC or system unit).

Note: Though the CEC DVD-RAM device can be an alternate IPL and save/restore device, it is not generally recommended as a save/restore device for performance reasons. The device's speed is very slow compared to physical tape or i5/OS support of virtual tape.

Also, when operating within the CEC, the DVD-RAM competes for SCSI bus resources with the first four disk drives in the CEC (system unit) and can significantly slow their performance.

If the DVD-RAM device within the CEC is to be used for save/restore, this should be done only during times when other workload does not require access to the first set of four CEC disk drives.

The DVD-RAM device competition with SCSI bus resources does not apply to:

- ▶ Any second set of four disk drives in the CEC
- ▶ When the DVD-RAM device is placed outside the CEC in a:
 - 5094/5294 I/O unit
 - Separate bridge box

Figure 2-2 shows System i5 520 desk-side and rack-mounted processor enclosures.



Figure 2-2 System i5 520 desk-side and rack-mounted processor enclosures

This System i5 model is available with the following configuration editions:

- ▶ Express configurations
- ▶ Value Edition
- ▶ Standard Edition
- ▶ Enterprise Edition
- ▶ High Availability (HA) Edition
- ▶ Solution Editions
- ▶ Collaboration Edition
- ▶ Capacity Back Up (CBU) Editions
- ▶ Telephony Express and Telephony Express High Availability (HA) packages

IBM Accelerator feature is available in the System i5 520 1.9 GHz processor models. This feature boosts the processor CPW as well the maximum number of partitions. The chart in Figure 2-3 shows the increase in CPW and the number of partitions for the Express and Value editions.

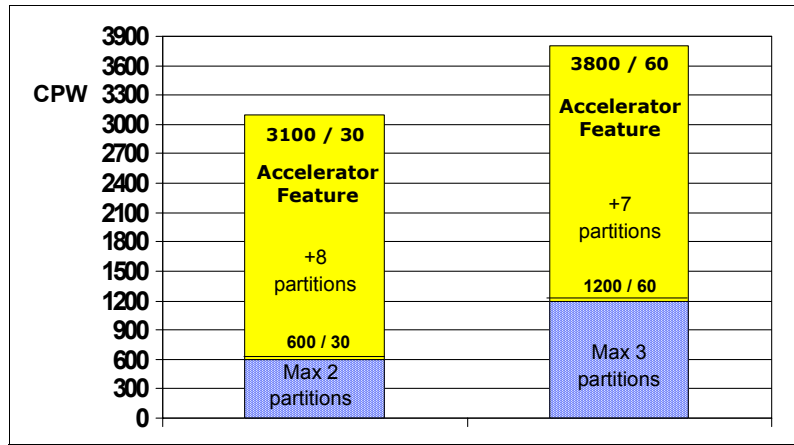


Figure 2-3 Improvement with the IBM Accelerator feature

The 600 to 3100 CPW model is in the P05 software pricing tier and the 1220 to 3800 CPW is in the P10 software pricing tier. The Accelerator feature is enabled either with a managing HMC or, if no HMC is available, through the Advanced System Management Interface (ASMI). For how to do this activation, see the IBM POWER5-based hardware information center Web site:

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>

Note that while activation of the IBM Accelerator feature enables a one-step increase to higher processor power, the lower CPW rating for 5250 OLTP applications remains at the lower, original value.

2.2.2 System i5 550 hardware changes

The IBM 550 model announced during January 2006 comes in 1-way to 4-way with 1.9 GHz processors capable of producing processing power from 3800 CPW up to 14000 CPW. This model has a maximum of 64 GB of DDR2 memory and 77 TB of disk space.

The new 550 model comes with the new faster speed memory cards can run on an i5/OS V5R3 with a V5R3M0 Licensed Internal Code (LIC) or later with a minimum required cumulative PTF level of C6045530.

The System i5 model 550 processor enclosure comes with the following:

- ▶ 1.9GHz POWER5+ processors
- ▶ Sixteen DDR2 memory slots
- ▶ Five PCI-X slots (four slots if optional HSL-2 adapter installed)
- ▶ Two imbedded 10/100/1000 Mb Ethernet LANs

This is either the #5706 or the #5707 2 port LAN adapter, depending on the order. With i5/OS V5R4 one of these may be used for Operations Console over the LAN.

- ▶ Embedded I/O controller (IOP-less capable)
- ▶ Optional (if mirroring) write cache/RAID-5 (IOP-less capable)

- ▶ No Base IOP
- ▶ Base WAN IOA for ECS (uses PCI-X slot)
- ▶ Two HMC ports
- ▶ One base and one optional HSL-2 loop
- ▶ Eight disk slots (four optionally activated)
- ▶ Two DVD drive slots
- ▶ One tape drive slot

Like the newly announced 520 model, the new 550 model comes with no base IOP. The 550 model 550 processor enclosure has eight disk slots or bays. With the base order, four of the slots are enabled. Use of the other four slots requires a chargeable enablement feature.

Like the i520, the i550 has one tape slot for QIC, LTO, or VXA tape drives and two slim line slots for DVD drives, of which OS/400 can drive only one slot. Both slim line slots can be driven only by AIX 5L or Linux if these operating systems are running in a partition that owns the processor enclosure (also referred to as the CEC or system unit).

Note: Though the CEC DVD-RAM device can be an alternate IPL and save/restore device, it is not generally recommended as a save/restore device for performance reasons. The device's speed is very slow compared to physical tape or i5/OS support of virtual tape.

Also, when operating within the CEC, the DVD-RAM competes for SCSI bus resources with the first four disk drives in the CEC (system unit) and can significantly slow their performance.

If the DVD-RAM device within the CEC is to be used for save/restore, this should be done only during times when other workload does not require access to the first set of four CEC disk drives.

The DVD-RAM device competition with SCSI bus resources does not apply to:

- ▶ Any second set of four disk drives in the CEC
- ▶ When the DVD-RAM device is placed outside the CEC in a:
 - 5094/5294 I/O unit
 - Separate bridge box

Figure 2-4 shows System i5 550 desk-side and rack-mounted processor enclosures.



Figure 2-4 System i5 550 desk-side and rack-mounted processor enclosures

Note: The easy access front cover, not the more commonly seen classic front door, is shown.

This System i5 model 550 is available in the following editions:

- ▶ Standard Edition
- ▶ Enterprise Edition
- ▶ HA Edition
- ▶ Solution Editions
- ▶ Domino Edition

2.2.3 System i5 570 hardware changes

The System i5 570 models announced in January 2006 come with a 2.2 GHz processor capable of producing processing power from 8400 CPW and up to 58500 CPW. This model comes in a 2-to-4-way, 4-to-8-way, and 8 -to-16-way processors. A 2-to-16-way processor is available for the Capacity Back Up (CBU) edition. In addition, the 570 has a maximum of 512 GB DDR2 memory and up to 193 TB of disk space.

This model can run on a i5/OS V5R3 with V5R3M0 LIC or later with a minimum of C6045530 cumulative PTF package installed.

Each System i5 model i570 processor enclosure (also referred to as the CEC or system unit) comes with:

- ▶ Two 2.2GHz POWER5+ processor cards
- ▶ Sixteen DDR2 memory slots
- ▶ Six PCI-X slots (five available when optional HSL-2 adapter installed)
- ▶ Embedded I/O controller and optional RAID adapter (with or without IOP)

- ▶ There must be one DVD-ROM or DVD-RAM per system
- ▶ One SP/FSP (Service Processor, frequently referred to also as a Flexible Service Processor) card with two HMC ports and two SPCN ports
- ▶ Six disk slots
- ▶ Two imbedded 1Gbps Ethernet ports
 - This is either the #5706 or #5707 2-port LAN adapter, depending on the order. With i5/OS V5R4 one of these may be used for Operations Console over the LAN.
- ▶ Two USBs and two serial ports (non-i5/OS use, AIX, or Linux use only)
- ▶ Two power supplies
- ▶ Two blowers for I/O cooling
 - Air flow is from front to rear.
- ▶ Two or three CPU regulators
- ▶ One system connector

Note: Though the CEC DVD-RAM device can be an alternate IPL and save/restore device, it is not generally recommended as a save/restore device for performance reasons. The device's speed is very slow compared to physical tape or i5/OS support of virtual tape.

Also, when operating within the CEC, the DVD-RAM competes for SCSI bus resources with up to six disk drives in the CEC (system unit) and can significantly slow their performance.

If the DVD-RAM device within the CEC is to be used for save/restore, this should be done only during times when other workload does not require access to the up to six disks in the CEC.

The DVD-RAM device competition with SCSI bus resources does not apply to when the DVD-RAM device is placed outside the CEC in a:

- ▶ 5094/5294 I/O unit
- ▶ Separate bridge box

Optional features can also be added to the base i570 configuration:

- ▶ Redundant Service Processor (two or more enclosures)
- ▶ Second HSL-2 adapter

The i570 comes in these available editions:

- ▶ Standard Edition
- ▶ Enterprise Edition
- ▶ HA Edition
- ▶ CBU Edition

Figure 2-5 shows a System i5 570 rack-mounted processor enclosure.



Figure 2-5 System i5 570 rack-mounted processor enclosure

2.2.4 System i5 595 hardware changes

The System i5 595 is a high-end enterprise server for large customer environments. With its new 1.9 GHz processors it offers processing power of 26700 to 184000 CPW. This model comes in 8-to-16-way, 16-to-32-way, and 32-to-64-way processors. The 4-to-32-way processor is available for the CBU edition.

The new 595 offers a huge CPW rating, 4.4 times larger than the previous model 890. Large memory and disk capacities are provided to support this CPW capacity along with either a Standard or Enterprise Edition, as well as High Availability (HA) and CBU editions. This model has a maximum memory of 2 TB and a disk storage space capacity of up to 381 TB.

As of August 2006, both the 1.65 GHz and 1.9 GHz 595 models support the higher speed DDR2 memory cards.

Figure 2-6 shows the System i5 - 595 model.



Figure 2-6 System i5 595 model

In addition, this model can support up to 31 HSL-2 loops offering increased I/O capability by linking a maximum of 96 I/O towers. This feature is also available in its 1.65 GHz model.

The i595 is available in the following editions:

- ▶ Standard Edition
- ▶ Enterprise Edition
- ▶ HA Edition
- ▶ CBU Edition

The System i5 - 570 requires a minimum operating system of i5/OS V5R3 with V5R3M0 LIC applied with cumulative PTF level C6045530 or later.

2.3 System i5 memory features

The new POWER5+ processors require faster memory — DDR2. The addition of the new Double Data Rate Dual Channel (DDR2) memory (533 MHz) to the new 520, 550, 570, and 595 models in combination with the CPU, distributed switch, and memory bus speeds helps contribute to the overall system performance.

DDR2 memory is required for the 1.9 or 2.2 GHz 520, 550, and 570 models announced in 2006. This DDR2 memory can be used on either the 1.65 GHz 5 or the 1.9 GHz 595 models. Note, however, that you cannot mix DDR1 and DDR2 memory cards anywhere within a 595 configuration.

The DDR1 memory cannot be used on the model 520, 550, and 570 with 1.9 or 2.2 GHz processors.

Figure 2-7 shows a DDR2 memory module.




Figure 2-7 DDR2 memory module

The System i5 models 520 and 550 with 1.9 GHz processors have PAIR plugging rules, allowing for more flexible and gradual memory configurations. This means that for each memory feature these models require two DIMMs. The model 570, however, will keep its QUAD plugging rules. For the model 595 there are 16 memory card slots per processor book. Memory in the 595 should be plugged in identical pairs with a minimum of four memory cards per processor book.

On the model i595, 1.65 and 1.9 GHz processors DDR1 memory features can still be used. However, even if DDR1 and DDR2 are supported in the 1.9 GHz i595, both cannot be used together on the same machine, and conversions from DDR1 to DDR2 types of memory are not available.

For more details see Appendix B, “IBM System i5 memory details” on page 499.

2.4 IBM System i5 maximums summary

Table 2-1 summarizes the System i5 maximums for processors, memory, and I/O attachments.

Table 2-1 System i5 hardware processor, memory, and I/O attachment maximums summary

Type # Model #	9405 520 1.9 GHz	9406								
		520 1.9 GHz		550 1.9 GHz	570 2.2 GHz			595 1.9 GHz		
N-Way	1	1	1/2	1/4	2/4	4/8	9/16	8/16	16/32	32/64
CPW range	1200 3800 / 7100	1200 3800 / 7100	1200 3800 / 7100	3800 / 14000	8400 / 16000	16700 / 31100	31000 / 58500	26700 / 505000	51000 / 92000	92000 / 184000
MCU range	2600 8200 / 15600	2600 8200 / 15600	2600 8200 / 15600	8200 / 30000	18200 / 34500	35500 / 67500	67500 / 130000	60500 / 114000	115000 / 213000	213000 / 405000
Main Store GB	32	32	32	64	128	256	512	512	1024	2048
HSL Loop	1	1	1	2	2	4	8	7	15	31
Feature I/O Towers	6	6	6	12	12	24	48	36	72	98
Disk Arms	278	278	278	548	546	822	1374	1620	2700	2700
Disk TB	39	39	39	77	77	116	193	228	228	381
PCI slots	90	90	90	172	173	353	692	504	1008	1152
Int Tape, CD, DVD	14/14	14/14	14/14	18/26	18/25	26/36	26/48	26/60	26/60	26/60
External Tape	18/36	18/36	18/36	18/36	18/36	26/48	26/48	26/60	26/60	26/60
External Optical, CD, DVD	18/36	18/36	18/36	18/36	26/48	26/48	26/48	26/60	26/60	26/60
Twinax	1920	1920	1920	5320	5360	7200	7200	7200	7200	7200
Comm lines	192	192	192	320	320	480	480	600	600	600
LAN lines	36	36	36	96	96	128	228	160	160	160
IXS Internal	18	18	18	36	36	48	48	60	60	60
IXA Ext.	8	8	8	16	16	32	57	48	57	57
Crypto CoProc (2)	8	8	8	8	8	8	8/32	8/32	8/32	8/32

Type # Model #	9405 520 1.9 GHz	9406								
		520 1.9 GHz		550 1.9 GHz	570 2.2 GHz			595 1.9 GHz		
Crypto Accelr (2)	2	2	2	4	4/8	4/8	4/8	4/16	4/16	4/16

2.5 System i5 models: the IOP-less approach and strategy

Originally, IOPs were built to off load the CPU of *device, including communications protocol, driver* type functions when the technology of the attached I/O controllers (I/O Adapters, or IOAs) was very basic and of limited processor and storage capacity.

As technology has evolved, the IOAs have progressed to have significant processor speeds and internal storage capacities. Therefore, on a *device type* basis, much of the functions that IOPs currently perform can now be performed by the newer and future technology IOAs. The goal is to eliminate, over time, the need for IOPs to manage I/O adapters.

Figure 2-8 shows an IOP-less approach diagram.

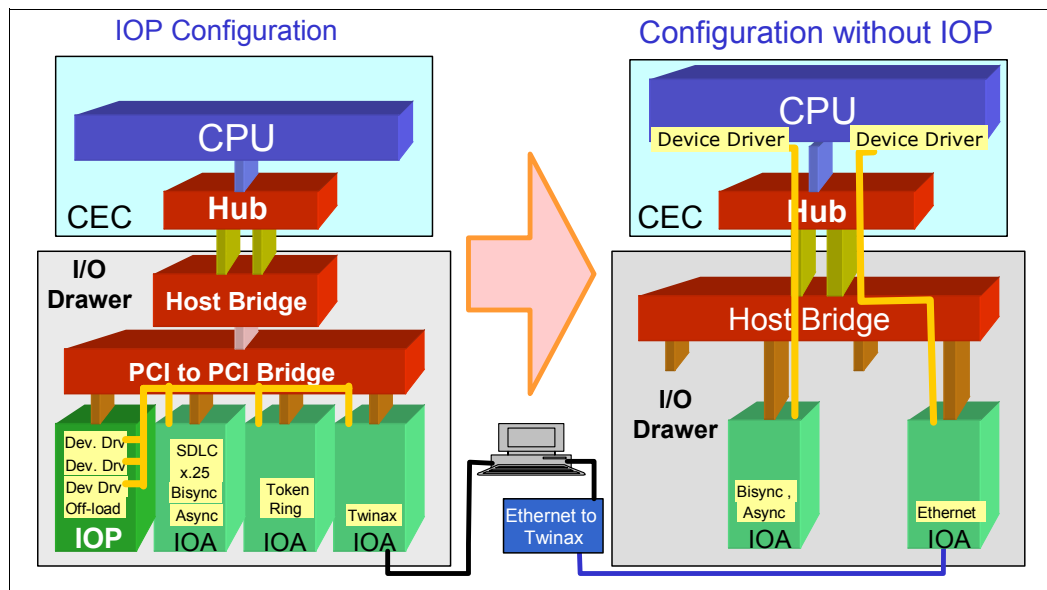


Figure 2-8 IOP-less approach diagram

Implementation of this IOP-less strategy involves:

- ▶ Move device drivers from the IOP to the host (SLIC).
- ▶ I/O Adapters (IOAs): Controllers (IOAs) perform many of the functions previously done by IOPs.
- ▶ Phase out non-strategic IOAs, devices, and protocols.
- ▶ Replace some IOA functions with non-IOA solutions.

The use of Enterprise Extenders is an example of moving IOP functions to host SLIC and i5/OS (non-IOA solutions) implementations.

In the past some SNA and SDLC protocol processing was implemented within a communications capable IOP. Over the years many iSeries customers have applications that have been written to use Systems Network Architecture (SNA) or its services such as SNA Distribution Services and High Performance Routing. SNA function remains quite robust, and years ago had more function relative to the then capabilities using Ethernet and TCP/IP protocols. As a result, customers relied on SNA support. Those customers have implicitly relied on the IOP (and corresponding network switches and routers) to understand and perform SNA-based protocol.

As the IOPs with this implementation are phased out over time and newer technology switches and routers no longer understand SNA protocol, functions like AnyNet® were originally implemented to handle the SNA protocol requirements over an IP-based network. AnyNet works but restricts some of the more advanced SNA functions. So, several years later IBM implemented Enterprise Extenders support that more fully supports most SNA capabilities, compared to AnyNet support. With V5R4, i5/OS includes Enterprise Extender support.

With simple i5/OS configuration parameters SNA applications can have their SNA protocol data streams *encapsulated* in IP frames. Thus, there is no dependency on the older technology IOPs and switches and routers to understand SNA. Only the initiating system and the target system must have Enterprise Extenders and SNA support implemented.

Customers using SNA and moving to or using IP networks should use Enterprise Extenders as an alternative to AnyNet services to achieve optimum performance and obtain broadest SNA application and protocol support.

Note that while there is movement away from relying on a hardware IOP, the new IOA-only technology for a particular I/O capability becomes available only when it has demonstrated sufficient performance to replace the current *IOP dependent* technology.

Adapter types

With the introduction of the *dual mode I/O* adapters (IOAs) for i5/OS in February 2007, the System i models now have three types of i5/OS adapters:

- ▶ IOP controlled IOAs
- ▶ IOP less only IOAs
- ▶ Dual mode IOAs (sometimes also referred to as *smart IOAs*)

Dual mode capability is based on the i5/OS level. To run a dual mode IOA in IOP less mode requires i5/OS V5R4. Dual mode adapters are capable of running with or without an IOP. If the dual mode IOA is placed in a slot that is under IOP control, then the IOA runs as IOP-based. If the slot is not controlled by an IOP, then the adapter runs in IOP less mode.

IOP control of a dual mode adapter always takes precedence over IOP less, except in specific Load Source situations. See the redpaper *PCI and PCI-X Placement Rules for IBM System i models with i5/OS V5R3 and V5R4*, REDP-4011-03, for Load Source and other card placement rules.

Note that an IOP in one partition cannot control an IOA in a different partition. To make it easier for users, dual mode IOAs are assigned separate feature codes based on whether the IOA is to be used as IOP controlled or as IOP less. It is the same physical adapter. The IBM marketing configurator uses the client feature selection to determine whether an IOP is needed with the IOA purchased.

Note: If you want to use a dual mode IOA and run it IOP less, i5/OS must be at V5R4 or later.

For the IOP less only 2 line Ethernet IOAs, such as the embedded Ethernet IOA in the 520, 550, and 570 processor enclosures and the #5706/#5707 PCI-X 1 Gbps Ethernet IOA, the adapter can be assigned to a partition. This means that both lines on that adapter are owned by that partition. You cannot assign one of the lines to another partition.

Also, there are no plans to enhance the twinax technology to an IOA-only feature. The 5250 display field manipulation capabilities are so elegant that producing a corresponding IOP replacement is developmentally cost prohibitive. This means that in order for the System i5 configuration to have twinax support, an IOP for the twinax feature continues to be required.

Given that twinax still remains a primary requirement for a number of installations, they should consider eventually purchasing an external box that converts twinax to Ethernet if they move to a configuration with no IOPs. Ethernet technology on i5 systems will continually be enhanced.

For an update on console options for i5/OS partitions, refer to 11.1, “V5R4 console support” on page 246.

Be aware that the Thin Console (machine type 9944) hardware and support became available in July 2006.

Note: The system hardware, including the twinax hardware, became ROHS-compliant during 2006. The Restriction of Hazardous Substances (RoHS) Directive is a set of criteria formulated by the European Union (EU) to regulate the use of toxic materials in electrical and electronic devices, systems, and toys. The Directive, also known as 2002/95/EC, became effective July 1, 2006.

The RoHS Directive applies to six specific substances:

- ▶ Lead
- ▶ Mercury
- ▶ Cadmium
- ▶ Hexavalent chromium
- ▶ Polybrominated Biphenyls (PBBs)
- ▶ Polybrominated Diphenyl Ethers (PBDEs)

These substances are very common in computer-related hardware. The RoHS Directive does not apply to batteries, tools, high-melting-point solders, the glass used in cathode ray tubes (CRTs) and fluorescent tubes, mercury-vapor light bulbs, ceramic components, and certain alloys for specialized applications. While 2002/95/EC applies specifically to the nations in the EU, similar measures have been proposed or adopted in several other countries.

For more information use any of the commonly accessible search engines and search with *RoHS*.

LPAR dynamic reconfiguration is defined as any action administered by a customer or support person that changes an OS/400 partition's IOP/IOA configuration while the partition is active. This includes concurrent maintenance and dynamic LPAR operations. For these operations the following rules apply:

- ▶ If a dual mode IOA is configured into a slot where an IOP capable of owning the IOA slot is present, the IOA will run IOP-based.
- ▶ If a dual mode IOA is configured into a slot where an IOP capable of owning the IOA slot is not present, the IOA will run IOP-less.
- ▶ If an IOP is configured into a slot that is capable of owning dual mode IOAs the dual mode IOAs must run IOP based. If the dual mode IOAs are currently active as IOP-less IOAs they must be transitioned to IOP-based before the IOP is allowed to power on.
- ▶ If an IOP is configured out of a slot that owned dual mode IOAs (this requires powering off the IOP and owned IOAs), and there is not another IOP capable of owning the IOAs, the IOAs will transition to IOP-less when they power on.

Concerning IOP-less implementation, there are four possible boot scenarios that can occur during PHYP IPL of an i5/OS partition:

- ▶ Pre-Search IPL (Directed/MSD/Last Used Devices) — Prior to searching for a Load Source, PHYP attempts to IPL from either the last used device from a prior IPL, or a particular device specified by SLIC. Instead of applying any upstream IOP logic, PHYP will instead continue to boot in the same mode (IOP-based or IOP-less) as the last time. In particular, SLIC writes information to PHYP that points to a slot that contains either an IOP or an IOA. IOP-based will be done if and only if the slot being pointed to contains an IOP. IOP-less boot will be done if and only if the slot being pointed to contains an IOA. Failing to IPL in the pre-search phase will cause the PHYP to attempt one of the three search cases listed below.
- ▶ Search IPL/IOP — If the LPAR Configuration for the partition points to a slot containing an IOP, then the PHYP will try all possible downstream adapters (Note: This includes IOAs that are not Load Source capable and IOP-less only IOAs) as IOP-based only. IOP-less boot will not be attempted.
- ▶ Search IPL/IOA/No Upstream IOP — If the LPAR Configuration for the partition points to a slot containing an IOA, and there is no upstream IOP, then the PHYP will only try to boot IOP-less from the specified adapter. IOP-based boot will not be attempted.
- ▶ Search IPL/IOA/Upstream IOP — If the LPAR Configuration for the partition points to a slot containing an IOA, and there is an upstream IOP, then the PHYP will try to boot all Load Source candidate devices under the adapter in both modes. First, IOP-based boot will be attempted. If that fails, then IOP-less boot will be attempted.

One ramification of the above boot scenarios is that once a partition's Load Source is IOP-less, then even if the customer adds an upstream IOP later, PHYP will not boot IOP-based. Reconfiguring the hardware or deleting and recreating the partition are ways to force PHYP to abort the pre-search IPL phase and enter the search phase, allowing a change from IOP-less to IOP-based Load Source boot.

This information is accurate up through December 2006.

2.5.1 Media Devices that can attach to adapters with no controlling IOP

In this section we summarize devices that can be attached to adapters when there is no controlling IOP configured:

- ▶ DASD
 - All supported internal 10 K and 15 K DASD devices supported by the specific processor or I/O tower or drawer enclosure
- ▶ Optical devices
 - #6330, 6331, 6333 DVD-RAM
 - #6336, 6337 DVD-ROM
- ▶ DVD devices in M/T 7212/7206
- ▶ Tape devices
 - VXA-2: #1889/6485, 7206-VX2, 7212-102(1104)
 - VXA-320: #6279, 7206-VXA3, 7212-102(1114)
 - SLR-60: #5753/4684/9653, 7207-330, 7212-102(1107)
 - SLR-100: #5754/4487/4687, 7212-102(1108)
 - HH LTO2: #5755, 7212-102(1109)

This means that other devices not listed cannot attach to the dual mode IOAs (smart IOAs), which can run with or without an IOP. See also Table 2-2 on page 25 for a view of this support from a tape attachment viewpoint.

2.5.2 RAID5 adapters with write cache

Originally, the 5709 RAID Enabler Card (controller) was announced and shipped as having 16 MB write cache. In 2005, there was an unannounced manufacturing line break in which 40 MB write cache was shipped instead of 16 MB. Replacements of existing 16 MB cache cards with 40 MB cache cards at no charge are not provided.

Feature 5726 was introduced to differentiate the same functional cache card for the model 570.

New RAID5 adapters with write cache became available February 2006 and replaced the existing 5709 and 5726 technologies. These new adapter feature codes are:

- ▶ #9510 (no-charge feature code for Express Configurations with RAID)
- ▶ #5728 for use in the Model 570
- ▶ #5727 for use in Models 520 and 550

Key characteristics include:

- ▶ 40 MB write cache
- ▶ Adapter can work with or without a controlling IOP in a 1.9 GHz 520, 1.9 GHz 550, or 2.2 GHz 570 model
- ▶ No PCI slot required, processor enclosure placement only
- ▶ RAID5 Adapter with 40 MB write cache
- ▶ No RAID6 support
- ▶ A minimum of three disk units required for a RAID disk array
- ▶ Supports up to eight disks

This requires the #6574 4-Disk Slot Expansion Base Controller feature to be installed in the 520 processor enclosure or the #6592 4-Disk Slot Expansion Base Controller feature installed in the 550 processor enclosure.

The cache helps boost performance even if RAID function is not used. It is conceptually a daughter card off the base controller. Although the electronics are identical between 5727 and 5728, there is different physical packaging between 570 and 520/550. Figure 2-9 shows a picture of one of the new RAID5 adapters.



Figure 2-9 Raid5 adapter with write cache

One can find the cache size of the RAID5 adapter in the processor enclosure unit using System i Service Tools by displaying the resource details of the IOA. A V5R3 PTF MF34406 is required for this function.

Note: When using the System i Service Tools interface, if you look at the 2757, 2780, or newer large write and read cache disk controllers you see a value that represents the actual cache size, not the larger size cited in marketing documents, which assume a significant level of data compression. For example, for a #2780 IOA you see a value of 235 MB for write cache, not the 757 MB size often discussed in marketing documentation.

2.5.3 Internet Small Computer System Interface controller

This new controller can be used with or without an IOP for internal and external Small Computer System Interface (SCSI) devices with Low Voltage Differential (LVD) SCSI ports, and is designed as a replacement for features #5702/5705/5712/5715.

This controller has two internal and two external ports that can support disk, tape, DVD, and CD devices. Ports can be active two at a time if attached to tape but only one if attached to disk. Mixed mode use of ports is not supported at this time. This means that you can only use the internal or the external ports but not a combination of both.

Figure 2-10 shows an entry-level SCSI controller.




Figure 2-10 Entry SCSI controller

Installation of this entry SCSI controller is supported in processor enclosure or I/O tower/drawer PCI slots as described below under specific feature numbers.

Entry-level SCSI controller features using an IOP (V5R3 and later) include:

- ▶ #5736 IOP-based PCI-X disk, tape attachments
 - Disk or tape controller usage to POWER5 configurator tools (#0290=tape “flag”)
 - Disk controller to pre POWER5 configurator tools
 - i5/OS V5R3 or later
 - Supported on system models 520, 550, 570, 595, 270, 820, 830, 840, 800, 810, 825, 870, and 890
 - Supported in processor enclosures or I/O towers/drawers 5074/5079 (tape only), 0595/5095, 5094/5294, 8294/9194, 0588/5088 (tape only), 0578/5078 (tape only), and 5790 (tape only)
- ▶ 5766 IOP-based PCI-X tape controller attachment
 - Tape controller usage to pre-POWER configurator tools
 - i5/OS V5R3 or later
 - Supported on models 270, 820, 830, 840, 800, 810, 825, 870, and 890
 - Supported in processor enclosures or I/O towers/drawers 5074/5079, 0595/5095, 5094/5294, 8294/9194, 5074/5079, 0588/5088, 0578/5078, and 5790

Entry-level SCSI controller features without an IOP (V5R4 on POWER5 models and a 520 1.9 GHz model with V5R3 and LIC level V5R3M5) include:

- ▶ 5775 PCI-X controller supporting disk, tape attachments
 - Disk or tape controller usage to POWER5 configurator tools (#0290=tape “flag”)
 - i5/OS V5R4 or later support on models 520, 550, 570, 595
 - i5/OS V5R3 with LIC level V5R3M5 installed on 520 1.9 GHz models
 - Specific tape devices only supported (See Table 2-2.)
 - Supported in processor enclosures or I/O towers/drawers 5074/5079 (tape only), 595/5095, 5094/5294, 8294/9194, 0588/5088 (tape only), 0578/5078 (tape only), 5790 (tape only)

- ▶ 0647 PCI-X controller disk, tape attachments to Linux or AIX partition
 - Models supporting Linux or AIX 520, 550, 570, 595
 - Models supporting Linux 820, 830, 840, 800, 810, 825, 870, 890
 - Specific tape devices only supported (See Table 2-2.)
 - Supported in processor enclosures or I/O towers/drawers 5074/5079 (tape only), 595/5095, 5094/5294, 8294/9194, 0588/5088 (tape only), 0578/5078 (tape only), 5790 (tape only)

Table 2-2 Tape attachment considerations for 5775 and 0647 features

I5/OS tape drives supported ^a	Feature number	Minimum release level
30 GB QIC (SLR60)	5735/4684/9653, 7207-330, 7212-102 (1107)	V5R3 ^b , V5R4
50 GB QIC (SLR100)	5754/4487/4687, 7212-102 (1108)	V5R3 ^b , V5R4
80 GB VXA-2	1889, 7206-VX2, 7212-102 (1104)	V5R3 ^b , V5R4
160 GB VXA-320	6279, 7206-VXA3, 7212-102 (1114)	V5R3 ^b , V5R4
200 GB LTO-2 (HH)	5755, 7212-102 (1109)	V5R3 ^b , V5R4

- a. This column lists the tape devices supported by the 5775 and 0647 (IOP-less adapter) features. This means that the following tape devices are not supported:
- o Tape drives that can attach via a Fibre Channel adapter/controller
 - o Tape drives that can attach via a Low Voltage Differential (LVD) SCSI controller (5702, 5712, 5736)
 - o IBM 3580, 3581, 3582, 3583, 3584, 3576, 3590, 3592
 - o Any tape device attached via a High Voltage Differential (HVD) SCSI controller (#2749)
- b. V5R3 only on 520 1.9 GHz model with LIC V5R3M5

2.5.4 Small Computer System Interface controller for disks with 90 MB cache

This SCSI controller can be used with (#5737) or without an IOP (#5776) for internal SCSI devices. A replacement for #5703 (with 40 MB write cache), this controller has 90 MB write cache functionality and RAID5 and RAID6 capabilities.

Like the 5703, this new controller can use internal ports only for up to 12 disk drives. This new controller is supported in processor enclosure or I/O tower/drawer PCI slots:

- ▶ Up to 12 disk drives (same as #5703)
- ▶ Reported resource or CCIN: 571B
- ▶ Supported in processor enclosure or I/O tower/drawer PCI slots: 0595/5095, 5094/5294, 8294/9194

Important: The IBM Redbook *IBM System i5, eServer i5, and iSeries Systems Builder IBM 5/OS Version 5 Release 4 - January 2006*, SG24-2155, contains a chapter that lists the CCIN number associated with the orderable feature number. The updated redpaper *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition)*, REDP-4011, also contains tables listing CCIN numbers associated with the orderable feature numbers. It also provides important information identifying a PCI slot or I/O device enclosure considerations for the various IOPs and IOAs.

It is common for multiple orderable feature numbers to all display the same CCIN number once installed on the system. Always refer to the IBM Redbooks Web site for the latest versions of these documents.

This controller is supported as three different feature codes:

- ▶ With IOP V5R3 and up:
 - #5737 Disk controller for system i5 models, and 270 and 8xx models
- ▶ Without IOP and V5R4:
 - #5776 Disk controller for i5 systems
 - #0647 Disk controller for Linux or AIX

As originally announced, RAID6 support is English only, running V5R3 microcode levels V5R3M0 and V5R3M5.

The original V5R3M0 PTFs supporting RAID6 are MF39611 and MF39933.

The original V5R3M5 PTFs supporting RAID6 are MF39614 and MF39926.

Check with the IBM e-configurator to determine whether the English-only support remains in effect.

The orderable feature numbers #5737, #5776, and #0647 offer good price performance but do not provide the highest capacity performance on a System i configuration.

There are some situations where we might recommend a #5737 or 0648 (direct attach Linux/AIX). Situations of this kind include:

- ▶ There is a small or modest number of disk drives being supported, especially if mirroring is to be used.
- ▶ Performance is not the primary consideration. Note that this comment is relative to high disk I/O operations per second environments. At physical disk I/O rates up to 600 per second, the #5737, #5776, and #0647 feature number controllers deliver very acceptable performance.
- ▶ Price is a very important consideration and there are a modest number of disk drives.

The #2757 and #2780 PCI-X Ultra4 RAID Disk Controller with 1 GB read cache and up 757 MB write cache, and the even higher speed capacity disk IOAs (several feature numbers including #5738 and #5777) with 1.6 GB read cache and 1.5 GB write cache offer even higher disk I/Os per second capacity.

The #2757, #2780, #5738, and #5777 also support the attachment of an auxiliary write cache IOA, which provides the added protection of redundant cache should a hardware failure occur involving the base write cache.

In general, these higher performance disk controllers (IOAs) should be considered first before selecting the #5737, #5776, or #0647 technology controller.

2.5.5 High performance disk controllers with 1.5 GB write cache, 1.6 GB read cache

During February 2007 these new PCI-X disk controllers, #5738 (IOP required) and #5777 (dual mode, with or without an IOP), were announced. The #5738 and #5777 disk controllers provide additional performance capacities with faster electronics and twice the write cache as their predecessor disk controller, the #2780.

i5/OS V5R3 or V5R4 is required. i5/OS V5R4 is required to run in IOP-less mode.

Note: Also announced in February 2007 were new disk IOAs with the same cache capacities as the #5738 and #5777, but for use only with the newly announced IBM TotalStorage® EXP24 disk enclosures. These controllers support disk drives plugged only in the EXP24 disk enclosure. The #5738 and #5777 support disk drives plugged into other disk I/O towers and enclosures.

A later topic discusses these *EXP24 disk enclosure only* disk IOAs.

Mirroring with the #5738/#5777 is supported via operating system support.

RAID-5 and RAID-6 capabilities are supported via the disk controller. However, these new controllers will not enable RAID protection unless a new, separate auxiliary write cache IOA (feature #5582 (for the #5738) and #5583 (for the #5777)) is also installed.

These disk controllers use a PowerPC® processor and double data rate 2 (DDR2) memory technology.

The #5738 is an IOP-based controller for 800, 810, 825, 870, 890, and later system models. You can install the #5738 and #5777 in a 825 system unit or a #5094/#5294 or #8294/#9194 I/O tower. Unlike the #2780 and #5580, you cannot install them in a model 800 or 810 system unit or in a #5074/#5079 or #0595/#5095 I/O tower or drawer.

The new auxiliary write cache IOA. #5582 (with the #5738) and #5583 (with the #5777) can also be the auxiliary write cache for a #2757 disk controller or a #2780 disk controller. These features are:

- ▶ #5590: #2780 disk controller and 1.5 GB auxiliary write cache IOA
- ▶ #5591: #2757 disk controller and 1.5 GB auxiliary write cache IOA

There is no performance advantage to using a larger auxiliary write cache, but it does provide future potential flexibility if the #2780 or #2757 is replaced.

It is only with the #5738 and #5777 that RAID support requires the auxiliary write cache to be installed.

In summary, four features represent the new disk controller either with or without an auxiliary write cache IOA, running under i5/OS:

- ▶ #5738: disk controller without auxiliary write cache IOA (requires IOP)
- ▶ #5777: disk controller without auxiliary write cache (IOP-less) IOA
- ▶ #5582: disk controller with auxiliary write cache IOA (requires IOP)
- ▶ #5583: disk controller with auxiliary write cache (IOP-less) IOA

New feature conversions of existing #2780/#5580 or #2757/#5581 disk controllers for the #5738, #5777, #5582, and #5583 disk controllers are available.

In the next topic the new IBM TotalStorage EXP24 disk enclosures are covered. These disk enclosures attach to the system via specific existing disk controllers and new (as of February 2007) disk controllers that are logically equivalent to the #5582 and #5583 described in this section.

For more information about the new hardware announced February 2007 refer to:

- ▶ *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition)*, REDP-4011 or later. There are important PCI plugging rules for these controllers.
- ▶ Planned redpaper REDP-3919.
- ▶ January 2007 or later System i Performance Capabilities Reference manual.

2.5.6 IBM TotalStorage EXP24 disk enclosures

The new System i #5786 or #5787 feature IOAs specifically connect an IBM TotalStorage EXP24 disk drawer or tower. The EXP24 disk enclosures offer an economical attachment of up to 24 disk drives in a small space — external to a System I/O tower or previously available enclosures.

This EXP24 enclosure does not have any PCI slots. You insert disk drives that are driven by one or more disk controllers located in a System i system unit or I/O tower or drawer. The EXP24 is attached via one or more external SCSI connections. i5/OS V5R3 or later is required.

Like previous internal (integrated) System i disk storage, the EXP24 disk tower or drawer and its disk drives are optional feature numbers. The disk drives that are placed in the EXP24 disk enclosure are identical to the drives provided for 520, 550, 570, and 595 models. Different disk features designate their placement in a #5786 or #5787 EXP24 disk drawer or tower to help IBM configuration tools.

The supported disk features include the i5/OS- supported 15,000 rpm drives #1266 (equivalent to #4326), #1267 (equivalent to #4327), and #1268 (equivalent to #4328), and the AIX 5L/Linux supported 10,000 and 15,000 rpm drives #1293 to #1299. Existing i5/OS 10,000 rpm disk drives are not supported in this enclosure.

The #5786 is a disk drawer that requires only 4U (or 4 EIA) units of rack space. The #5787 is a tower or deskside configuration. Conversions from rack-to-tower or tower-to-rack are not provided. Each #5786 and #5787 has redundant power supplies and dual power cords for higher availability.

Each EXP24 disk unit has four 6-packs of disk slots. A disk controller can drive either one or two 6-packs per SCSI connection. The choice between 6 or 12 disk slots is determined by the use of the #5742, EXP24 6/12 Disk Slot Enabler feature, or a #5741, EXP24 6 Disk Slot Enabler. One Disk Slot Enabler feature must be used for each 6-pack into which a disk drive is placed. At least one #5742, 6/12 Disk Slot Enabler feature must be used to create a 12-pack, and each 12-pack requires two disk slot enablers.

This new disk enclosure attaches to specific supporting System i disk controllers (IOAs) via SCSI cables. Figure 2-11 is a logical representation of the connections and lists the supporting disk IOAs, which include disk controllers available prior to February 2007.

You see the HSL/SPCN cable indicating where the supporting disk controller can be placed.

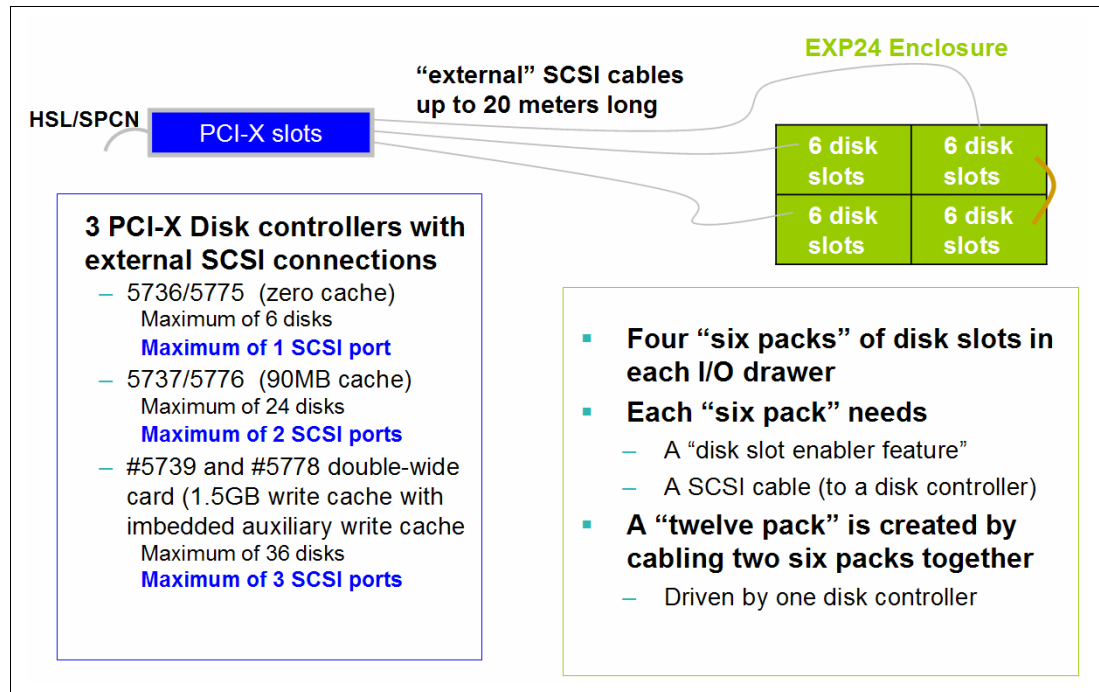


Figure 2-11 EXP24 example attachment to System i adapters

You see the SCSI cables coming out of the supporting disk controller with each SCSI port connection supporting up to 6 disks or 12 disks if two 6-packs are cabled together (short SCSI cable). A cable from the disk controller connects to an *enabler port* (#5741 or #5742) on the EXP24 enclosure.

Disk controllers that support the EXP24 disk drawer or tower include:

- ▶ Entry-level disk controller features #0647, #5736, and #5775 disk controllers with zero write cache
- ▶ Disk controller features #0648, #5737, and #5776 with 90 MB write cache
- ▶ New as of February 2007 disk controller features #5739 (IOP controlled) and #5778 (IOP-less) with 1.5 GB write cache and 1.6 GB read cache

There are many configuration and cabling possibilities when ordering and setting up usage of the EXP24 disk enclosure. You should work with your IBM or authorized business partner to ensure proper planning, ordering, and set up.

2.5.7 Disk Controllers for EXP24 disk enclosures

The disk controller features #5739 (IOP controlled) and #5778 (IOP-less) with 1.5 GB write cache and 1.6 GB read cache are explicitly for use with the IBM TotalStorage EXP24 disk drawer or tower. The write and read caches are built into its double-wide packaging, eliminating the need for a separate auxiliary write cache IOA and cable. This controller can provide mirroring, RAID-5, and RAID-6 capabilities equivalent to the #5582 and #5583 disk controllers with auxiliary write cache features installed.

A single one of these controllers can be located in a 5xx or 825 model system unit, or in a #0588/#5088, #5094/#5294, #8094/#9094, or #8294/#9194 expansion unit. It cannot be located in a #5074/#5079 or #0595/#5095 expansion unit, or a model 800 or 810 system unit. It supports only disk drives physically installed in an EXP24 disk enclosure. If the EXP24 1.5 GB disk controller is installed in a system unit or #5094/#5294, #8094/#9094, or #8294/#9194 expansion unit, it does not run any disk drives in these enclosures.

For more information about the new hardware announced in February 2007 refer to:

- ▶ *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition)*, REDP-4011 or later. There are important PCI plugging rules for these controllers.
- ▶ Planned redpaper REDP-3919.
- ▶ January 2007 or later System i Performance Capabilities Reference manual.

Note: The cache sizes cited in most documentation are actually the *up to* values that assume a significant amount of data compression. The following list shows the actual cache size as well as the *up to* values for the fastest disk IOAs supported on System i configurations:

- ▶ 2757:
 - Write cache: 235 MB with up to 757 MB
 - Read cache: 0
- ▶ 2780:
 - Write cache: 235 MB with up to 757 MB
 - Read cache: 256 MB with up to 1 GB
- ▶ CCIN 571B (#5737, #5776, #0648):
 - Write cache: 90 MB
 - Read cache: 0
- ▶ CCIN 571E/574F(auxiliary write cache) (#5738, #5777, #5582, #5583):
 - Write cache: 390 MB with up to 1.5 GB
 - Read cache: 415 MB with up to 1.6 GB
- ▶ CCIN 571F/575B (built-in auxiliary write cache) (#5739, #5778, #5781, #5782, #5799, 5800):
 - Write cache: 390 MB with up to 1.5 GB
 - Read cache: 415 MB with up to 1.6 GB

When using the System i Service Tools interface, if you look at the 2757, 2780, or newer large write and read cache disk controllers you see a value that represents the actual cache size, not the larger size cited in marketing documents, which assume a significant level of data compression. For example, for a #2780 IOA you see a value of 235 MB for write cache, not the 757 MB size often discussed in marketing documentation.

See “Auxiliary write cache overview” on page 44 for more information about the physical placement of the auxiliary write cached.

See 16.2.4, “Disk hardware performance - smaller cache disk controllers” on page 417.

2.5.8 Four Gbps fibre channel adapters

The 4 Gbps Fibre Channel Disk and Fibre Channel Tape Adapter have two different cards that can be only used with an IOP. The two different cards support disk and tape devices and are designed to replace the 2 Gb FC Adapters #2766 and #2787 (disks controller) and the 2 Gbps FC Adapters #2765 and #5704 (tape controller).

FC #5760 for i5 systems and Model 270 or 8xx is used for disks, while FC #5761 is used for tape devices.

Both cards can connect to a 4 Gbps, 2 Gbps, or 1Gbps SAN and are supported on systems running i5/OS V5R3 or later. These are also supported on IBM System i5 models and on models 270 and 8xx.

Figure 2-12 shows a picture of the adapter.



Figure 2-12 4 Gbps fiber channel adapter

These new adapters provide better throughput and better response times. See Table 2-3 for details.

Table 2-3 Card feature codes and performance

Card link speed	Card feature	Maximum throughput over a single connection (MBps)	Collection services disk response time (milliseconds)	Performance data collection utility response time (milliseconds)
1Gbps	#2766	62	8.3	2.9
2 Gbps	#2787	86	5.7	1.3
4 Gbps	#5760	148	3.5	1.2

Performance Data Collection Utility (PDCU) is an internal IBM DS tool that provides measured response times on the storage subsystem device only. Collection services is system response times. These measurements are taken while running the same I/O workloads.

As shown in Table 2-3, the #2787 can become very busy when attached to a 2 Gbps SAN fabric. If this happens, you may see serious throughput improvement by going to a #5760, even if the SAN fabric stays at 2 Gbps.

Table 2-4 shows a summary of fibre channel disk and tape features for reference purposes.

Table 2-4 Disk controller reference

i5/OS disk controller Feature Codes (FCs)		
4 Gbps	PCI-X	#5760
2 Gbps	PCI-X	#2787
1Gbps	PCI	#2766

Table 2-5 is a tape controller reference.

Table 2-5 Tape controller reference

i5/OS tape controller FCs		
4 Gbps	PCI-X	#5761
2 Gbps	PCI-X	#5704
1 Gbps	PCI	#2765

2.5.9 Existing communication adapters without input/output processor

There are existing adapters for both Ethernet and WAN that can be used without any input/output processor (IOP) provided they are on a System i5 server running on i5/OS Version 5 Release 4. This section lists those not previously discussed, such as the 2-port Ethernet LAN adapters 5706/5707.

The following are the Ethernet and WAN adapter feature codes that can be used without IOP and their special considerations:

- ▶ Feature #6800 and #6801 Ethernet adapters without IOP
 - #6800 is #5700 Optical single port 1Gb Ethernet adapter.
 - #6801 is #5701 UTP single port 10/100/1000 MGb Ethernet adapter.
- ▶ Feature #6803 and #6804 WAN adapters without IOP
 - Without IOP only integrated modem lines will work.
 - #6803 is #2793 (non-CIM) Adapter, #0614 Adapter for Linux and AIX.
 - #9493 is #9793 (non-CIM) Base ECS Adapter.
 - #6804 is #2794 (CIM) Adapter, #0615 for Linux and AIX.
 - #9494 is #9794 (CIM) Base ECS Adapter.

Communication functions on RS-232 port are not available when used without IOP.

The notable advantage of dual mode cards is that they can be used by IOP-based 8xx customers (with the proper level of software).

The 5700/6800 technically supports auto-negotiation, but will only negotiate to 1000 Mbps/Full duplex. This is the gigabit fiber standard, which is only 1000 Mbps/Full.

The 5701/6801 supports auto-negotiation if configured with *Auto in the i5/OS line description. The parameter value *Auto makes it possible to negotiate to the highest capability of the link partner. This is usually 100 Mbps/Full or 1000 Mbps/Full.

If the link partner is not capable of auto-negotiation then hard coded values maybe entered (that is, 100 Mbps/Half, 1000 Mbps/Full, and so forth). The auto negotiation in this configuration is turned off and we force the link to run at the configured speed.

The #5701/#6801 PCI 1 Gbps Ethernet UTP IOA allows an iSeries server to attach to IEEE standard 802.3Z high-speed (1 Gbps) Ethernet LANs. It may also be used to directly connect to existing 10 Mbps or 100 Mbps Ethernet LANs. However, it will not run at gigabit speeds in this configuration. Crossover cables are not supported. The #5701/6801 adapter supports a UTP CAT 5 media interface.

You should always configure the i5/OS LAN line description for auto-negotiation for gigabit interfaces. If that fails then you must understand what the link partner is configured for and set the i5/OS line description to the supported speed.

Important: These Ethernet LANs do not support SNA. SNA protocol flows through them successfully when the SNA data and flow control are packaged within AnyNet or Enterprise Extender packets.

2.6 AIX and Linux

There are several enhancements to AIX 5L and Linux on System i5 partitions in the areas of TCP/IP and network connectivity.

Network connectivity enhancements

One enhancement for AIX 5L/Linux (not i5/OS) is the support for 4 Gb Fibre Channel Adapter, single port feature #5758 (physically the same as i5/OS #5761 announced January 2006) and the 4 Gb Fibre Channel Adapter with dual port - feature #5759.

In addition, another enhancement is that the 10 Gb Ethernet adapters #5721/#5722 can now be shared with i5/OS.

There are also changes improving configuration flexibility in terms of PCI slot usage and IOA placement. The following are examples of these configuration flexibility improvements for System i5 models 520, 550, 570, and 595:

- ▶ Can now place an HMC-designated, AIX 5L/Linux-owned IOA in the middle of i5/OS IOP/IOAs. This was not possible in the early releases, as the AIX 5L/Linux IOA placement was less flexible.
- ▶ Allows better and easier usage of PCI slot characteristics (speed, length slot, and so on).

TCP/IP and TCP/IP Offload Engine

TCP/IP Offload Engine (TOE) is a technology implementation that is increasing in popularity. It is a way to offload the CPU and some I/O subsystems of TCP/IP protocol when handling the higher speeds (1 Gbps and higher) of communication.

In the last three years, the communication speed in Ethernet systems has increased faster than computer processor speed. This produces an input/output (I/O) bottleneck. The processor, which is designed primarily for computing and not for I/O, cannot keep up with the data flowing through the network. As a result, the TCP/IP flow is processed at a rate slower than the speed of the network. TOE addresses this problem by removing the burden (offloading) from the microprocessor and I/O subsystem.

The manner in which TOE is implemented depends on your needs. Considerations include flexibility, scalability, and performance.

The following is a summary of these features:

- ▶ PCI-X iSCSI TOE (TCP/IP Offload Engine) adapters for attaching external disk/tape
 - Copper #5713
 - Fibre #5714
- ▶ Four displays (primarily for AIX 5L console)
 - #3641, #3643, #3642, and #3645
- ▶ #5740 PCI-X 10/100/1000 Gbps 4 port Ethernet adapter
- ▶ #1898 SCSI Ultra-320 146.8 GB 15k rpm disk
- ▶ New hardware features also available for i5/OS
 - PCI-X Disk/Tape controller #0647
 - PCI-X Disk Controller 90MB #0648
 - VXA-320 tape device #6279

Supported IXS/IXA Linux versions

No new Linux (Intel®) versions were announced for iSeries integrated xSeries server support as of October 2006. As of this date, the supported releases are:

- ▶ Red Hat Enterprise Linux 3.0 ES Edition (IXS and IXA)
- ▶ Red Hat Enterprise Linux 3.0 AS Edition (IXS and IXA)
- ▶ Red Hat Enterprise Linux 4 ES Edition (IXA only)
- ▶ Red Hat Enterprise Linux 4 AS Edition (IXA only)
- ▶ SUSE Linux Enterprise Server 8 (IXA only)
- ▶ SUSE Linux Enterprise Server 9 (IXA only)

There is a planning statement for Linux on iSCSI. Refer to the following Web site for the latest information including specific Windows operating system and Linux distributions and xSeries and BladeCenter models supported:

<http://www.ibm.com/systems/i/bladecenter/windows/>

2.7 iSCSI support

The iSCSI Host Bus Adapter (HBA) is new and is supported in the i5/OS partition running on System i5 models. This adapter enables the integration xSeries and BladeCenter systems to the System i5 via industry-standard Ethernet technology, as shown in Figure 2-13. iSCSI support is discussed in detail in 5.2, “System x and BladeCenter support via iSCSI” on page 121.




Figure 2-13 iSCSI xSeries: System i5 server connectivity

Full mixing and matching of Integrated xSeries Server (IXS), Integrated xSeries Adapter (IXA), and iSCSI attachment of xSeries server will be supported on V5R4 sometime during 2006. iSCSI attachment greatly expands the number and power of xSeries servers that can be integrated with i5/OS and for the first time includes BladeCenter support.

2.8 V5R4 POWER5-based systems and i5/OS console support

The V5R4 Power5 and the i5/OS offer many new enhancements for a variety of console configurations. Enhancements like the embedded ports (part of the processor enclosure) for Operations Console, IOP-less adapters, and new specify codes are just an example. More information about these topics is available in Chapter 11, “i5/OS-based consoles and logical partitioning” on page 245.

In V5R4 System i5 model 520 and 550 servers, there are currently seven console options available to choose from.

Figure 2-14 shows the seven consoles available.

Option	PCI slots used	Ethernet port used	User Interface	Announced
1 5250 Twinax console	IOP + twinax IOA	none	Green screen	2004
2 Ops Console*	IOP + WAN IOA	none	GUI / Ops Navigator	2004
3 LAN Console*	IOP + LAN IOA	No CEC port, uses LAN IOA port	GUI / Ops Navigator	2004
4 LAN Console*	LAN IOA	No CEC port, uses LAN IOA port	GUI / Ops Navigator	Jan 2006
5 LAN Console*	none	Uses one of CEC's two user ports	GUI / Ops Navigator	Jan 2006
6 HMC 5250	none	Uses one of CEC's two HMC ports	Green screen	2004
7 Thin console (5250)	none	Uses one of CEC's two HMC ports	Green screen	July 2006

Figure 2-14 System i5 520 and 550 console options

The thin console is a significant additional console option announced mid 2006. The following topics highlight some of these console options. Chapter 11, “i5/OS-based consoles and logical partitioning” on page 245, contains more 5250 console information.

2.8.1 Hardware Management Console

It is now possible that a single i5 system can be managed by up to two Hardware Management Consoles (HMCs) at a time, and a single HMC can manage up to 48 non-590/595 Servers or thirty-two 590/595 Servers and up to 254 partitions. Refer to 11.1.2, “Hardware Management Console” on page 249, for details on HMC console in V5R4.

The following HMC devices and models are supported by System i configurations:

- ▶ 7310-C04
- ▶ 7310-C03
- ▶ 7310-CR3
- ▶ 7310-CR2
- ▶ 7310-CR4

Refer to 11.1.2, “Hardware Management Console” on page 249, for details on the HMC console in V5R4.

2.8.2 Thin console

This newly introduced hardware is made by Neoware for IBM. This console can attach to a single partition System i5 model 520 or 550.

Marketed as IBM machine type 9944 model 100, this machine’s target market is the 520/550 system environments.

Neoware provides hardware warranty and software support. At the time this book was written, this product is supported in the U.S., Canada, Italy, France, Germany, Spain, UK, Netherlands, Belgium, Switzerland, Sweden, Austria, Denmark, Portugal, Norway, and Finland.

For additional information about this topic see 11.1.3, “Thin Console” on page 251.

2.9 VXA-32 Tape Drive support

This VXA-320 Internal Tape Drive is only supported on POWER5 and POWER5+ systems. It is supported with i5/OS V5R3 with V5R3M0 LIC or later. The device is driven by the internal SCSI controller with LVD interface. Figure 2-15 shows the different tape devices, their speed, and cartridge capacity.

	520 / 550 FC	5094 / 5294 FC	Native Speed	Cartridge Capacity	
				Native	Compressed
LTO-2	#5755	N/A	24 MBps	200GB	400GB
VXA-320	#6279	N/A	12 MBps	160GB	320GB
VXA-2	#1889	#4685	6 MBps	80GB	160GB
QIC 50	#5754	#4687	5 MBps	50GB	100GB
QIC 30	#5753	#4686	4 MBps	30GB	60GB

Figure 2-15 Tape speeds and capacity chart

VXA-320 devices in #5094/#5294 tower are driven by PCI Disk Controller, with IOP in V5R3, but can be without IOP in V5R4. Figure 2-16 shows the media compatibility chart for various cartridges.

Media Compatibility	VXA-320 drive	VXA-2 drive	VXA-1 drive
40 GB VXA cartridge	Read	Read/Write	Read/Write
80 GB VXA cartridge	Read/Write	Read/Write	NA
160 GB VXA cartridge *	Read/Write	NA	NA

Figure 2-16 Media compatibility chart

A VXA tape drive located in an external bridge box is supported.

2.10 IBM System Storage TS3400 Tape Library

The IBM System Storage TS3400 Tape Library is a new tape library product that combines IBM tape drive and automation reliability with entry-level pricing. The new model was introduced during February 2007 and is supported by many operating systems including System i i5/OS, System p AIX 5L, System x™, and servers running HP-UX, Sun™ Solaris™, Windows, and Linux.

You can attach the TS3400 via a supported System i Fibre Channel adapter to System i models running i5/OS V5R3 or later:

- ▶ 820/830/840
- ▶ 800/810/825/870/890
- ▶ 520/550/570/595

The TS3400 supports high-performance drives, unattended backup and brings archive, open system attachment flexibility, and high capacity and performance to the marketplace. The TS3400 Tape Library capabilities include:

- ▶ Support for up to two IBM System Storage TS1120 Tape Drives Model E05 with 4 Gbps dual-ported switched fabric Fibre Channel attachment, ordered separately
- ▶ Support for the 3592 Tape Cartridge media (no 3592 tape device support)

- ▶ 18 cartridge slots physical capacity in two removable magazines (Additional magazines can be ordered.)
- ▶ Encryption capabilities designed to work with the IBM Encryption Key Manager component supporting System Managed and Library Managed encryption, as well as support for Application Managed encryption
- ▶ Control path and data path failover, bar code reader, dual power supplies, remote management as standard features
- ▶ Sequential and random access mode automation
- ▶ Optional rack mount kit

You partition the library into two logical libraries:

- ▶ Each logical library comprised of one drive and one magazine.
- ▶ Run any single-drive logical library in either sequential (autoloader) mode or random (library) mode.

Figure 2-17 shows a picture of the TS3400, and Figure 2-18 shows a functional view of where the TS3400 fits with other IBM tape system library products.



Figure 2-17 IBM System Storage TS3400 Tape Library

Tape Library	TS3100	TS3200	TS3310	TS3500	TS3400	TS3500
Number of drives	1 LTO	1-2 LTO	1 - 10 LTO	1 - 192 LTO	1-2 TS1120	1 - 192 3592 or TS1120
Drive data rate	80 MBps	80 MBps	80 MBps	80 MBps	104MBps	104MBps
Max number of cartridges	22	44	218	6887	18	6260
Max native cartridge capacity	400GB	400GB	400GB	400GB	700GB	700GB
Native library capacity	up to 8.8TB	up to 17.6TB	up to 87.2TB	up to 2.75PB	up to 12.6TB	up to 4.3 PBs
Encryption support	No	No	No	No	Yes	Yes
Interface	4Gbps FC 160Mbps SCSI	4Gbps FC 160Mbps SCSI	4Gbps FC SCSI-LVD	4Gbps FC SCSI-LVD	4Gbps FC	4Gbps FC

Figure 2-18 The TS3400 Library among other Library products

2.11 CPU power redundancy

There is no change in physical redundant power supply options available on IBM System i5 models. The following are the features that can be ordered:

- ▶ The #5158 is an optional 850W power supply that provides redundant power for the Model 520 system unit.
- ▶ The #5159 provides an optional 850W power supply that installs in a 520 system unit with processor.
- ▶ The #7889 is an optional 1475W power supply, which provides redundant power for the Model 550 system unit. A second line cord is required.
- ▶ The #7768 provides a redundant processor power regulator for the Model 570 single enclosure system.
- ▶ The #7875 is a processor power regulator for the Model 570. One #7875 is required per processor card. A single redundant #7875 can be ordered to provide redundant power regulation for the Model 570.
- ▶ All 595 configurations include several redundant power supplies.

The redundant power features for 520, 550, and 570 all require an additional processor line cord feature to be ordered.

Minimum operating system levels should be i5/OS V5R3 with V5R3M5 LIC or later, AIX 5L for POWER V5.2 for IBM eServer, Red Hat Enterprise Linux AS for POWER Version 3, or SUSE Linux Enterprise Server 9 for POWER.

See *IBM eServer i5, iSeries, and AS/400e System Builder IBM i5/OS Version 5 Release 3*, SG2155-12, for more information.

2.12 Rack-related changes during 2006

This section summarizes the rack-related changes made during 2006:

- ▶ All refreshed with new 1.9 GHz and 2.2 GHz processor models get new rack specify codes.
- ▶ Rack modem tray #6586 was introduced.
- ▶ The #7307 enclosure feature for the #5790 I/O unit was introduced. It includes adjustable depth rails, replacing #7311 to hold up two 5790 PCI-X units.
- ▶ The 5094 i/O tower in a rack RPQ became available in a wider number of countries and added support for 800, 810, 825, 870, and 890 models.
- ▶ The #0588 I/O tower was withdrawn from marketing. The #5790 is the closest counterpart, though it has fewer PCI slots per I/O drawer.
- ▶ The #6664 shorter (7 US feet versus 12 US feet) Power Distribution Unit (PDU) power cord for the 5790 I/O unit was introduced.

2.13 ytty

Announced during February 2007, the #5096 and #5296 (two stacked #5096s) are responses to customers asking for I/O towers for non-disk I/O at a reduced price. The #5096 has

14 PCI-X slots and zero disk slots for 520, 550, 570, or 595 and 800, 810, 825, 870, or 890 clients who need a large number of PCI slots using a minimum number of HSL loop positions.

i5/OS V5R3 or later with PTFs is required. The #5096 is essentially a #5094 that has no disk slots and thus has a lower maintenance price. The #5096 does not support disk units or removable media devices such as tape or DVD drives located inside the #5096, but does support PCI adapters attaching disk units, for example, the EXP24 disk enclosure, and tape or optical drives external to the #5096. The #5096 can be upgraded via MES to a #5094.

The #5296, which has 28 PCI-X slots and zero disk slots, is essentially two #5096 disk controllers placed in a 1.8m rack. The #5296 can be upgraded to a #5294.

Some features of the #5096 and 5296 that are identical to those of the #5094 and 5294 include cabling (HSL,SPCN, and electrical power) and power supply redundancy features.

For more information see Announcement letter 107-083.

2.14 IBM System i5 code administration and terminology

The IBM eServer i5 layered code structure makes the IBM eServer i5 platform even more flexible. It also enables easy accommodation of different operating systems. This section is intended to lessen the confusion in terminology among different levels of *microcode* on the Hardware Management Console (HMC), the POWER5-based hardware system, and the *operating system* microcode interface.

First, a definition of terms shown on the slide. For i5/OS, TIMI and the layers above the POWER5 Hypervisor are still in place. OS/400 and i5/OS interface through the TIMI, which eliminates the need to understand the hardware details of the system underneath the operating system. TIMI interfaces to the microcode, called System Licensed Internal Code (SLIC) prior to POWER5 technology. On POWER5, SLIC is changed and enabled for interfacing with the POWER5 Hypervisor. In Figure 2-19 we name this *updated SLIC* "Licensed Internal Code (LIC)". The POWER Hypervisor™ code is unique to POWER5 hardware and includes LPAR configuration and management capabilities.

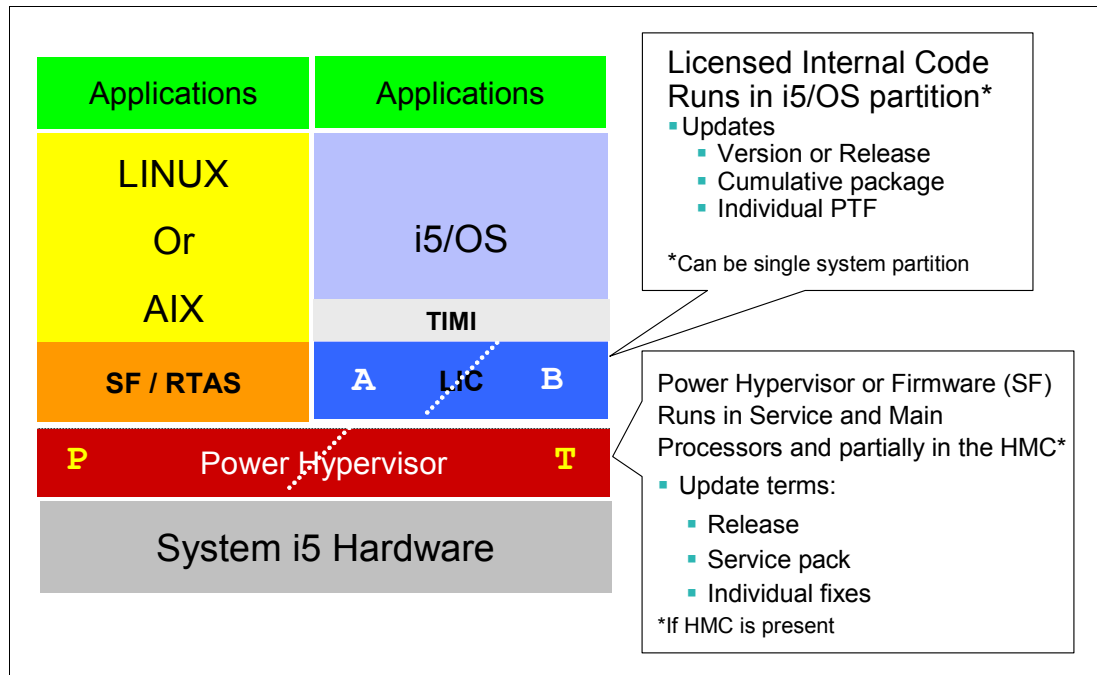


Figure 2-19 System i5 terminologies

A and B within the LIC layer indicate *sides* or copies of the microcode. A represents the permanent copy of the microcode. B represents the microcode containing fixes that have not been made permanent. Typically the system operates from the B side.

For the AIX-5L and Linux operating systems, the layers above the POWER5 Hypervisor are similar, but their content is characteristic for each operating system. The layers of code supporting Linux and AIX-5L consist of System Firmware and Run-Time Abstraction Services (RTAS):

► System Firmware (SF)

This is separate and distinct from the POWER Hypervisor System Firmware. This System Firmware is composed of Low Level Firmware and Open Firmware. Low Level Firmware is code that performs server unique input/output (I/O) configurations such as high-speed link (HSL)-2/RIO-G loops and PCI-X bridges. Open Firmware contains the boot time drivers (for example, SCSI, SSA, Token Ring, and Ethernet), the boot manager, and the device drivers required to initialize the PCI adapters and attached devices. For the remainder of this description SF does not refer to this level. Rather, SF refers to microcode levels of the Power Hypervisor.

► Run-Time Abstraction Services (RTAS)

This consists of code that supplies platform-dependent accesses and can be called from the operating system. The intent of this code is to minimize the need for an operating

system to understand hardware unique details. These calls are passed to the POWER Hypervisor that handles all I/O interrupts.

POWER5 technology-based hardware is the code layer called the *POWER Hypervisor* (sometimes referred to as *PHYP*).

This code is part of the System Firmware (SF) shipped with the IBM eServer i5 hardware. The POWER Hypervisor resides in flash memory on the Service Processor as well as in a small portion of the main processors. If the Hardware Management Console is managing the system, there is HMC System Firmware as well. Both the HMC firmware and the POWER Hypervisor firmware have levels of code. The HMC firmware level and the POWER Hypervisor firmware need to be within a range of levels in relation to each other for maximum efficiency. This firmware performs the initialization and configuration of the IBM eServer i5 hardware, as well as the virtualization support required to implement capacity on demand. Limited LPAR and capacity on demand management is supported without an HMC. However, full multiple operating system LPAR creation and management and up to 254 partitions concurrently on the eServer i5 servers require an HMC at fairly recent levels of code.

P represents the permanent copy of the firmware, while *T* represents the temporary copy of the firmware. T means fixes can be removed from this copy.

In summary, System Firmware has multiple components that run in multiple physical locations/processors. The four key processors are:

- ▶ POWER5 system main processor
- ▶ Service processor (also known as FSP)
- ▶ Bulk power processor (available with System i5 595 models)
- ▶ HMC (if part of the configuration)

2.15 V5R4 RAID protection

With RAID5 the disk array in the RAID set can continue to be usable when only one disk within the array set has failed. With RAID6 the disks can continue to be usable when up to two disks within the array set have failed. Note that neither RAID5 or RAID6 solutions help protect against the failure of three disks, the attached disk controller, or the attached IOP. Refer to Figure 2-20 on page 43 for a summary of disk data protection options.

Mirroring can help protect to a higher level than RAID protection, starting at the disk controller (IOA) level, and move up to the IOP level and to higher protection at the bus level.

Auxiliary write cache can help protect against extended outages. Auxiliary write cache for the #2757 or #2780 was introduced during 2005 and can be used with i5/OS releases V5R2, V5R3, and V5R4. Auxiliary write cache allows the primary write cache contents to be recovered much more quickly than a rebuild. If the primary write cache fails, you must stop the system and have an authorized service person recover the data and install a replacement primary write cache.

Although not dependent on RAID protection and not new with V5R4, we review basic auxiliary write cache support available on the System i family of models in “Auxiliary write cache overview” on page 44.

Mirroring should be the first protection strategy considered for all internal disk storage, which is important to running the customer’s business.

Note that this book provides a high-level review of protection mechanisms available with external disks in 2.16, “SAN and TotalStorage DS6000 and DS8000” on page 48.

Figure 2-20 shows RAID5, RAID6, and Mirroring disk protection schemes prior to the availability of the new disk IOAs announced February 2007. Figure 2-21 on page 44 shows an updated depiction of the schemes available considering the February 2007 disk IOAs, which support RAID5 or RAID6. Both figures include commentary on using the auxiliary write cache.

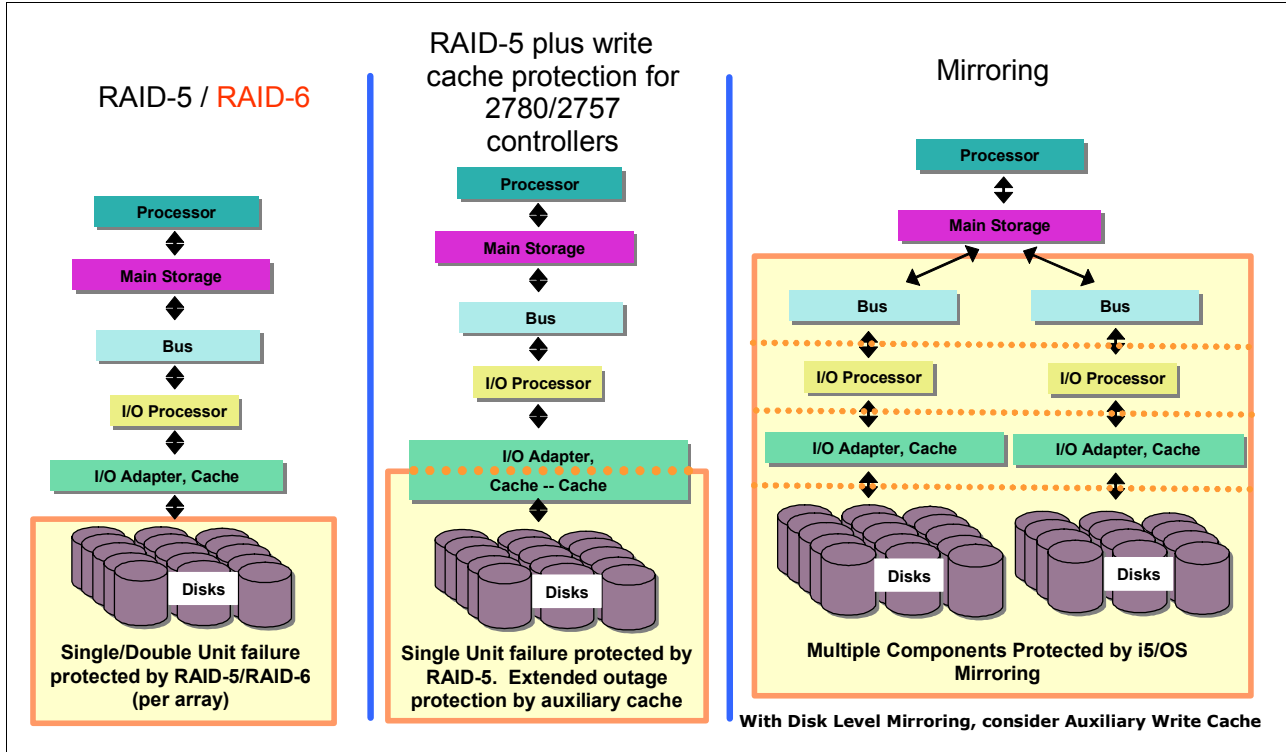


Figure 2-20 DASD protection schemes diagram with #2757, #2780 disk IOAs

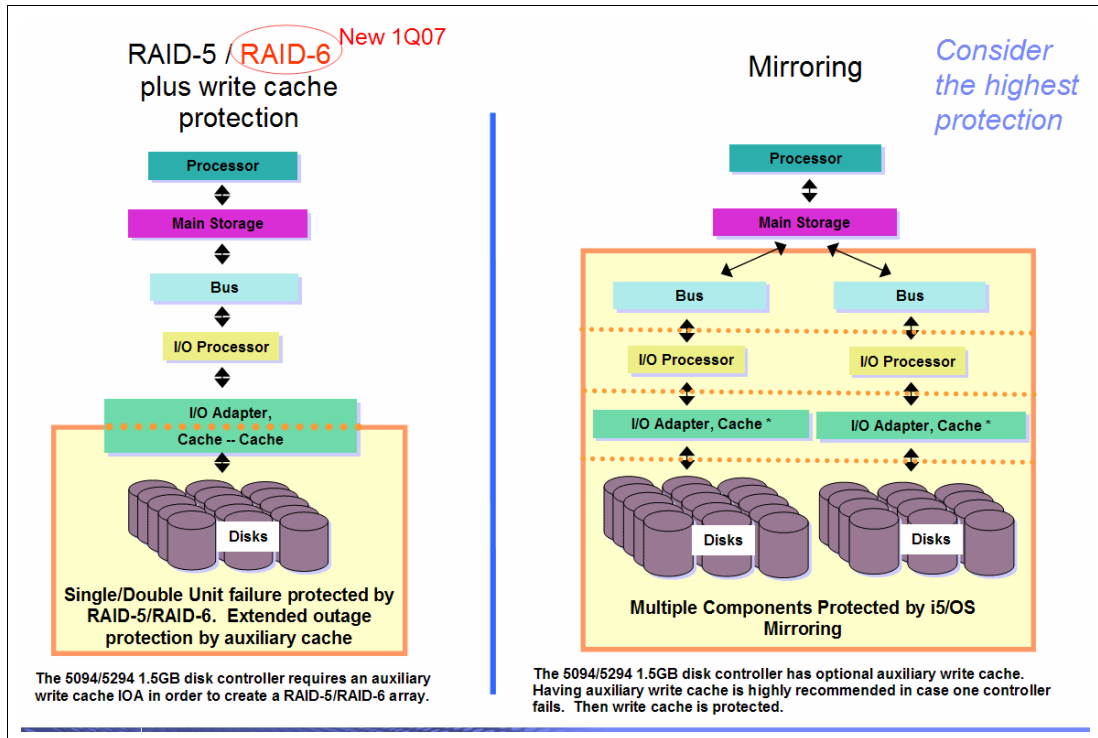


Figure 2-21 DASD protection schemes diagram with #5738, #5777, #5739, #5778 IOAs

Auxiliary write cache overview

The auxiliary write cache option, introduced during 2005, applies to the #2757, #2780, #5738, #5777, #5739, and #5778 IOAs and is supported on i5/OS V5R2 (#2757, #2780 only), V5R3, and V5R4. It provides a way to have redundant write cache for the disk controller. This cache does not protect against the possibility of an adapter failing, but it does protect against the contents of the adapter being lost if the disk I/O adapter fails.

Though a disk controller failing will very likely take the i5/OS partition it supports down (unless the failing controller is part of a user dependent or independent auxiliary storage pool (ASP)), restarting the affected partition is fairly straightforward if the write cache contents are undamaged and available. If the primary write cache contents are lost, then with the large cache sizes available today, it can take significantly longer to get the i5/OS partition back up and productive.

A duplicate (*auxiliary or secondary*) write cache helps protect against the extended outage scenario. You need an authorized System i hardware service professional to recover the data in the auxiliary write cache when that person replaces the failed disk controller or just its (primary) write cache.

Figure 2-22 shows a logical representation of how the auxiliary write cache plugs into the disk controller with the built-in primary write cache and lists important software levels and disk controller and auxiliary write cache hardware feature numbers.

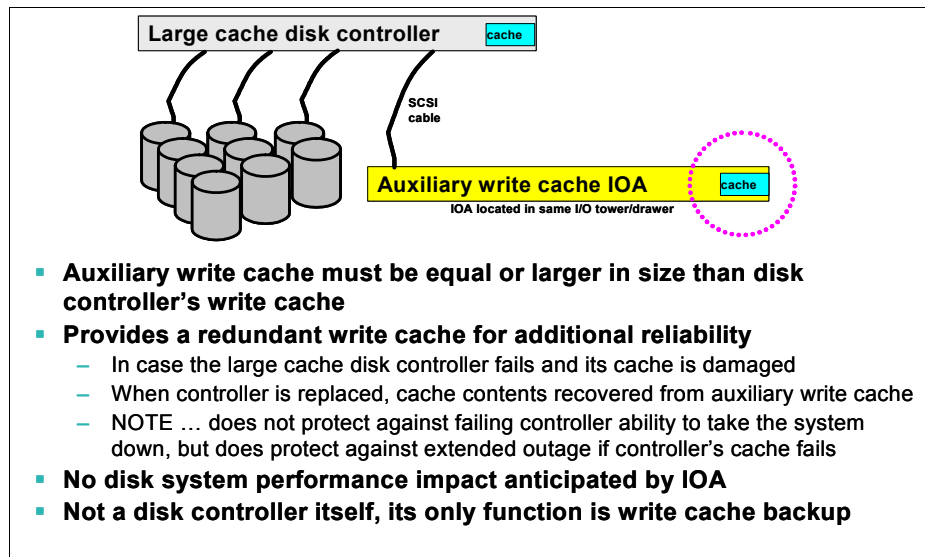


Figure 2-22 Auxiliary write cache overview

The physical dimensions of the auxiliary write cache card are the same as the disk IOA card itself (the #5739/#5778 IOA specifically for the EXP24 disk enclosure is a double-wide card with the auxiliary cache built in). This is true except for auxiliary write cache features #5590 and #5591 for the #2780 and #2757 IOAs, as listed below.

Where system uptime is heavily required and mirroring at an IOP or an IOA (controller) level is not possible, you should order an auxiliary write cache card.

As you can see from the figure, for the disk IOAs that have the optional auxiliary write cache the IOA cache is attached via a SCSI cable connected to the slot where four additional disk drives could be attached.

Remember that there are MES options to upgrade a current configuration to the available auxiliary cache card IOAs, as well as upgrading the #2757/#2780 disk IOAs to the newer, higher I/O capacity disk IOAs.

The auxiliary write cache IOA features include:

- ▶ #5580 (757 MB) for the #2780 IOA
- ▶ #5581 (757 MB) for the #2757 IOA
- ▶ #5590 (1.5 GB) for the #2780 IOA
- ▶ #5591 (1.5 GB) for the #2757 IOA
- ▶ #5582 (1.5 GB) for the #5738 IOA
- ▶ #5583 (1.5 GB) for the #5777 IOA

As stated earlier in this book, there is no performance gain in attaching the larger auxiliary write cache to the #2757/#2780.

Supported System i5 models include 270, 520, 550, 570, 595, 800, 810, 820, 825, 830, 840, 870, and 890 models. Card slot and device connectivity requirements include:

- ▶ Two PCI Slots are required.
 - Slots must be in the same I/O drawer/tower or 8xx processor enclosure. Adjacent slots are not required.
 - #2757/#2780/#5738 IOAs must attach to an IOP. Using the same IOP is not required unless associated with the load source disk controller.
 - The SCSI cable attaches to the fourth port in the IOAs.
- ▶ The cable to the auxiliary write cache card takes the place of up to five disk drives.

Important: You must review the *PCI and PCI-X Placement Rules for IBM System i models with i5/OS V5R3 and V5R4*, REDP-4011-03 or later, redpaper to ensure proper planning for card placement and specific I/O tower or enclosure support.

Watch for a planned redpaper, RED-3919, that will include additional details on the new disk controllers announced February 2007.

2.15.1 RAID level 6

RAID level 6 uses two independent parity computations and algorithms to provide protection against failure of two disks in the array. It essentially writes two *parity data stripes* versus one parity data stripe used with RAID5. Since this is much better than RAID5 in terms of protection, this is a very good solution for mission-critical applications when mirroring cannot be used.

RAID Level 6 or RAID6 requires a minimum of four drives.

These two parity stripes, as shown in Figure 2-23, are what provide the protection with two disk failures in a RAID6 array set of disks versus the single disk failure protection with RAID5.

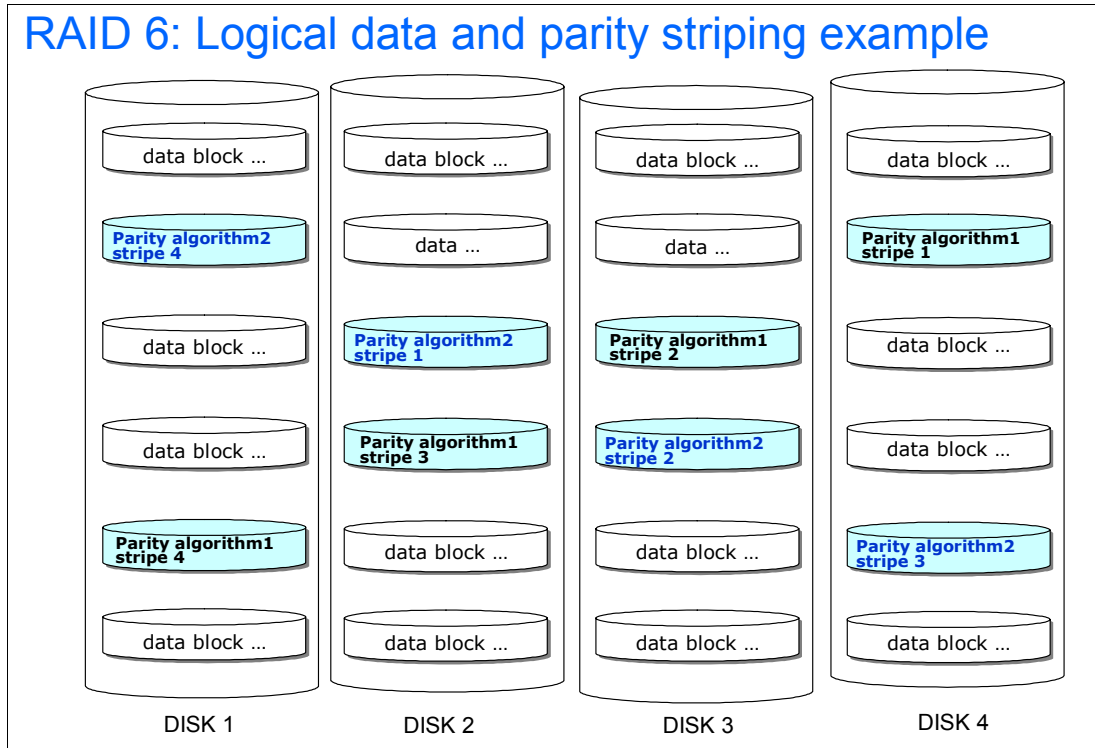


Figure 2-23 RAID6 logical and parity striping example

While offering a better DASD protection, RAID6 has some performance considerations:

- ▶ Controller overhead to compute parity addresses.
- ▶ Heavy write to disk application environments may see slightly less throughput when compared to RAID5 with same technology disk hardware.

RAID5 and RAID6 comparison

Figure 2-24 compares RAID5 and RAID6. The Performance Capabilities Reference guide can be found on the Web at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzahx/sc410607.pdf>

	RAID-5	RAID-6
Reserved Capacity for striping	Capacity of 1 DASD per RAID array	Capacity of 2 DASD per RAID array
Minimum DASD/array	3	4
Maximum DASD/array	18	18
RAID Controllers	Many existing controllers + #5727,#5728	#5737 #0648 #5776
Required i5/OS	All supported	V5R3 or later
Performance impact	Known	Similar to 5703 running RAID-5*

* See Performance Capabilities Reference Guide

Figure 2-24 Comparison of RAID5 and RAID6

2.16 SAN and TotalStorage DS6000 and DS8000

Storage Area Network (SAN) switches are supported on the System i models for connecting multiple host servers with storage servers and devices, creating a network of storage. There are a variety of SAN switches supported by the System i5, including:

- ▶ IBM TotalStorage SAN Switches 2005-H08/-H16, 3534-F08, and 2109-F16/-F32 with 4-32 ports and 2 Gbps fabric switching
- ▶ IBM TotalStorage SAN Switches SAN12M-1 2026-E12 and SAN24M-1 f2026-224 2 Gbps fabric switches with FlexPort scalability from four to 24 ports
- ▶ Cisco MDS 9216 Multilayer Fabric Switch 2062-D01 and 9120/9140 2061-020/-040, with 16 to 48 ports and 2 Gbps performance
- ▶ IBM TotalStorage SAN Switch SAN32M-1 2027-232 and Director SAN140M 2027-140 with 8 and 140 ports and 2 Gbps performance
- ▶ CNT UltraNet Multi-service Director (UMD) 2042-N16, CNT FC/9000 Fibre Channel Director 2042-256/128/001/
- ▶ Cisco MDS Multilayer Director 9509 2062-D07 and 9506 2062-D04 which support 1 and 2 Gbps Fibre Channel and up to 128 and 224 fibre channel ports
- ▶ IBM TotalStorage SAN Director 2109-M14/M12 and Cabinet 2109-C36 with up to 128 ports 2 Gbps performance

These external storage servers can attach via the following System i5 Fibre Channel support hardware:

- ▶ Remote Load Source boot from an external disk device: This requires the 2847 IOP with either a single 2766, 2787, or 5760 Fibre Channel disk adapter (controller) attached, as illustrated in Figure 2-25. This configuration supports a Load Source disk unit and up to 31 additional LUNs.




Figure 2-25 Remote Load Source implementation

This IOP cannot be used as an alternate IPL device for booting from any other devices such as a DVD-ROM, CD-ROM, or any integrated internal Load Source. Also, the 2847 IOP cannot be used as a substitute for a 2843 or 2844 IOP to drive non-Fibre Channel storage, LAN, or any other iSeries adapters.

Extending SAN's connectivity, the TotalStorage DS6000 and DS8000 have been supported in the System i family since V5R2. This support requires an internal boot device or Load Source unless you are in V5R3M5 SLIC or later where you can use a #2847 IOP Fibre Channel i5/OS Load Source feature. It supports features #2766, 2787, and very recently the 5760 PCI-X Fibre Channel Disk Controllers.

- ▶ Non-remote Load Source external disk connections: The IBM SystemStorage disk subsystems attach via the 2844 IOP and either 2766, 2787, or 5760 PCI-X Fibre Channel Disk Controller. A single 2844 IOP can attach only one 2766, 2787, or 5760 controller. For I/O-intensive processing that requires high performance, use the System i5 integrated storage or the high-end DS8000 product (ESS2107).

Base DS6000 and DS8000 remote Load Source support considerations include:

- ▶ I5/OS V5R4M0
- ▶ PTFs for i5/OS V5R3 (or PTF supersedes)
 - MF33328, MF33845, MF33437, MF33303
 - SI14690, SI14755, SI14550
- ▶ *PTFs for OS/400 V5R2 (or PTF supersedes)
 - MF33327, MF33301, MF33469, MF33302
 - SI14711, SI14754

2.16.1 Flashcopy

FlashCopy® is a point-in-time function of the additional charge Copy Services available with the IBM Enterprise Storage Subsystem (ESS) servers, including the DS6000 and DS8000. There is a specific iSeries Copy Services Toolkit that has unique extensions to the general FlashCopy capabilities available with the DS6000 and DS8000 servers. Independent Auxiliary Storage Pools (IASPs) are the focus of these iSeries extensions.

The point-in-time copy is typically used where you need a copy of production data to be produced with minimal application downtime while accessing the source volumes of data.

FlashCopy copies source volumes to a target volume within the same ESS Logical Subsystem (LSS). FlashCopy has COPY and NOCOPY options. The copy is performed in two stages, that is, logical completion and then physical completion. The logical completion occurs regardless of the COPY/NOCOPY option selected. This frequently takes less than five seconds, as it creates a bit map representing the source volume rather than actually copying the data.

At logical completion, the data has been cloned, the source volume access can be resumed, and both the source and target volumes are available for read and write access. At this point, the contents of the source and target volume are exact duplicates. You can back up, recover to, and resume processing from this point.

The physical completion depends on what you want to do with the target volumes. If you specified COPY, then a background task copies the source volume data to the target volume. This must then complete for physical completion to occur. While this background copy is taking place, the source volume can be accessed and updated. Any source volume updates are mapped to the bit map and factored into the copy process.

When the physical copy has completed, the relationship between source and target volumes is forgotten.

Figure 2-26 shows FlashCopy with i5/OS and iASP.

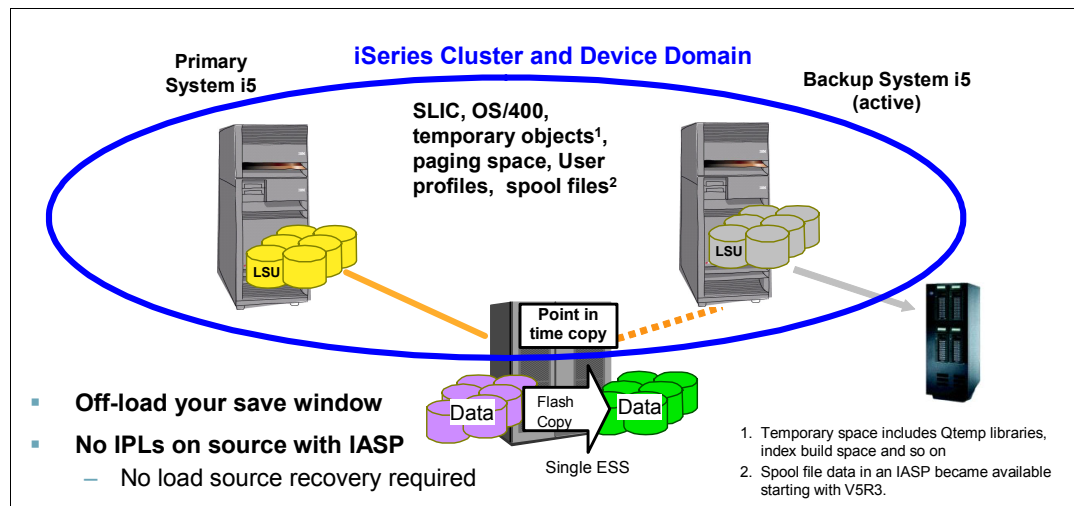


Figure 2-26 Flashcopy with i5/OS and iASP

If you only want to use the target volume data as input to backup, choose the NOCOPY option. Then no background copy is performed. When a target volume is needed only for a short time (such as for input to backup), copying tracks of unchanged data from the source volume to the target volume is a waste of resources. The second system can process (via pointers) data on the corresponding source volume data.

This sounds good but causes additional disk I/Os on the source system. Depending on the current source system activity, this may impact performance. This also running FlashCopy has planning considerations:

- ▶ Must end all active access on the IASP (vary off IASP).
- ▶ Customer applications must be in an IASP environment to use it.

- ▶ Requires iSeries Copy Services Toolkit for IASP.
- ▶ Disk sizing for system ASP is very important. This requires the fastest disk on the system, as this is where memory paging, index builds, and so forth, happen.
- ▶ IPL the backup system after the save has taken place.
- ▶ Separate feature code IOP/IOA for each IASP on the target is required.

FlashCopy, as described here, can only be done within the same external disk subsystem. Therefore it is not suitable for Disaster Recovery. However, using FlashCopy with iASP has many benefits to the System i5 environment. Here are some of the advantages:

- ▶ No IPL to attach second copy of data to the backup server (IPL at the end of process).
- ▶ Saves can start approximately 5 to 10 minutes after the flash.
- ▶ No special System i SLIC microcode or i5/OS issues. The second system has its own copy of this software.
- ▶ Saves will be done through i5/OS. This enables taking advantage of performance using i5/OS save techniques including parallel saving and BRMS in a network.
- ▶ Replication on the production server (in cooperation with High Availability Business Partner (HABP) solution) to another system can continue to run without getting behind while the save is running.
- ▶ Multiple IASPs can be attached and saved in a single process. This is a huge benefit compared to the current remote Load Source mirroring technique.

2.16.2 Peer-to-Peer Remote Copy

Peer-to-Peer Remote Copy (PPRC) is the acronym representing copying data from volumes on one External Storage Server (ESS #1) to volumes on a second External Storage Server (ESS #2). This is now a generic term that applies to IBM TotalStorage server-to-server capacities that have several classes of functions.

Metro Mirror (synchronous PPRC)

Metro Mirror is a remote data mirroring technique for all supported servers, including z/OS® and open systems. It is designed to constantly maintain an up-to-date copy of the local application data at a remote site, which is within the metropolitan area (typically up to 300 km away using DWDM). With synchronous mirroring techniques, data currency is maintained between sites, though the distance can have an impact on performance. Metro Mirror is used primarily as part of a business continuance solution for protecting data against disk storage system loss or complete site failure.

Because the copying function occurs within the disk subsystem level, applications have no knowledge of its activity.

IBM TotalStorage Global Copy (PPRC Extended Distance, PPRC-XD)

Global Copy is an asynchronous remote copy function for z/OS and open systems for longer distances than are possible with Metro Mirror. It operates over high-speed Fibre Channel communications links and is designed to maintain a complete and consistent remote mirror of data asynchronously at virtually unlimited distances with almost no application response time delays.

This function is appropriate for remote data migration, off-site backups, and transmission of inactive database logs at virtually unlimited distances.

PPRC with iASP

PPRC with iSeries iASP is a good high-availability and disaster recovery solution and has a number of benefits over a remote Load Source solution. However, here are some important planning considerations:

- ▶ Applications must be in an IASP environment to use it.
- ▶ An IASP Copy Services Toolkit is required, an optional chargeable service offering.
- ▶ System objects are not replicated. User profiles/system values need to be replicated on a secondary/target system.
- ▶ Disk sizing for system ASP is very important. This requires the fastest disk on the system, as this is where memory paging, index builds, and so on happen.
- ▶ Separate FC IOP/IOA for each IASP on target server.

Here are the benefits of using PPRC with System i5 IASP:

- ▶ Only replicates database/application changes. System and temporary I/O are not replicated. This can provide a huge performance improvement over a hardware solution that replicates the system ASP.
- ▶ Data can be made available to the DR system in as little as 10 minutes (may be longer in case of abnormal failure on production as database recovery required on DR system).
- ▶ Could be combined with FlashCopy with the saves taken from the remote site.

2.17 10 GB Ethernet LAN IOA

The #5721 and #5722 10 GB Ethernet LAN IOAs shown in Figure 2-27 are high-performance cards requiring a significant portion of the bus capacity to run at a good speed. With that requirement rules are established to help ensure good performance.

These cards do not run at slower speeds (do not negotiate down) and cannot be connected to mere 10/100/1000 Ethernet networks.




Figure 2-27 10GB Ethernet LAN IOAs

For performance considerations, the following placement rules have been implemented:

- ▶ Supported only in 64-bit PCI-X slots. They are not supported in PCI slots or in #5074/5079 I/O towers.
- ▶ Dual mode IOAs (smart IOAs) only, which can run with or without an IOP.
- ▶ Needs i5/OS V5R4 on POWER5 and higher.

- ▶ Limit of six cards per HSL-2 loop or limit of three cards per HSL-1 loop.
- ▶ For 5094/5294/5088/0588 only a maximum of two cards, and they must be installed in the highest speed slots (C5 or C9).
- ▶ For 5095/0595 only a maximum of four cards, and they must be spread out between the two EADS.
- ▶ For 5790 a maximum of four cards only, and they must be spread out between the two EADS.
- ▶ For 520 - 1.9 GHz, only a maximum of three cards requiring one per EADS and then one in C04.
- ▶ For 520 - 1.5/1.65 GHz, only a maximum of two cards.
- ▶ For 550, maximum of two cards.
- ▶ For 570, maximum of two cards per processor drawer.

Note: Using all of these cards with significant bandwidth being used concurrently, response time would be greatly impacted. The maximum card statement allows additional bandwidth to cover potential card failures.

2.18 IBM Configurator for ebusiness

The IBM Configurator for e-business or eConfig is used to configure and upgrade IBM eServer logo systems and subsystems. It is a tool with which to configure multiple product lines. Many IBM customers have been using this tool as a means of creating their systems and eventually seeing the cost.

Very recently eConfig has been enhanced to support IOA-level mirroring and the new Custom Data Protection Specify, #0269 option for 520/550/570/595.

2.18.1 IOA-level mirroring #0308

In the past most customers who wanted to mirror disk controllers also wanted to mirror the IOP. Mirroring at the IOP level has been available for years, but with the introduction of the dual mode IOAs (smart IOAs), which can run with or without an IOP, IBM has received requests for eConfig to support this specific mirroring level. Figure 2-28 shows the IOA-level mirroring configuration.

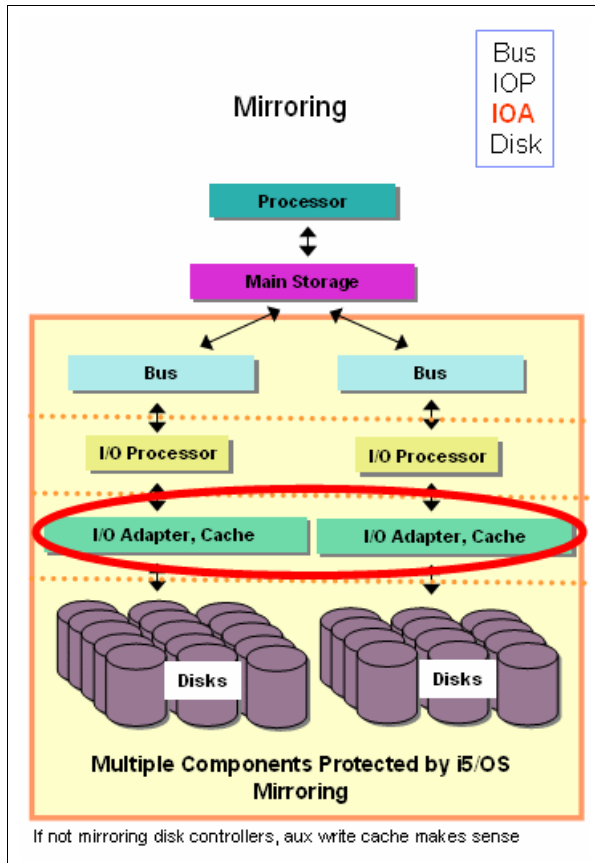


Figure 2-28 IOA-level mirroring

2.18.2 Custom Data Protection Specify #0269

This is a great enhancement to eConfig as more customers are doing partitioning and having each i5/OS partition run different business objectives. Running partitions with different business objectives, you have the choice to implement a different level of data protection depending on the need, and of course the cost to implement. Previously only a single total-system specify was understood by eConfig. This enhancement allows the eConfig tool to better support more complex, multi-partition configurations like the one shown in Figure 2-29.




Figure 2-29 Partitions having different levels of data protection

2.18.3 5372-IS5 Partition Preload Specifies

This enhancement makes it possible to specify the configuration of each partition. Instead of just specifying one hardware that includes all the necessary specifications of a multi-partitioned system, this 5372-IS5 Partition Preload Specifies makes it easy for you to design the partitions individually (like choosing what data protection should be implemented on a target partition). Figure 2-30 shows how this feature works.

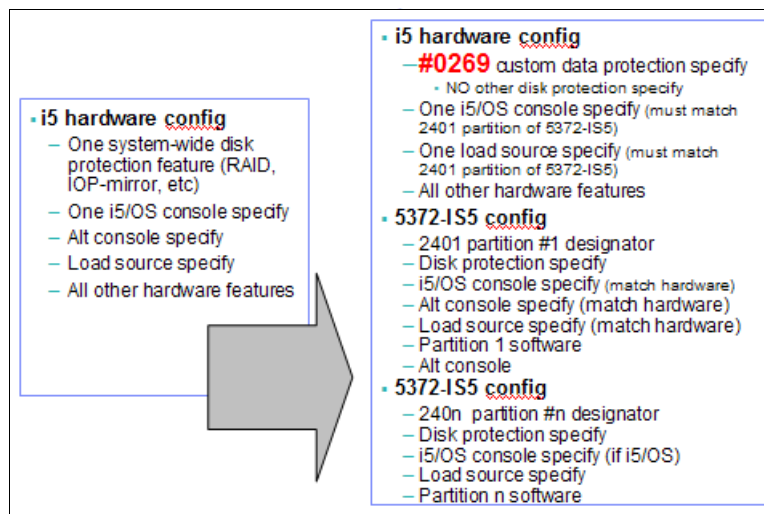


Figure 2-30 Partition Preload Specifies example

2.19 Customer Specified Placement and LPAR preload

IBM Business Partners and customers can take advantage of the Customer Specified Placement (CSP) and LPAR preload enhancements. The clear benefit to both IBM Business Partners and customers is that they can save on installation and setup time. Even more significant is the savings of having the System i5 server delivered with multiple partitions

using the Custom Data Protection and software preload rather than doing the partitioning and data protection configuration themselves.

Enhancements to CSP and LPAR preload:

- ▶ Model 520 and 550 — new box orders are now part of Customer Specified Placement (CSP):
 - New 550 chargeable feature #0453
 - New 520 chargeable feature #0456
 - Joins existing #8453 (no charge) feature for 570 and 595
 - Saves customer time and money as well as improving quality
 - Reliable, as it is done properly
- ▶ Enhancements to Software Preload

In addition to i5/OS or AIX 5L, you can now specify to load any IBM licensed program that is preloaded by IBM Manufacturing (by partition)
- ▶ Support of custom data protection in configurations

Data protection by partition delivered from the factory using #0269 with #0453, #0454, and #0455 specifies

Custom manufacturing option summary

To summarize:

- ▶ Customer Specified Placement (CSP)
 - Based on System Plan Tool (SPT) output (or previously available LPAR Validation Tool (LVT) output), place cards/disk drives in requested physical location.
 - One feature per system (new box, not MES order).
 - #8453 for models 570/595.
 - No charge from IBM Manufacturing.
 - #0453 for model 550 per box (chargeable fee required).
 - #0456 for model 520 per box (chargeable fee required).
- ▶ #0454 LPAR Partitioning Initialization models 595, 570, 550, and 520
 - Configures partition and assigns correct resources based on LVT output.
 - Chargeable. One \$300* feature per partition initialized.
- ▶ #0455 LPAR OS Preload models 595, 570, 550, and 520
 - Load i5/OS or AIX 5L into partition created by #0454.
 - Load any program products normally preloaded by IBM Manufacturing.
 - Data protection by partition, preloaded by IBM Manufacturing.
 - Chargeable. One \$600* feature per partition loaded.

For more information about CSP see the IBM Web site at:

<http://www.ibm.com/eserver/power/csp/>

2.20 System i5 editions

System i5 editions are specifically configured software and hardware configurations that address the constantly changing and often unpredictable business requirements. They are

designed with the power and capacity to run core business applications as well as new e-business applications simultaneously on a single machine.

System i5 editions are flexible packages of options that help reduce the complexity of choosing the best server to match your business need. Each System i edition combines a set of software licensing and hardware features to help meet the particular demands of any enterprise.

For additional information about System i editions, refer to redpaper *IBM System i Editions: Buying and Selling Guide*, REDP-3916. This redpaper is updated periodically.

2.20.1 Editions at a glance

The System i5 comes with several different editions designed to meet the requirements of any enterprise. Each of these editions is unique in such a way that each one targets to match a specific business IT requirement.

The following System i5 editions differ in the following areas:

- ▶ Value Edition: Entry-level offering and *a la carte* choices that come with fixed amounts of 5250 CPW.
- ▶ Express Configurations: Attractively priced entry-level offering with seven fixed configurations offered. Each comes with a fixed amount of 5250 CPW.
- ▶ Standard Edition: Tailored for applications that do not require 5250. This edition is upgradeable to the Enterprise Edition.
- ▶ Solution Edition: Based on the Enterprise Edition, but more on running a specific ISV solution. This edition requires the purchase of a qualified ISV solution.
- ▶ Enterprise Edition: Basically includes all of the packaged options, as well as datacenter management tools, additional e-business tools, and an IBM services voucher to help speed up implementation.
- ▶ High Availability (HA) Edition: Contains the software required to support the System i5 HA hardware environment. The base model hardware, the Capacity on Demand capabilities (both temporary and permanent), and the upgrade paths are all identical to the Enterprise Edition.
- ▶ Capacity BackUp (CBU) Editions: These are similar to HA Editions but available in specific primary system and available backup system configurations.
- ▶ Domino/Collaboration Edition: Combines excellent price/performance targeted for Lotus workloads with the reliability, manageability, and low cost of ownership that have made the System i family so successful as a Domino server platform.

Figure 2-31 shows a summary of the System i5 editions and the server models that support each edition.

	9405		9406						
	Express	Value	Standard	Enterprise	Solution	Domino + Collaboration	High Availability	Capacity BackUp	Telephony Express ¹
520	x	x	x	x	x	xx	x	X ¹	x
550			x	x	x	x	x	x	
570			x	x			x	x	
595			x	x			x	x	

¹ October 2006. Detail information not yet integrated.

Express Configurations

- Attractively priced entry offering
- Seven specific packages

Value Edition

- Entry "a la carte" flexible options

Standard Edition

- Targeted for applications that do not require 5250 OLTP
- Upgradeable to Enterprise

Enterprise Edition

- Most flexible option for dynamic business environments
- Full 5250 OLTP capability
- Includes additional software & service/education vouchers

Solution/Domino Edition

- Application targeted editions with specific application software prerequisites

High Availability Edition

- Attractive pricing for 2nd system in business continuity environment
- Specific hardware and application prerequisites
- Full 5250 OLTP

Capacity BackUp Edition

- For business continuity & emergency/disaster scenarios
- Full or zero 5250 OLTP

Figure 2-31 System i5 Edition summary

During July and October 2006 new Lotus Domino, Workplace, and other product offerings and expanded Capacity Backup edition offerings were announced.

For the latest edition information refer to the following Web site:

<http://www-03.ibm.com/systems/i/editions/>

Similar to the edition summary table in Figure 2-31 on page 58, the table in Figure 2-32 shows the different models, their processors, and other hardware characteristics for the Standard and Enterprise Editions and Value and Express Editions.

Model	N-Way {1}	Server Feature	Edition Feature	Processor Feature	Processor CPW {2}	5250 CPW	SW Tier	Base i5/OS Licenses	Base EE Features {3}	Accelerator for System i5	Processor Activation Feature	Max HSL loop / I/O Drawer	Max IXS / IXA / iSCSI
520	1	0975 0970	7350 (Value) 7140-7142 (Express)	8325 1.9GHz L2 cache	600/3100 600/3100	30 30	P05	1	NA	7355 7680-7682	NA	1 / 6	18 / 8 / 21
520	1	0975 0970 0970	7352 (Value) 7143 & 7148 (Express) 7144 & 7152 (Express)	8327 1.9GHz L2 & L3 cache	1200/3800 1200/3800 3800	60 60 60	P10	1	NA	7357 7354 NA	NA	1 / 6	18 / 8 / 21
520	1	0906	7734 (Enterprise)	8327 1.9GHz L2 & L3 cache	1200	1200	P10	1	NA	NA	NA	1 / 6	18 / 8 / 21
	1	0906	7735 (Enterprise)	8327 1.9GHz L2 & L3 cache	2800	2800	P10	1	NA	NA	NA	1 / 6	18 / 8 / 21
520	1	0906	7784 (Standard)	8327 1.9GHz L2 & L3 cache	3800	0	P10	1	NA	NA	NA	1 / 6	18 / 8 / 21
520	1/2	0906	7785 (Standard) 7736 (Enterprise)	8330 1.9GHz L2 & L3 cache	3800-7100	0 (Std Ed) Up to max	P20	1	1	NA	Base 8410 Opt 7320	1 / 6	18 / 8 / 21
550	1/4	0910	7154 (Standard) 7155 (Enterprise)	8312 1.9GHz L2 & L3 cache	3800-14000	0 (Std Ed) Up to max	P20	1	1	NA	Base 8413 Opt 7323	2 / 12	36 / 16 / 42
570	2/4	0934	7757 (Standard) 7747 (Enterprise)	8338 2.2GHz L2 & L3 cache	8400-16000	0 (Std Ed) Up to max	P30	1	1	NA	Base 7738 Opt 7618	2 / 12	36 / 16 / 42
570	4/8	0935	7758 (Standard) 7748 (Enterprise)	8338 2.2GHz L2 & L3 cache	16700-31100	0 (Std Ed) Up to max	P40	1	1	NA	Base 7738 Opt 7618	4 / 24	48 / 32 / 84
570	8/16	0936	7759 (Standard) 7749 (Enterprise)	8338 2.2GHz L2 & L3 cache	31100-58500	0 (Std Ed) Up to max	P40	1	1	NA	Base 7738 Opt 7618	8 / 48	48 / 57 / 126
595 {4}	8/16	0940	7480/7900 (Standard) 7481/7901 (Enterprise)	8966/8973 1.9GHz L2 & L3 cache	26700-50500	0 (Std Ed) Up to max	P50	4	4	NA	Base 8457 Opt 7815	7 / 36	60 / 48 / 168
595 {4}	16/32	0942	7482/7903 (Standard) 7483/7904 (Enterprise)	8966/8973 1.9GHz L2 & L3 cache	51000-92000	0 (Std Ed) Up to max	P50	4	4	NA	Base 8457 Opt 7815	15 / 72	60 / 57 / 168
595 {4}	32/64	0943	7486/7906 (Standard) 7487/7907 (Enterprise)	8966/8973 1.9GHz L2 & L3 cache	92000-184000	0 (Std Ed) Up to max	P60	4	4	NA	Base 8457 Opt 7815	31 / 96	60 / 57 / 168

Figure 2-32 Standard and Enterprise Edition feature summary

2.20.2 IBM System i5 520 Collaboration Edition


The Collaboration Edition is a powerful, flexible, and cost-effective alternative to running Lotus Domino, IBM Workplace, or WebSphere Portal solutions and applications on Windows servers utilizing the flexibility and security of i5/OS.

The Collaboration Edition consists of 1-way 520 with the minimum memory and disk needed for performance to run collaboration applications. In this edition, you will get the following System i5 model 520 configuration:

- ▶ 1-way POWER5+ 1.9 GHz with L3 cache
- ▶ 8200 MCU (Mail and Calendar Users Benchmark)
- ▶ One i5/OS License
- ▶ 4-32 GB Memory (minimum of 4 GB)
- ▶ 4-278 Disk Drives, 35 GB or more per disk (minimum of four drives)
- ▶ Either RAID or mirroring disk protection

Details to qualify for this edition are shown in Figure 2-33.

520 Collaboration Edition Pre-Requirements



Customers can qualify for the Collaboration Edition by either

- 1. Purchasing select new IBM Collaboration software licenses (see below)**
- 2. Transferring existing IBM Collaboration software licenses to the new Collaboration Edition 520**

11 Ways to qualify	1= Lotus CEO Lotus Communications - 70 Licenses
	2= Lotus Domino Collaboration Express - 125 Licenses
	3= Lotus Sametime - 350 User Licensee
	4= Lotus QuickPlace - 350 User Licenses
	5= IBM WebSphere Portal Express – 10 x 20 User Packs
	6= IBM Lotus Document Manager – 160 User Licenses
	7= IBM Workplace Services Express – 7 x 20 User Packs
	8= IBM Software Solution on Demand Workplace - 40 User licenses
	9= IBM WebSphere Portal Enable – 1 Processor Based License
	10= Workplace Collaboration Services – 1 Processor Based License
	11= Sametime Extranet – 1 Processor Based License

Figure 2-33 520 Collaboration Edition prerequisites

For more details on this Collaboration Edition for the System i5, see the following Web site:

<http://www.ibm.com/servers/eserver/iseriess/domino/edition.html>

2.20.3 IBM System i5 Solution editions

IBM offers specific System i5 520 and 550 model configurations that are part of a customized application solution. Through October 2006, the customized solution providers include:

- ▶ Clear Technologies
- ▶ SAP®
- ▶ Oracle® JD Edwards®

For the latest information see the System i5 Web site:

<http://www-03.ibm.com/systems/i/editions/solutionedition.html>

Figure 2-34 shows a summary of available Solutions Editions as of October 2006.

Customized Solution Editions at a glance						
Editions include:	Clear Technologies for C2CRM	SAP Solution Edition 1-way	mySAP ERP 2-way	mySAP ERP 4-way	Oracle JD Edwards EnterpriseOne	Oracle JD Edwards EnterpriseOne
System i5 Model	550	520	550	550	520	550
5250 CPW	Yes	No	No	No	No	Yes
Capacity on Demand	Yes	Yes	Yes	Yes	No	Yes
Active processors	1	1	2	4	1	1
Disk Pack Options	No	No	No	No	Yes (16 disk drives)	No
IBM i5/OS™	Yes (1 processor license)	Yes (1 processor licenses)	Yes (2 processor licenses)	Yes (4 processor licenses)	Yes (1 processor license)	Yes (1 processor license)
IBM Director for iSeries (5733-DR1)	Yes	Yes	Yes	Yes	Yes	Yes
IBM Web Enablement (5722-WE2)	Yes	Yes	Yes	Yes	Yes	Yes
Domino Utility Server Express	Yes (2 processor licenses)	No	No	No	No	No
IBM DB2® Query Manager and SQL Development Toolkit	No	No	No	No	No	Yes
Lotus Sametime	Yes (100 client licenses)	No	No	No	No	No
IBM WebSphere Development Studio	No	No	No	No	No	Yes (1 server license)
Performance Tools	No	No	No	No	No	Yes (1 server license)
IBM Advanced Function Printing	No	No	No	No	No	Yes (1 server license)
Customer care						
Education vouchers	No	No	No	No	No	No
Service vouchers	No	No	No	No	No	No

Figure 2-34 Available Solutions Editions as of October 2006

2.20.4 Capacity Backup (CBU) Editions

Capacity Backup editions offered through October 2006 for a primary system should be considered by customers:

- ▶ Wanting to get started with Disaster Recovery (DR), then deploy to a High Availability (HA) environment
- ▶ Considering starting with low-cost DR, then expanding to a full HA solution
- ▶ Using third-party disaster recovery services
- ▶ Whose cost considerations are oriented toward in-sourcing HA/DR
- ▶ Are considering an HA project
- ▶ Who have a wide set of business objectives and are considering a broader set solution options
- ▶ Considering using 8xx systems for backup

- ▶ Considering cost-effective alternatives for upgrades from an 8xx configuration
- ▶ Are currently using the Standard Edition and considering a Capacity Backup configuration

As of August 2006, we announced a full set of 550, 570, and 595 CBU editions. During October 2006 we announced a 520 CBU edition.

For information beyond what is covered here refer to the following Web site:

<http://www.ibm.com/systems/i/hardware/cbu>

The supported Capacity Backup edition options available were significantly increased with August 2006 announcements. The following figures show the Capacity Backup models available for specific IBM System i5 primary system models and for specific IBM System i 8xx primary system models. See Appendix C, “System i5 Capacity Backup edition software-based details” on page 509, for more details on the differences between the basic CBU Edition and the CBU Standard Edition.

Figure 2-35 shows what CBU Editions a customer is eligible to order given their 5xx primary system.

Primary system	Enterprise Edition FC	Standard Edition FC	Eligible n-way CBU	CBU Edition FC (with 5250 OLTP)	CBU Standard Edition* FC
550 1/4-way 1.65GHz (3300-12000 CPW)	7463	7462	550 1/4w*	7921	7920
550 1/4-way 1.9GHz (3800-14000 CPW)	7155	7154	550 1/4w*	7921	7920
570 1/2-way 1.65GHz (3300-6000 CPW)	7491 / 7489	7490 / 7488	570 1/4w*	7915	7914
570 2/4-way 1.65GHz (6350-12000 CPW)	7495 / 7470	7494 / 7469	570 1/4w*	7915	7914
570 2/4-way 2.2GHz (8400-16000 CPW)	7747	7757	570 1/4w*	7915	7914
570 4/8-way 2.2GHz (16700-31100 CPW)	7748	7758	570 1/4w* & 1/8w*	7915, 7917	7914, 7916
570 5/8-way 1.65GHz (15200-23500 CPW)	7472	7471	570 1/4w* & 1/8w*	7915, 7917	7914, 7916
570 8/16-way 2.2GHz (31100-58500 CPW)	7749	7759	570 1/4w* & 1/8w* & 2/16w	7915, 7917, 7760	7914, 7916, 7918
570 9/12-way 1.65GHz (25500-33400 CPW)	7474	7473	570 1/4w* & 1/8w* & 2/16w	7915, 7917, 7760	7914, 7916, 7918
570 13/16-way 1.65GHz (36300-44700 CPW)	7476	7475	570 1/4w* & 1/8w* & 2/16w	7915, 7917, 7760	7914, 7916, 7918
595 8/16-way 1.65GHz (24500-45500 CPW)	7497	7496	595 2/16w*	7911	7910
595 8/16-way 1.9GHz (26700-50500 CPW)	7481 / 7901	7480 / 7900	595 2/16w*	7911	7910
595 16/32-way 1.65GHz (46000-85000 CPW)	7499	7498	595 2/16w* & 4/32w	7911, 7909	7910, 7912
595 16/32-way 1.9GHz (51000-92000 CPW)	7483 / 7904	7482 / 7903	595 2/16w* & 4/32w	7911, 7909	7910, 7912
595 32/64-way 1.65GHz (86000-165000 CPW)	7985	7984	595 2/16w* & 4/32w	7911, 7909	7910, 7912
595 32/64-way 1.9GHz (92000-184000 CPW)	7487 / 7907	7986 / 7906	595 2/16w* & 4/32w	7911, 7909	7910, 7912

Figure 2-35 IBM System i5 primary systems and available secondary CBU edition models

Figure 2-36 is a draft showing 5250 CBU editions announced October 2006.

Primary System	Eligible CBU System
810 1w (750 CPW) Enterprise or Standard	520 1w (1200) Enterprise, (3800) Standard
810 1w (1020) Enterprise or Standard	520 1w (1200) Enterprise, (3800) Standard
810 1w (1470) Enterprise or Standard	520 1w (1200) Enterprise, (3800) Standard
810 2w (2700) Enterprise or Standard	520 1w (1200 or 2800) Enterprise, (3800) Standard
810 2w (2700) Enterprise or Standard	520 1w (1200 or 2800) Enterprise, (3800) Standard
520 1w (1000 CPW) Enterprise or Standard	520 1w (1200) Enterprise, (3800) Standard
520 1w (2400) Enterprise or Standard	520 1w (1200 or 2800) Enterprise, (3800) Standard
520 1w (3300) Enterprise or Standard	520 1w (1200 or 2800) Enterprise, (3800) Standard
520 2w (6000) Enterprise or Standard	520 1w or 1/2w (1200 or 2800 or 3800-7100) Enterprise, (3800 or 3800-7100) Standard
520 1w (1200) Enterprise	520 1w (1200) Enterprise
520 1w (2800) Enterprise	520 1w (1200 or 2800) Enterprise
520 1w (3800) Standard	520 1w (3800) Standard
520 1/2w (3800-7100) Enterprise or Standard	520 1w or 1/2w (1200 or 2800 or 3800-7100) Enterprise, (3800 or 3800-7100) Standard

Figure 2-36 Draft: IBM System i5 primary systems and available secondary 520 CBU edition models

Figure 2-37 shows what IBM System i5 CBU Editions (does not include 520 CBU Editions announced October 2006) the customer is eligible to order given their 8xx primary system.

Primary system	Enterprise Edition FC	Standard Edition FC	Eligible n-way CBU	CBU Edition FC (with 5250 OLTP)	CBU Standard Edition FC
825 3/6-way POWER4 (3600-6600 CPW)	7418	7416	550 1/4w	7921	7920
870 5/8-way POWER4 (7700-11500 CPW)	7433	7431	570 1/4w	7915	7914
870 8/16-way POWER4 (11500-20000 CPW)	7421	7419	570 1/8w	7917	7916
890 16/24-way POWER4 (20000-29300 CPW)	7424	7422	570 1/8w	7917	7916
890 24/32-way POWER4 (20000-29300 CPW)	7427	7425	570 2/16w & 595 2/16w	7760, 7911	7918, 7910
890 16/24-way pre editions *	n/a	n/a	570 1/8w	7917	7916
890 24/32-way pre editions *	n/a	n/a	570 2/16w & 595 2/16w	7760, 7911	7918, 7910

Figure 2-37 IBM System i primary systems and available secondary CBU edition models

The asterisk (*) in this figure indicates that, as of August 2006, IBM believes that all pre-edition 890 systems have been upgraded to edition-based 890 systems. If this is determined to be untrue for a specific 8xx configuration and that customer with a pre-edition 890 wishes to acquire a CBU system, contact Mark Olson (at olsonm@us.ibm.com) or Steven Finnes (at finnes@us.ibm.com) of IBM Rochester for exception processing. Normal eligibility approval processes check against the edition feature code numbers and will fail the check if an edition feature code is not found.

There are important requirements inherent in taking advantage of these CBU edition offerings:

- ▶ You must register the CBU system and primary system. Register at:

<http://www.ibm.com/systems/i/hardware/cbu>

Since the CBU Edition has new, more flexible terms and conditions for transferring licenses and Enterprise Enablement features between systems, the customer must sign a new license agreement (electronically in most cases or, if necessary, by fax).

The order will only be scheduled for manufacture after the system has been registered and the signed license agreement has been transmitted to IBM.

- ▶ The primary system must be a matching model 550/570/595 or designated 825, 870, or 890.
- ▶ The primary system must be equal to or larger than the CBU Edition system.
- ▶ The primary system must match 5250 OLTP capability (CBU-to-Enterprise or CBU-Standard-to-Standard Edition).
- ▶ You must declare how many role swap processors are involved.

With these offerings you can temporarily transfer i5/OS and Enterprise Enablement processor entitlements for:

- ▶ Planned outage roll swaps
- ▶ Unplanned outage failover
- ▶ Testing
- ▶ A whole system or part of a system

For example, you can move i5/OS and EE for just one partition assuming associated i5/OS or 5250 CPW resources for that partition are not in use on the primary system.

For temporary activations it is important to note that:

- ▶ Processor activations cannot be temporarily transferred.
- ▶ Processor activations require either:
 - On/Off Capacity on Demand (which includes i5/OS licensing and Enterprise Enablement and does not require temporary entitlement transfers)
 - Permanent processor activations (for which temporary i5/OS and Enterprise Enablement entitlement transfers may be applied)

Important: Here is information about the mechanics of temporary software license transfer:

- ▶ Every System i machine comes with an i5/OS license. The customer also has a specific number of i5/OS processor license entitlements (commonly called i5/OS licenses or i5/OS processor licenses).
- ▶ A software key tells each system's license manager how many i5/OS processor license entitlements each system has. If the customer ever runs/assigns more i5/OS work than the system has entitlements, then frequent license entitlement problem messages are issued to the operator indicating that the system is out of license compliance. By the terms of the IBM licensing agreement, the customer must then bring the system into compliance by either reducing i5/OS workload or adding i5/OS processor license entitlements.

Signing the additional contract for CBU temporary transfers documents that IBM and the customer understand when entitlements can be transferred from the primary to the CBU system.

- ▶ When entitlements are temporarily transferred to the CBU system, the software key information is not altered on either machine. Thus, the CBU system will continue to issue frequent out-of-compliance messages indicating that i5/OS is out of compliance. The additional contract gives permission to the customer, under these specific situations, to ignore these messages. The customer agrees to monitor the i5/OS usage on both the primary and CBU systems to ensure that the i5/OS usage on the pair of machines does not exceed the total number of i5/OS processor license entitlements.

2.21 Upgrade paths from one edition to another edition

With announcements through August 2006, IBM has significantly expanded the System i5 model-based upgrade paths, giving more choices as well as flexibility to the customers who want to upgrade their system.

In this section we provide a set of generalized figures of possible upgrade options, organized according to specific hardware and software edition feature numbers. However, for complete information about supported upgrade paths, refer to the following sources of documentation:

- ▶ IBM System Hardware Information Center

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>

Search with *system i5 upgrade* to find a list of available documentation.

- ▶ The IBM System i5 support Web site. topic: Planning - Migration and upgrades

<http://www-03.ibm.com/servers/eserver/support/series/planning/migrationupgrade.html>

This page contains links that provide assistance to System i clients with planning, upgrading, and maintenance activities from system to system software upgrades and aids for hardware upgrade and migration.

Important: We strongly recommend that you use information from these Web sites for all up-front planning for migration to minimize problems later when performing the upgrade.

- ▶ IBM Redbook: *IBM eServer iSeries Migration: A Guide to Upgrades and Migrations to POWER Technology*, SG24-72000

This book was last updated in March 2005. It covers upgrading to IBM System i5 models as well as the then available software edition-to-edition upgrades. It does not include System i5 models or edition options announced during 2006. A 2006 update to this book is planned.

Important: Good sizing and capacity planning processes should be an inherent part of any configuration upgrade being considered.

Figure 2-38 shows generalized upgrade path options for 8xx to 5xx and 5xx to 5xx models, using Standard and Enterprise Editions and model 520 Value and Express Editions.

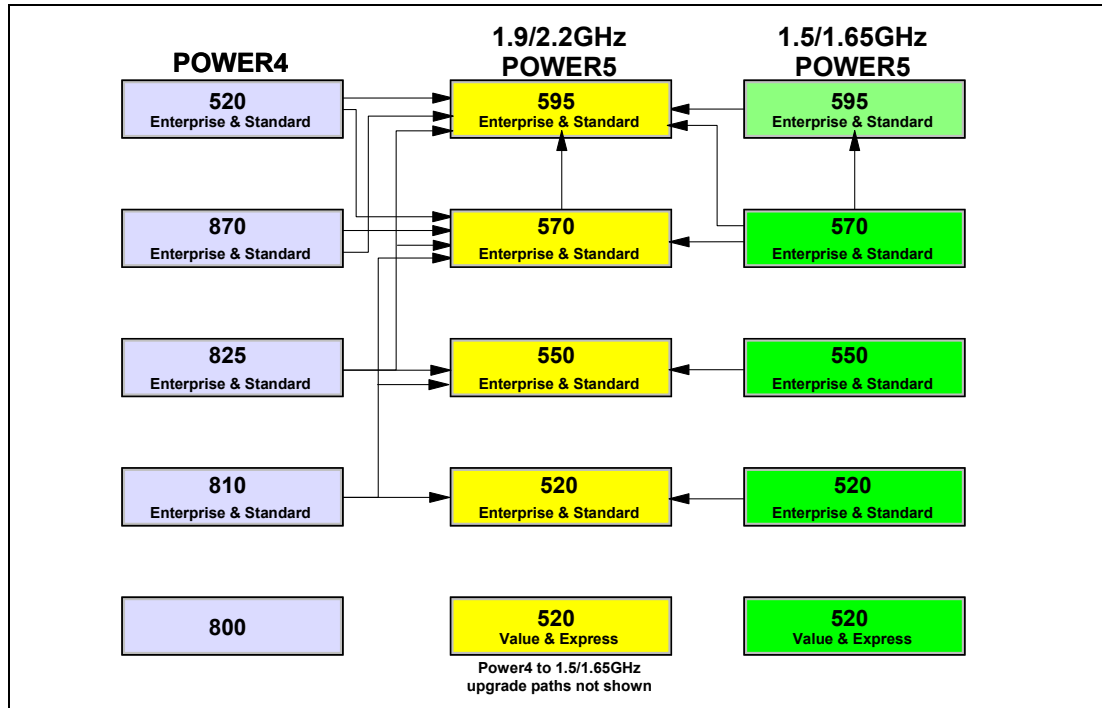


Figure 2-38 Upgrade path options for Standard, Enterprise, Value, and Express Editions

Figure 2-39 shows possible upgrade path options for special edition offers *before* the CBU edition offerings announced in August 2006.

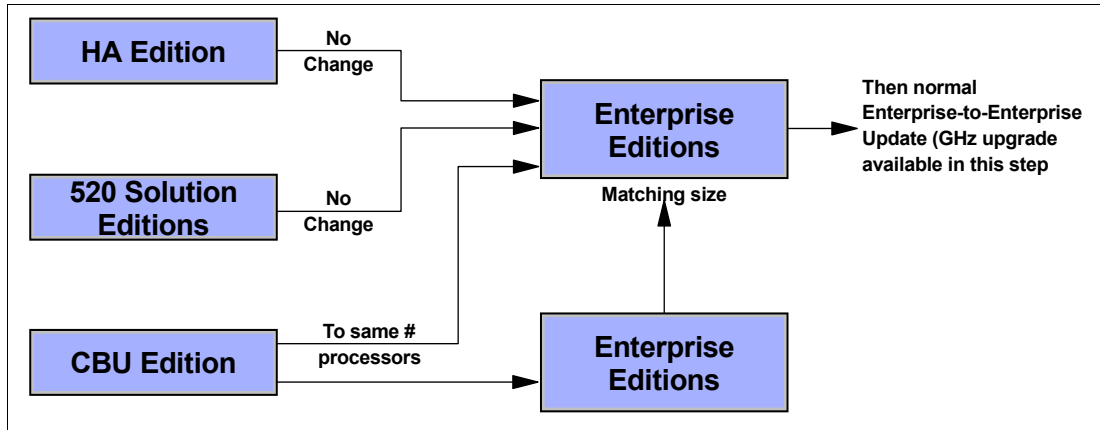


Figure 2-39 Special edition offering upgrade path options

Figure 2-40 shows upgrade options for another set of special edition options that were announced in May 2006.

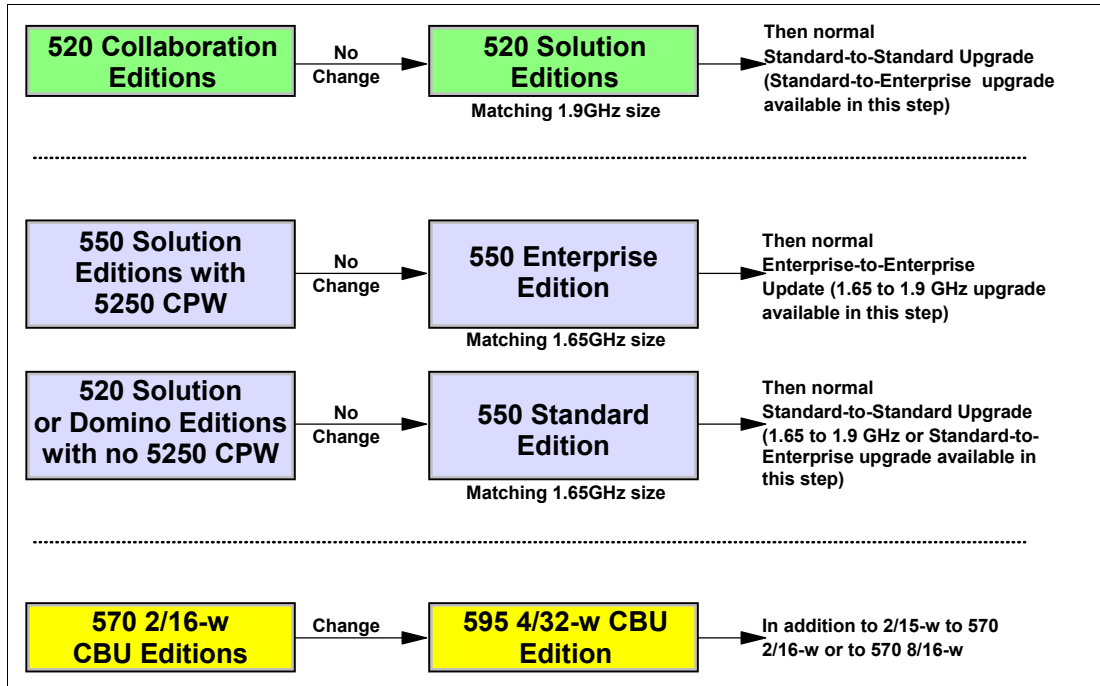


Figure 2-40 Special Edition offering upgrade paths available as of May 2006

The following are additional details for the upgrade options shown in Figure 2-40.

Model 520

Like the existing 520 Solution Edition, which can be upgraded to an equivalent horsepower model 520 Enterprise Edition and then upgraded to a more powerful 520, the new 520 Collaboration Edition can be upgraded to an equivalent horsepower model 520 Standard Edition and then use the existing Standard Edition upgrade paths.

Model 550

Even though most 1.65 GHz 550s will not take advantage of the relatively modest 15% performance boost by upgrading to a 1.9 GHz 550, for those customers who need this capability, the edition feature conversion structure is now available.

Model 570

The 1.65 GHz and the 2.2 GHz 2/16-way 570 CBU Editions now have an upgrade path to the 595 CBU edition. Previously, only upgrades from a 570 2/16-way CBU to a 8/16-way Enterprise or a 1.65 GHz 2/16-way to a 2.2 GHz 2/16-way was available.

Figure 2-41 shows the expanded upgrade options for Capacity Backup configurations after the Capacity Backup options announced August 2006.

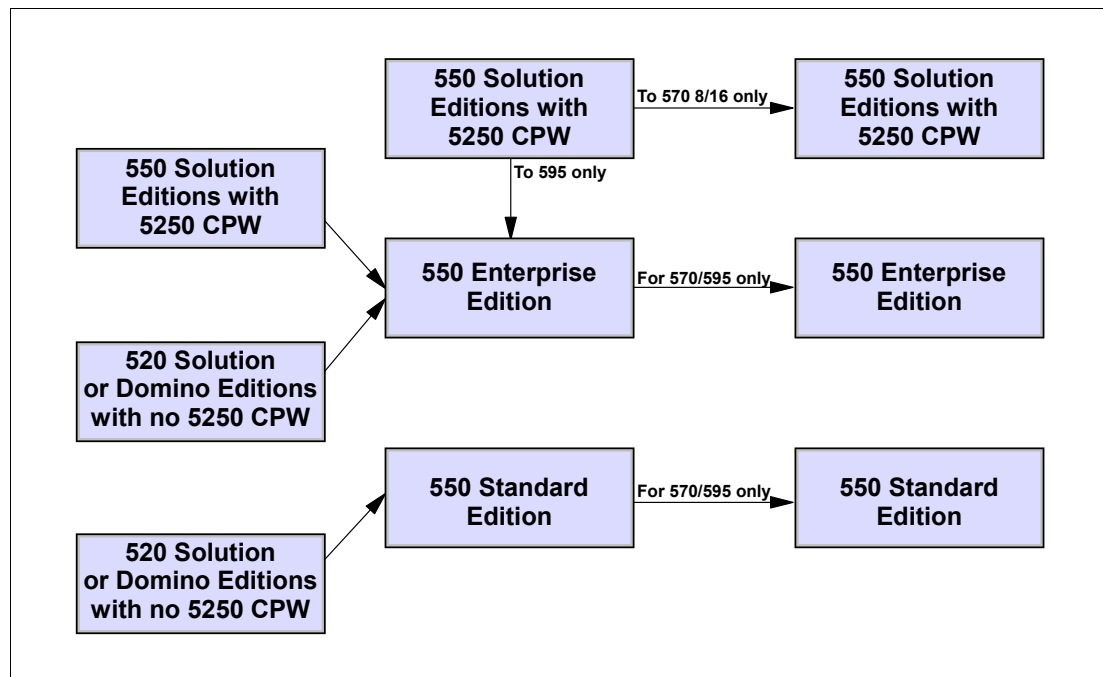


Figure 2-41 Capacity Backup upgrade options as of August 2006

Upgrades from the 825/870/890 are now available (assuming that the customer is eligible for a 570/595 CBU Edition machine). With the addition of permanent processor activations and temporary transfers announced August 2006, there is no longer a reason to provide an upgrade path from the CBU Edition to the Enterprise Edition.



i5/OS work management

This chapter covers V5R4 enhancements to managing your work under i5/OS. It includes new US Daylight Saving Time support to be effective March 2007.

We also review memory affinity and provide updates on adding memory to a partition.

3.1 Work management enhancements

This section is devoted to the tools that you use to interact with jobs running on your system. There are several improvements that can help reduce the time you spend doing these tasks.

3.1.1 Job log management

The i5/OS job description object and the End Job (ENDJOB) and End Subsystem (ENDSBS) commands have parameters to minimize the production of a job log unless a severe job error occurs. However, history has shown, for whatever the reason, that some customer environments end up having dozens of job logs being produced in a short time window, which can possibly impact overall system performance or require extra job log (spooled job log file) cleanup time and effort.

With V5R4 there are additional enhancements that you can take advantage of to reduce job log management.

Changes affecting job log spooling

Scenarios can occur today where hundreds to thousands of spooled job logs are generated that are often not wanted. Occurrences such as the unexpected ending of server jobs due to networking failure, or the immediate ending of jobs (End Subsystem, End Job, End Prestart Job) cause them to be spooled. The i5/OS end commands for jobs and subsystems contain parameters to eliminate or reduce job log production, but many users do not take advantage of them, so V5R4 has some additional enhancements.

Generating job logs consumes resources, including CPU and storage, and can cause output queue contention. It can also increase job termination time and delay recovery.

The i5/OS job description has a message-related parameter in which you can specify LOG(4 0 *NOLIST). This is effective for not producing a job log when the job ends normally. Jobs that end abnormally still generate a job log.

With V5R4, you have improved control of job logs. We have added new interfaces to better manage the spooling of job logs:

- ▶ Additional ways to avoid the spooling of job logs altogether
- ▶ Deferring the generating of job logs to a background server job

There previously existed situations where there was extended system use of job log pending state. This happened when the job had ended, but the job log had not yet been written. Job logs were not written until after the next IPL in certain situations, such as immediate power down (PWRDWNSYS *IMMED), as well as certain abnormal job terminations.

Controlling the output of job logs

With V5R4 there is a new job attribute, Job Log Output (LOGOUTPUT). It has the following parameters:

- ▶ *JOBEND: Job logs will be generated as they are today, during job termination.
- ▶ *JOBLOGSVR: Job logs will be generated by a server job after the job has completed its activity. This new option should be used if you still want job logs but can have their productions generated in a lower priority time frame.
- ▶ *PND: Job logs will not be generated, but will be put into a pending state.
- ▶ *SYSVAL: Use the job log output (QLOGOUTPUT) system value. This can apply to all jobs bs using *SYSVAL.

- ▶ An individual job attribute defaults to use the value in the associated job description object. The job description defaults to use the value in the system value. The system value defaults to *JOBEND for compatibility with prior releases.

We recommend that you set this value to *JOBLOGSVR for best performance.

Job log pending status cleanup

Automatic cleanup (GO CLEANUP) includes pending job logs under *Job logs and other system output*. Job logs never have to be generated, and cleanup will remove the pending job logs from the system.

Customizing your own cleanup is easy with the Remove Pending Job Log (QWTRMVJL) API.

If you are detaching spooled files (SPLFACN(*DETACH)) to free up job structures, jobs in job log pending will continue to use job structures until the job log is printed or otherwise removed from the system.

Other job log considerations

Default behavior for IBM-supplied server jobs will generally be:

- ▶ LOG(4 0 *NOLIST): Only generate job logs if the job abnormally ends.
- ▶ LOGOUTPUT(*PND): When job logs are generated, keep the job log in a pending state.

The new spooled file recovery (SPLFRCY) IPL attribute can prevent deletion of job logs and other spooled files in certain disaster recovery situations.

Pending job logs are now viewable like job logs for active jobs. You can use either of these commands to view them:

- ▶ Display job log (DSPJOBLOG)
- ▶ Work with job (WRKJOB), option 10

New work with job logs (WRKJOBLOG) command

We provide this green screen interface for working with pending and spooled job logs, as seen in Figure 3-1.

```
Work with Job Logs (WRKJOBLOG)

Type choices, press Enter.

Job log state . . . . . *PENDING      *PENDING, *SPOOLED

Time period:
  Start time and date:
    Beginning time . . . . . *AVAIL      Time, *AVAIL
    Beginning date . . . . . *CURRENT     Date, *CURRENT, *BEGIN
  End time and date:
    Ending time . . . . . *AVAIL      Time, *AVAIL
    Ending date . . . . . *CURRENT     Date, *CURRENT, *END
  Job name . . . . . *ALL           Name, generic*, *ALL
  User . . . . . *ALL             Name, generic*, *ALL
  Number . . . . . *ALL           000000-999999, *ALL

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 3-1 WORKJOBLOG command

Note: We would like to note that you can show the spool file with a status of *PENDING and *SPOOLED. Both values are allowed at the same time.

New interface for working with job logs

The Work with Job Logs (WRKJOBLOG) display has been changed, as shown in Figure 3-2.

```
Work with Job Logs                                RCHASM27
                                                06/19/06 17:38:41
Type options, press Enter.
  4=Delete  5=Display  8=Work with spooled files  12=Work with job

Opt  Date      Time      State Job      User      Number  File  User
      06/19/06 09:52:37 *SPL  ADMIN   QEJBSVR   027924   1  ADMIN
      06/19/06 09:52:45 *SPL  QHTTSPVR QSYS     027920   1  QHTTSPVR
      06/19/06 09:54:30 *SPL  ADMIN   QEJBSVR   028323   1  ADMIN
      06/19/06 09:54:31 *SPL  QHTTSPVR QSYS     028320   1  QHTTSPVR
      06/19/06 12:41:40 *SPL  ADMIN   QEJBSVR   028329   1  ADMIN
      06/19/06 12:49:24 *SPL  QHTTSPVR QSYS     028326   1  QHTTSPVR
      06/19/06 12:49:28 *SPL  ADMIN   QEJBSVR   028337   1  ADMIN
      06/19/06 12:53:44 *SPL  QHTTSPVR QSYS     028363   1  QHTTSPVR
      06/19/06 12:53:49 *SPL  ADMIN   QEJBSVR   028366   1  ADMIN
      06/19/06 12:57:10 *SPL  QHTTSPVR QSYS     028371   1  QHTTSPVR
      06/19/06 13:07:01 *SPL  QHTTSPVR QSYS     028378   1  QHTTSPVR
                                                More...

Parameters or command
====>
F3=Exit      F4=Prompt  F5=Refresh  F9=Retrieve  F11=Display output queue
F12=Cancel   F14=Find   F15=Repeat find  F17=Top     F18=Bottom
```

Figure 3-2 WRKJOBLOG display

For information about some general spool file management enhancements with V5R4, refer to Chapter 15, “i5/OS printing and printer output enhancements” on page 377.

3.1.2 Subsystem enhancements

Before V5R4, there was the potential for a performance bottleneck from serial flow of control. There was a recommended limit of 250–300 devices allocated to an interactive subsystem. Exceeding this limit could result in performance slow downs and problems with device error recovery processing done in the subsystem monitor job.

Device I/O constituted a single point of failure. A device I/O error could hang the entire subsystem, which requires subsystem restart to correct.

These locking conditions are not normal, and not everyone has this as a problem on their system. Still, we made enhancements in the V5R4 subsystem.

Subsystem monitor jobs now have multiple threads. Every subsystem has a primary thread and at least one secondary thread. Subsystems with workstation entries will have additional secondary threads to handle device I/O. This has no impact on user applications, except improved performance and recovery. Multiple threads now handle sign-on displays and device recovery. Device I/O no longer is a single point of failure for a subsystem.

3.1.3 Call stack display enhancements

The green screen and iSeries Navigator call stack displays previously supported were:

- ▶ Original Program Model (OPM)
- ▶ Integrated Language Environment® (ILE)

We have added additional call stack display support added for:

- ▶ Portable Application Solution Environment (i5/OS PASE)
 - i5/OS PASE kernel entries require Service (*SERVICE) special authority.
- ▶ Java
- ▶ Licensed Internal Code (LIC) entries
 - LIC entries require Service (*SERVICE) special authority.

Fields have been renamed and reorganized to display a greater variety of data. The command Retrieve Call Stack (QWVRCSTK) API is updated with new formats.

Figure 3-3 shows call stack with i5/OS PASE entries.

Display Call Stack				System:	SYSTEM5
Job:	QPADEV000R	User:	USERONE	Number:	211957
Thread:	00000006				
Type	Program	Statement	Procedure		
	QP2USER QSYS	4	Qp2RunPase		
	QP2USER2 QSYS	7	__Qp2RunPase		
	QP2USER2 QSYS	6	runpase_main__FPi		
	QP2USER2 QSYS	2	runpase_common__FiPvT2		
P	sh	0000008C	__start		
P	sh	00000A44	main		
P	sh	00000680	exfile		
P	sh	000002F8	sh_exec		
	P sh	00000930	xec_switch		
P	sh	00000158	job_wait		
P+	unix		*N		
P	libc.a(shr.o)	00000170	waitpid		
P	unix	00000008	<syscall32>:kwaitpid		
					Bottom
F3=Exit	F5=Refresh	F11=Display activation group	F12=Cancel	F17=Top	
F18=Bottom	F21=Include	F22=Display entire field			

Figure 3-3 Call stack with i5/OS PASE entries

Figure 3-4 shows call stacks with Java entries.

```

                                Display Call Stack
                                System:  SYSTEM5
Job:  QPADEV000R      User:  USERONE      Number:  211959
Thread:  00000002

Type  Program          Statement      Procedure
-----
QJVATools  QSYS              3              main
QJVAQSHCMD QSYS             434            QjvaQshCmd
QJVAJVM    QSYS             199            RunJava
QJVAJVM    QSYS              2              CallStaticVoidMethod__7JNIE >
QJVAJNICLS QSYS              1              CallStaticVoidMethodV__FP7J >
J  hworld          0000119E      main
J  hworld          00001124      main
J  java/io/BufferedReade > 0000109A      readLine
J  java/io/BufferedReade > 00000FB0      readLine
J  java/io/BufferedReade > 00000EC6      readLine
J  java/io/BufferedReade > 00000DDA      fill
J  java/io/InputStreamRe > 00000D14      read

More...
F3=Exit      F5=Refresh    F11=Display activation group  F12=Cancel  F17=Top
F18=Bottom   F21=Include   F22=Display entire field

```

Figure 3-4 Call stack with Java entries

Figure 3-5 shows call stacks with licensed internal code entries.

```

                                Display Call Stack
                                System:  SYSTEM5
Job:  QPADEV000C   User:  USERONE   Number:  020353
Thread:  00000005

Type  Program          Statement          Procedure
L     pminitiateprocess  00000604  pmInitiateProcessUnderTarge >
L     aimach_program_call > 00000054  aimach_program_call_portal
L     cblabbranch       000001CC  cblabbranch
      QCMD      QSYS          /04F3
L     aiuser_program_call > 000000C4  aiuser_program_call_portal
L     cblabbranch       000001CC  cblabbranch
1    QCMD      QSYS          /01C7
L     aiuser_program_call > 000000C4  aiuser_program_call_portal
L     cblabbranch       000001CC  cblabbranch
      QWCCDSAC QSYS          /0568
L     aiuser_program_call > 000000C4  aiuser_program_call_portal
L     cblabbranch       000001CC  cblabbranch
2    QUIMGFLW  QSYS          /04D4
                                           More...

F3=Exit      F5=Refresh  F11=Display activation group  F12=Cancel
F16=Job menu F17=Top    F18=Bottom  F21=Exclude  F22=Display entire field

```

Figure 3-5 Call stack with licensed internal code entries

3.1.4 New job interrupt exit program support

We now provide the ability for the source job to initiate the invocation of a program in a target job.

Use of function and default settings controlled by the Allow jobs to be interrupted (QALWJOBITP) system value, which can be one of these three:

- ▶ Function is disabled and new jobs default to uninterruptible
- ▶ Function is enabled, but new jobs default to uninterruptible
- ▶ Function is enabled and new jobs default to interruptible

Note: We must note that programs that will be allowed to run must be registered in the Job interrupt program exit point (QIBM_QWC_JOBITPPGM). The Call Job Interrupt Program (QWCJBITP) API causes exit program to be run in the initial thread of the target job. The Change Job Interrupt Status (QWCCJITP) API can set the interrupt status of the current job.

3.1.5 WRKACTJOB display now shows job's current user profile

Prior to V5R4, the Work with Active JOBS (WRKACTJOB) command displayed simple job name and job user on the initial panel. Job user is the user profile under which the job started running. In V5R4 the WRKACTJOB initial panel changed to display the current user instead of the job user. The current user is the user profile under which the thread is currently running.

For many jobs, the current user and the job user are the same user profile. But when the initial thread has swapped to a different user profile, it is generally more interesting to know the current user. The current user information may be different when the panel is refreshed.

The job user will still be displayed on the third panel available via F11. Both current user and job user are presented in printed output.

In Figure 3-6 you see the V5R4 Work with Active Jobs command 5250 display, showing the Current User column, where in previous releases you saw the *job started* user profile. Note that in earlier releases this current user profile name was already displayed on most iSeries Navigator job-related windows.

In Figure 3-6, you see user profiles AS0301 for job name JOB0301, JCOOL for job name QPADEV004 (same as job started user profile), and AS0309 for the ODBC/JDBC™ jobs starting with QZDASOINIT. In this example, user AS0309 was using iSeries Navigator - Database Run SQL Scripts.

```

Work with Active Jobs                                RCHAS55B
                                                    06/26/06 17:10:29
CPU %:      9.2      Elapsed time:  00:03:10      Active jobs:  252

Type options, press Enter.
  2=Change  3=Hold  4=End  5=Work with  6=Release  7=Display message
  8=Work with spooled files  13=Disconnect ...

Current
Opt Subsystem/Job User      Type CPU % Function      Status
-  CHAINBCH      QSYS      SBS   .0      DEQW
-  JOB0301      AS0301    BCH   .5      PGM-CHAINUP_CW  RUN
-  QBATCH       QSYS      SBS   .0      DEQW
-  QCMN        QSYS      SBS   .0      DEQW
-  QCTL        QSYS      SBS   .0      DEQW
-  QINTER      QSYS      SBS   .0      DEQW
-  QPADEV0004   JCOOK     INT   .0      CMD-WRKACTJOB   RUN
-  QSERVER     QSYS      SBS   .0      DEQW
-  QPWFSERVSD   QUSER     BCH   .0      SELW
:
:
-  QUSRWRK     QSYS      SBS   .0      DEQW
-  QSCLICEV    QSYS      BCH   .0      DEQW
-  QZDASOINIT  AS0309    PJ    .0      TIMW
-  QZDASOINIT  AS0309    PJ    .0      TIMW
-  QZDASOINIT  AS0309    PJ    8.1     TIMA

More..

Parameters or command
===>
F3=Exit  F5=Refresh  F7=Find  F10=Restart statistics

```

Figure 3-6 Work with active jobs (WRKACTJOB) display - current user

3.2 Daylight saving time support for March 2007

During 2007, the US is changing when daylight saving time (DST) starts and ends. Based upon the Energy Policy Act of 2005, DST is extended one month. Starting in 2007 DST will begin on the second Sunday in March and end on the first Sunday of November.

This change could impact current applications running on any operating system using time dependent values. Consult the following IBM support Web site for information about this subject:

<http://www.ibm.com/support/alerts/daylightsavingtimealert.html>

To support these changed DST dates, i5/OS is introducing the following new time zone values (objects) that can be entered into system value QTIMZON on releases V5R3 and V5R4:

- ▶ QN0500EST3: Eastern Standard Time/Eastern Daylight Time
This corresponds to current value QN0500EST: Eastern Standard Time (EST)/ Eastern Daylight Time (EDT).
- ▶ QN0600CST2: Central Standard Time/Central Daylight Time
This corresponds to current value QN0600CST: Central Standard Time (CST)/Central Daylight Time (CDT).
- ▶ QN0700MST3: Mountain Standard Time/Mountain Daylight Time
This corresponds to current value QN0700MST: Mountain Standard Time (MST)/Mountain Daylight Time (MDT).
- ▶ QN0800PST2: Pacific Standard Time/Pacific Daylight Time
This corresponds to current value QN0800PST: Pacific Standard Time (PST)/Pacific Daylight Time (PDT).
- ▶ QN0900AST2: Alaska Standard Time/Alaska Daylight Time
This corresponds to current value QN0900AST: Alaska Standard Time (AST)/Alaska Daylight Time (ADT).
- ▶ QN1000HAS2: Hawaii-Aleutian Standard Time/Hawaii-Aleutian Daylight Time
This corresponds to current value QN1000HAST: Hawaii-Aleutian Standard Time (HAST) Hawaii-Aleutian Daylight Time (HADT).

This new daylight saving time support is being made available for V5R3 and V5R4 in two stages with sets of PTFs for each stage and release level:

- ▶ Stage 1: PTFs make the new time zone objects available for testing purposes only. Production systems should not be used for testing. Once the PTFs have been installed you can change the system value QTIMZON with the appropriate new value appropriate for your time zone. The PTFs available are:
 - V5R3: SI24716
 - V5R4: SI24717
- ▶ Stage 2: PTFs planned to be available November 2006 are intended to be final and enable the system to automatically update the time values in March and November, starting in 2007. This is if you have specified one of the new QTIMZON values. These PTFs are immediate and thus do not require a partition shutdown and restart to take effect. The PTFs available are:
 - V5R3: SI24906
 - V5R4: SI24908

Note: QTIMZON supports user-defined time zone objects. Any modification to the IBM shipped time zone objects or use of a user-defined time zone object prohibit any automatic update to a partition time value.

We recommend that you review IBM Fix Central information and the PTF cover letters before installing these PTFs. Fix Central can be found at:

<http://www-03.ibm.com/servers/eserver/support/series/>

3.3 Memory affinity

Memory affinity has been available for POWER4™ and POWER5-based technologies on System i (and System p) models since i5/OS V5R3. When there are one or two processors on your hardware configuration, memory affinity has no significance.

Memory affinity is important as the number of processor cards within your system reaches two or more, as memory cards are physically connected to the processor cards. Positioning of L2 and L3 cache within the processor cards and the memory cards can affect performance during heavy peak workload time periods.

When a new thread/task (we use the term *thread* here) is activated it is dispatched to an available processor according to rules not described here. Memory affinity affects both how threads are dispatched to processors and how storage is allocated on their behalf. As the number of processor cards and attached memory cards increases, memory affinity becomes more important.

Essentially, if thread xxx is assigned initially to processor y on card zzz, then the associated L2 and L3 caches and memory cards for that processor contain instructions and data for that task. Typically, a thread runs and then goes into a wait state for something such as a disk operation or a wait on a queue. When that thread is next ready to run, the system tries to assign that thread to the same processor, thereby increasing the chances that the correct instruction stream and data are already in the same caches and main memory as the previous time this thread was active on a processor.

In a simple environment with a processor for each thread ready to run, this is simple to manage. In the real world when there are hundreds or thousands of threads ready to run, things get more complicated.

First, a short review of cache structure on the POWER4 and POWER5. Instructions and data work most efficiently when the instruction stream and data is already in Level 1 cache. In a busy partition, the needed data or instruction stream may not be the L1 cache. The processor effectively waits for a number of cycles while a block of storage is loaded into an L1 cache from the L2 cache. The L2 cache is considerably larger than the L1 cache, so the instructions or data could be in L2 cache. These systems also support a Level 3 (L3) cache, an equal portion of which can be thought of as being associated with each processor chip within the DCM or MCM.

Eight-processor Multi-Core Modules (MCMs) on 595 (and 890) models (four 2-processor cards within the module) or 2-processor Dual-Core Modules (DCMs) on 550 (two 2-processor DCMs) or 570 (up to eight 2-processor cards) models are where memory affinity comes into play — when there are hundreds and perhaps thousands of threads being dispatched.

Both processors on a core share the same Level 2 (L2) cache, while each processor has separate Level 1 (L1) instruction and data caches. A processor executes fastest when the accessed data and instruction stream are in the L1 cache.

For best performance in a heavy workload environment, running a thread in the same processor it last run on would increase the chances of finding associated physical cache and main memory cards instantly ready to process.

Figure 3-7 depicts two 8-processor MCMs and introduces the terms *remote node* memory and *local node* memory. Remote node memory is the memory not physically plugged into the same set of processor cards in which a thread has just been dispatched. This means that the memory needed by the thread is physically plugged into a separate processor card structure.

Memory affinity means to try to dispatch a thread in a processor so that local node memory has the higher chance of having most or at least some of the needed instructions and data if not already in L1, L2, and L3 cache.

Remote and local node memory also applies two processor DCMs, except memory is now remote if it is physically plugged into a DCM different from the one to which the thread has just been dispatched.

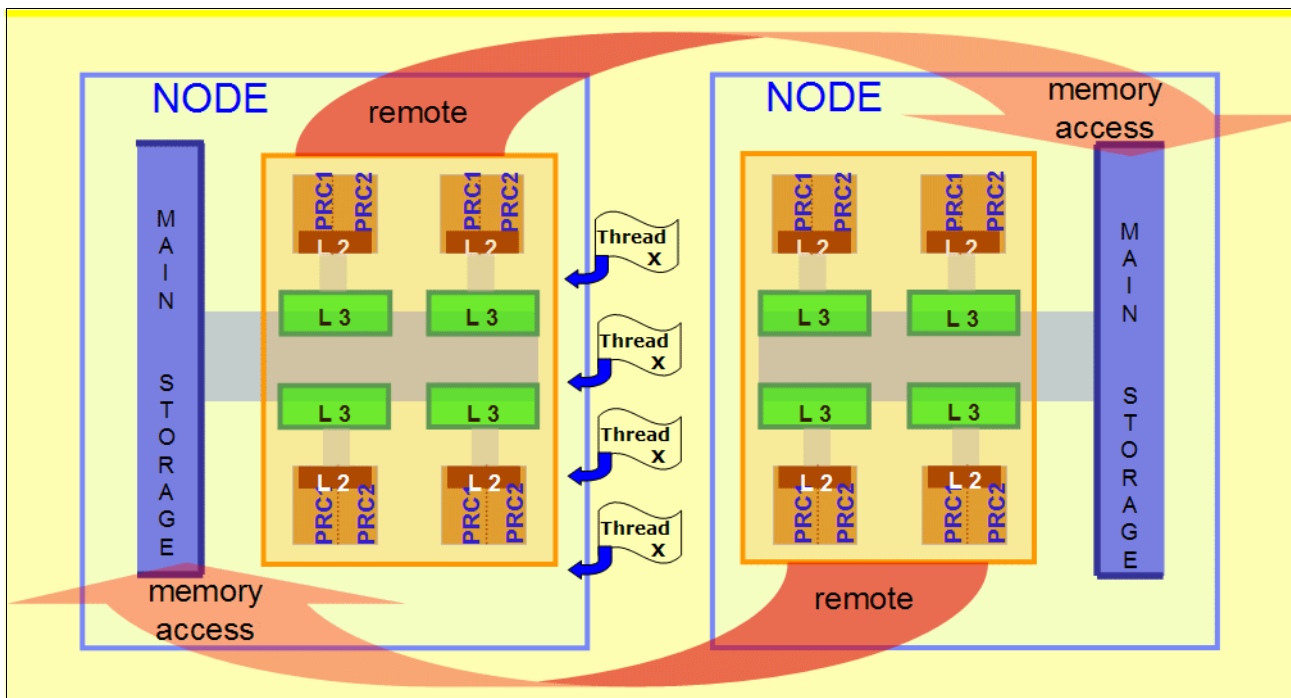


Figure 3-7 Memory affinity view - MCM example

By default, the system tries to dispatch a thread into a processor according to the following priority choices:

- ▶ Same as previous processor used for the thread
- ▶ Processors on the same processor card as the previously used processor
- ▶ Processor on a processor card within the same node as the previously used processor
- ▶ Processor on a processor card and that thread has to access memory in a separate (remote) node

In almost all cases the default action taken by the system delivers the best possible performance. Of course this presumes that you have memory cards placed in the best performance locations. For best memory card placement see Appendix B, “IBM System i5 memory details” on page 499.

i5/OS has system values and subsystem routing entry and prestarted job entry parameters that use default values, but permit changes to the default values.

We list the system values and routing entry and prestarted job entry parameters that tell the system how to manage threads and associated local node memory:

- ▶ System Value QTHDRSCAFN with two parameters
You can use this system value to specify whether secondary threads (after the primary thread has initially been dispatched) are grouped with the initial thread (thread level). It also specifies the degree to which the system tries to maintain the affinity between threads and the subset of resources they are assigned to (level of affinity):
 - Group: *Group or *Nogroup (default)
 - Level: *Normal (default) or *High
- ▶ System Value QTHDRSCADJ
Controls whether microcode (SLIC) is enabled to change the home node of existing tasks/threads to achieve a better nodal balance
 - Default value: 1 = dynamic adjustment
 - System uses a node’s target weight, current weight
The weight values are calculated by the task dispatching code’s memory of where threads have run most recently and the current memory configuration. Note that dynamic LPAR movement of processor and memory resources and physical card placement affects these calculations and when weight values are adjusted.
- ▶ Subsystem routing entry and prestarted job entry memory affinity parameters
The parameters optionally override the system values for the job processed by the specific routing or prestarted job entry:
 - THDRSCAFN
 - Group: *SYSVAL (default), *Group, or *Nogroup
 - Level: *Normal or *High
 - RSCAFNGRP
 - *Yes or *No (default)

Discussion of when to consider choosing a parameter value different from the default value is beyond the scope of this book.

A thorough review of memory affinity can be found in the white paper *Enhanced eServer iSeries Performance: Memory Affinity*. This document can be found at the performance management Web site:

<http://www.ibm.com/eserver/series/perfmgmt>

Select the **Resource Library** link and the heading **Performance Papers and Articles**.

The V5R3 Technical Overview presentation set, Performance Update presentation, also discusses memory affinity and simultaneous multi-threading available on POWER5 technology. This presentation can be found at:

<http://www-03.ibm.com/servers/eserver/support/series/index.html>

1. Select the **i5/OS V5R4 Technical Overviews** link.
2. Select the **System i Technical Overviews** link.
3. Select **V5R3** link.



DB2 Universal Database for i5/OS

This chapter covers the strategic initiatives of and V5R4 enhancements to IBM DB2 Universal Database for iSeries.

In this chapter we use the terms *DB2 Universal Database for iSeries*, *DB2 Universal Database for i5/OS*, and *DB2 for i5/OS* interchangeably. This is because the IBM System i Web site for database (<http://www.ibm.com/series/db2>) has introduced the term DB2 Universal Database for i5/OS and DB2 for i5/OS, but almost all Information Center documentation and the formal product name still uses *for iSeries*. Over time *for i5/OS* will be reflected in all documentation and product names.

Four excellent redbooks describing V5R4 DB2 for i5/OS became available during 2006. See the following redbooks for more detailed information than included in this chapter:

- ▶ *SQL Performance Diagnosis on IBM DB2 Universal Database for iSeries*, SG24-6654
- ▶ *Preparing for and Tuning the SQL Query Engine on DB2 for i5/OS*, SG24-6598
- ▶ *The Ins and Outs of XML and DB2 for i5/OS*, SG24-7258
- ▶ *Stored Procedures, Triggers, and User-Defined Functions on DB2 Universal Database for iSeries*, SG24-6503

4.1 Strategic initiatives

There are three strategic initiatives that drive DB2 Universal Data Base development:

- ▶ Openness: industry standard support
 - Accommodating Independent Software Vendors (ISVs)
 - Portability/compatibility.
 - Flexibility
- ▶ Continued *leadership* in database technologies
 - Consistency across DB2 family
 - Shared R&D across IBM labs
- ▶ Continued leveraging of System i strengths
 - Availability
 - Scalability
 - Usability: total cost of ownership
 - Application flexibility

As you progress through this chapter, relate the V5R4 enhancements back to these strategic initiatives.

4.2 Application flexibility and portability

With features such as Free Format embedded SQL in ILE RPG, Scalar Fullselect, and Standard Flagger included in V5R4, DB2 Universal Database for i5/OS continues its position as a leader in ease of use and integration. V5R4 contains many new capabilities that continue to make DB2 Universal Database for i5/OS a database that is simple to use for programmers, analysts, and administrators.

Additionally, DB2 Universal Database for i5/OS will continue to leverage and build on the strengths the System i family of products. Those strengths include reliability, availability, total cost of ownership, scalability, and usability.

4.2.1 Integration of free-format RPG and SQL

Many RPG programmers have improved their productivity in recent years with the usage of the free-form syntax. V5R4 improves this productivity by allowing SQL requests to be embedded into free-format RPG for the first time, as demonstrated in Example 4-1.

Example 4-1 Integration of free-format RPG and SQL

```
d get Order Count    pi          10i 0
d inCustNumber      9b 0 const
d outSqlState       5a
d outSqlMsg         256a

d ordercount       s          10i 0

/free
  outSqlMsg = *blanks;
  exec sql SELECT count(*) into :orderCount FROM orders
          WHERE cust_no = :inCustNumber;
```

```

outSqlState = sqlState;
if %subst(sqlState:1:2) <> '00';
    exec sql GET DIAGNOSTICS CONDITION 1
                :outSqlMsg = MESSAGE_TEXT;
    orderCount = 0;
endif;
return orderCount;
/end-free

```

4.2.2 Standards compliance

SQL is the industry standard for database access and programming. While the heritage of application development on the System i platform has been to use native interfaces, more and more application developers are using SQL as a programming language as the result of the DB2 Universal Database for i5/OS strategic investment in SQL.

You can see the trend toward greater use of SQL in Figure 4-1. SQL is the strategic database interface for DB2 Universal Database for i5/OS to accommodate all the porting of third-party applications and tools — almost all of which third-party solutions utilize SQL-based interfaces for their data access.

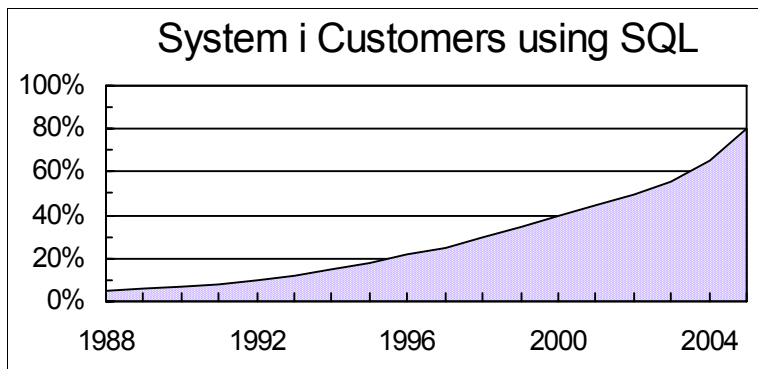


Figure 4-1 The industry trend toward off-the-shelf software results in a move to SQL

Enhanced SQL standards support

With V5R4 inclusion of Subquery and Scalar Fullselect and Core Flagger support, DB2 for i5/OS became the first heavy duty DataBase Management Systems (DBMS) to fully comply with the core level of the SQL 2003 Standard.

The beauty of the IBM DB2 UDB family is that we can all share from this, and DB2 UDB for iSeries will continue to leverage the technology leadership position of IBM and maintain close compatibility with the other DB2 UDB family members.

Note: SQL examples of the V5R4 enhancements are in 12.7, “Structured Query Language (SQL)” on page 290.

DB2 UDB for iSeries Version 5 Release 4 complies with the following IBM and Industry SQL Standards:

- ▶ International Standards Organization (ISO) 9075: 1992, Database Language SQL - Entry Level
- ▶ ISO 9075-4: 1996, Database Language SQL - Part 4: Persistent Stored Modules (SQL/PSM)

- ▶ ISO 9075: 2003, Database Language SQL - Core
- ▶ American National Standards Institute (ANSI) X3.135-1992, Database Language SQL - Entry Level
- ▶ ANSI X3.135-4: 1996, Database Language SQL - Part 4: Persistent Stored Modules (SQL/PSM)
- ▶ ANSI X3.135-2003, Database Language SQL – Core

4.2.3 Improved DB2 family compatibility

There are a number of V5R4 enhancements that have benefitted from DB2 family development. These include:

- ▶ iSeries Navigator database management capabilities such as the DB2 On Demand Performance Center that simplifies SQL performance analysis for database analysts and administrators.
- ▶ i5/OS programmers get a richer SQL toolbox.
- ▶ The addition of ANS-ISO Timestamp format.
- ▶ An increase of up to 2 MB SQL statements, and 32 K keys.
- ▶ SQL programmers will be able to solve a wider range of business problems with additions such as Recursive SQL, Triple DES encryption, and the integration of SQL with Free Format RPG syntax.
- ▶ Improved Materialized Query Table support, which enables the database to take on an even broader range of workloads.
- ▶ Performance and scalability are improved on a number of fronts boosted by the SQL Query Engine (SQE) handling more DB2 requests, and its On Demand Index capability, most SQL procedures, functions, and triggers will receive a significant performance boost with V5R4.

4.3 Database application development enhancements

This topic focuses mainly on key changes, allowing the application developer to integrate his application within a heterogeneous environment.

Refer to Chapter 12, “i5/OS-based application development languages and tools” on page 269, where specific items, such as changes in the SQL language, are explained.

In addition to Embedded SQL, many iSeries developers are using industry-standard interfaces such as JDBC and .NET to build innovative solutions. Enhancements to these programming interfaces that are part of iSeries Access for Windows include customization for better integration of iSeries-specific capabilities along with keeping current with industry standards.

4.3.1 iSeries .NET provider

.NET is Microsoft’s latest language-neutral environment for writing programs that can easily interoperate, especially with a variety of databases from Windows and non-Windows operating systems. .NET is also the collective name given to various software components built upon the .NET platform. The components that make up the .NET environment are referred to as the .NET Framework.

The .NET Framework is a comprehensive, object-oriented collection of reusable classes that you can use to develop applications ranging from traditional command-line or graphical user interface (GUI) applications to applications based on the latest enhancements or changes provided by ASP.NET, such as Web forms and XML Web services. Typically, .NET applications use the Microsoft ADO.NET classes to access and manipulate database objects locally or remotely.

In many ways sets of .NET capabilities are analogous to IBM WebSphere-base application interfaces and application server management capabilities, based upon the WebSphere Application Server.

The .NET Framework running on a Windows operating system includes a common language Runtime (CLR) component. CLR is shipped with Windows 2003 Server.

The common language runtime manages memory, thread execution, code execution, code safety verification, compilation, and other system services. These features are intrinsic to the managed code that runs on the common language runtime. Similar to the Java Virtual Machine (JVM™) environment, .NET programs are not compiled directly into executable code, but are compiled into an intermediary language known as Microsoft Intermediary Language (MSIL, or IL). Later, at program execution, the CLR loads the code into the runtime environment and a just-in-time compiler (JIT) compiles the IL into native executable code. This native code is then executed by the runtime's execution engine.

Figure 4-2 is a high-level depiction of using .NET interfaces to access database objects under i5/OS. Figure 4-3 on page 89 shows ADO.NET as a follow-on to the still popular Microsoft Open DataBase Connectivity (ODBC) functions and interfaces.

Under i5/OS, the .NET provider is named IBM.Data.DB2.iSeries. It allows applications using Microsoft's .NET framework to access DB2 UDB for iSeries databases.



Figure 4-2 .NET provider for Windows

For complete documentation of the .NET Data Provider, see the IBM DB2 UDB for iSeries .NET Provider Technical Reference. This technical reference is part of the Programmer's Toolkit, an optionally installed component of the iSeries Access for Windows product. Be sure that you have the Programmer's Toolkit component installed on your PC, then to get to the IBM DB2 UDB for iSeries .NET Provider Technical Reference select **iSeries Access for Windows → Programmer's Toolkit → Database → .NET Framework Classes**.

Improvements to the .NET provider support under V5R4 i5/OS include:

- ▶ LOB column support
Available also with the latest V5R3 Service Pack for iSeries Access for Windows (SI23492 or later).
- ▶ System naming and library list support
Available also with the latest V5R3 Service Pack for iSeries Access for Windows (SI23492 or later).
- ▶ MS FW 2.0 compatibility
Available also with the latest V5R3 Service Pack for iSeries Access for Windows (SI23492 or later).
- ▶ IntelliSense® support to aid programmers
IntelliSense is a feature introduced and popularized by the Microsoft Visual Studio® Integrated Development Environment. It involves a form of *autocomplete* for variable names, functions, and member functions (for C++ classes). Using IntelliSense is a convenient way to access descriptions of functions and their parameter lists. It facilitates faster software development by reducing the amount of keyboard input required. For more information go to the Microsoft Web site:
<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/vsintro7/html/vcovrautomaticstatementcompletion.asp>
- ▶ Multiple active result sets on a connection
- ▶ Customizable string processing

Table 4-1 shows the large set of new .NET to i5/OS capabilities in V5R4 and also highlights V5R3 .NET enhancements.

Table 4-1 .NET enhancements V5R3/V5R4

V5R3 enhancements	V5R4 enhancements
Requires OS/400 V5R2 or later <ul style="list-style-type: none"> ▶ SQL (INSERT, UPDATE, DELETE) ▶ Commitment control ▶ Connection pooling ▶ SQL naming ▶ Unicode ▶ Tracing ▶ Threads ▶ IASPs (multiple databases) ▶ Stored procedure support ▶ iSeries-specific properties ▶ User-defined types Supported after GA by SP SI15176 or later <ul style="list-style-type: none"> ▶ System naming (/) ▶ Library list ▶ Large objects (LOBs) ▶ LOB column support ^a ▶ System naming and library list support ^a ▶ MS FW 2.0 compatibility ^a 	Requires i5/OS V5R4 <ul style="list-style-type: none"> ▶ 128-byte column names ▶ Maximum SQL statement lengths of 2,097,152 bytes or 1,048,576 characters ▶ Support for IBM Enterprise Workload Manager (eWLM) correlator Requires OS/400 V5R2 or later <ul style="list-style-type: none"> ▶ System naming and library list ▶ Support for LOB data types ▶ Support for multiple active result sets per connection ▶ Customizable String processing for Char for bit data, date, time, timestamp, decimal, and numeric data types ▶ Support for IntelliSense ▶ Additional sample programs ▶ Improved support for delimited names ▶ LOB column support ▶ System naming and library list support ▶ MS FW 2.0 compatibility ▶ Multiple active result sets on a connection ▶ Customizable string processing ▶ JDBC – Version 3 currency and performance ▶ Optimization Goal connection attribute ▶ IBM EWLM support added to CLI, DRDA®, .NET, ODBC, JDBC ▶ Driver support for Windows Vista™ (depending on availability) ▶ OLE DB Driver ▶ System naming and library list support ▶ ODBC driver ▶ Optimization goal connection attribute

a. Supported after GA with Service Pack - S23492 or later.

These V5R4 enhancements bring i5/OS more closely in line with Microsoft's technology roadmap, shown in Figure 4-3.

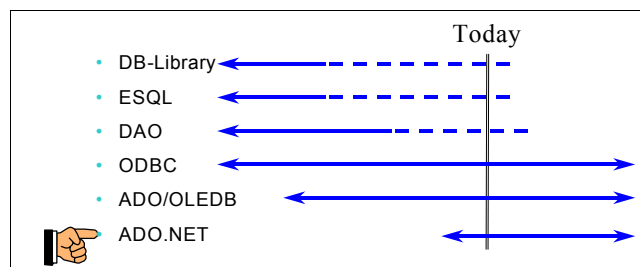


Figure 4-3 Microsoft's technology roadmap

i5/OS IBM Redbook - SG24-6440

The relatively new IBM Redbook *Integrating DB2 Universal Database for iSeries with Microsoft ADO .NET*, SG24-6440, describes the .NET database access to i5/OS integrated DB2 UDB for iSeries.

That book addresses Windows applications using the Microsoft ADO.NET interface to access the i5/OS database (DB2 UDB for iSeries). Applications using ADO.NET with the iSeries can work with several different .NET providers:

- ▶ The IBM.Data.DB2.iSeries provider, a .NET-managed provider new to iSeries Access for Windows with V5R3
- ▶ The IBM.Data.DB2 provider, a .NET provider that works with all operating systems supporting IBM DB2 Universal Database in conjunction with DB2 Connect™
- ▶ The Microsoft System.Data.OleDb provider, as a bridge to one of the OLE DB providers included with iSeries Access for Windows (IBMDA400, IBMDASQL, and IBMDARLA)
- ▶ The Microsoft System.Data.Odbc provider, as a bridge to the ODBC driver included with iSeries Access for Windows

Important: This book shows customers how to use ADO.NET effectively to harness the power of DB2 UDB for iSeries, showing examples, best practices, pitfalls, and comparisons between the different ADO.NET data providers.

4.3.2 DB2 Connect Unlimited Edition for iSeries

The recent introduction of DB2 Connect Unlimited Edition for iSeries also makes DB2 plug-ins for Microsoft Visual Studio more accessible to .NET developers wanting to process DB2 UDB for iSeries databases. DB2 Connect Unlimited Edition is a special version of IBM DB2 Connect offering packaged and priced for the iSeries marketplace.

There is a new customized package for the iSeries developer community, available since 3Q05. It provides affordable access to:

- ▶ DB2 add-ins for Visual Studio
- ▶ Federated data access
- ▶ Joining of DB2 tables in different iSeries partitions
- ▶ Joining of DB2 tables across iSeries and LUW server
- ▶ Extend iSeries data to mobile devices: DB2 Mobility on Demand

For additional details see:

<http://www-306.ibm.com/software/data/db2/db2connect/edition-uei.html>

In addition to the .NET plug-ins, this product also provides heterogeneous access to non-DB2 servers with the federated database support and the ability to make DB2 UDB for iSeries data available on mobile devices. IBM DB2 UDB for several operating systems, including i5/OS, supports the basic constructs and accessibility of a federated database.

A federated database is one that presents a logical single database of several independent databases running on independent servers with no sharing of physical resources, such as disks, connected over a communication network. Formally, the data is horizontally partitioned among the database servers.

The applications accessing the data via *distributed SQL* see a single logical database. Microsoft calls this technology Distributed Partitioned Views (DPVs). Federated database support is included in Microsoft SQL Server software.

Technically, a federated database architecture is different from a shared disk cluster architecture used by, for example, Oracle9i Real Application Clusters.

The DB2 Connect offering also simplifies application development with Web services support and provides a portable interface for applications that need to access any of the DB2 UDB family members. The DB2 Connect Unlimited edition does not run natively on i5/OS. It connects to i5/OS from Windows, UNIX®, and Linux clients. More details about this new product can be accessed on the referenced Web site.

4.3.3 IBM EWLM support added to CLI, DRDA, .NET, ODBC, JDBC

IBM Enterprise Workload Manager (EWLM) allows you to define business-oriented performance goals for an entire domain of servers, and then provides an end-to-end view of actual performance relative to those goals. With EWLM, you can ensure that work requests across a mixed-platform, multi-tiered application environment are performing according to your expectation, and the V5R4 enhancements make it easy to instrument your DB2 workloads for EWLM.

CLI enhancements

The SQL Call Level Interface (CLI) is a popular server-side programming interface used by many applications for data access. An increased handles limit and new connection and statement attributes give application developers greater flexibility when enhancing their applications. Changes include:

- ▶ Maximum handles limit doubled to 160,000
- ▶ SQLFetchScroll block-fetch and column-wise binding
- ▶ Column-wise blocked insert binding
- ▶ Optimization goal connection attribute
- ▶ Cursor sensitivity statement attribute
- ▶ Improved XA documentation
- ▶ New SQLGetInfo and SQLColAttributes options
- ▶ Max rows attribute (SQL_ATTR_MAX_ROWS)

XA over DRDA

Extended Architecture (XA) is an open systems architecture intended to handle distributed transactions. XA over Distributed Relational Database Architecture™ (DRDA) support was restricted in i5/OS releases prior to V5R4.

XA and JTA background information

Before summarizing the new V5R4 support we first provide background information about XA. Initial support under i5/OS (OS/400) was provided in V5R2.

IBM DB2 Universal Database (UDB) has supported distributed transactions and two-phase commit for a distributed set of database operations for years with the COMMIT and ROLLBACK operations under DRDA support. The initial implementation was over SNA-based networks. When multiple database sources were being accessed, the *two-phase* term is applied to:

- ▶ Phase 1
Ask the distributed databases to *prepare to commit* or *prepare to rollback* an in-progress transaction.
- ▶ Phase 2
Based upon the phase 1 responses, phase 2 tells the database managers to actually commit or roll back.

The industry-defined distributed transaction-oriented paradigm is a TCP/IP-based extension of earlier IBM DRDA implementations under the IBM DB2 UDB implementations. The distributed transaction architectures and implementations include two-phase commit capabilities. Because the work being performed (within the bounds of such a transaction) can occur on many different platforms and involve many different databases from different vendors, the TCP/IP-based standard has been developed to allow a manager process to coordinate and control the behavior of the databases. X/Open is a standards body that developed the Distributed Transaction Processing Model (commonly referred to as XA) to solve accessing heterogeneous databases. X/Open applications run in a distributed transaction processing (DTP) environment. In an abstract model, an X/Open application calls on resource managers (RMs) to provide a variety of services. All of this is managed under a Transaction Manager (TM).

The X/Open Distributed Transaction Processing Model consists of:

- ▶ Definitions of the communications between an application, a Transaction Manager (TM), and one or more Resource Managers (RM). This includes the protocol for transaction coordination, commitment, and recovery between a TM and one or more RMs.
- ▶ An application program that specifies actions that constitute a transaction. This application can use Java Transaction APIs (JTAs) to request functions and obtain feedback, especially important when updating, deleting, or adding records (rows) of tables within a logical database within a transaction.
- ▶ Resource Managers (RMs) that provide access to shared resources. The most common RM is a database (Relational DataBase Management System, RDBMS), for example, IBM DB2 UDB for xxx (for example, DB2 for i5/OS), Oracle, SQL Server, and Sybase. However, a print spooler could be an RM if it were to support commitment control through the X/Open XA interface.
- ▶ A Transaction Manager that assigns identifiers to transactions, monitors their progress, and takes responsibility for transaction completion (COMMIT) and for failure recovery (ROLLBACK). The transaction resources are typically thought of as possibly different Relational DataBase Management Systems housed on a single server. They could also include several instances of a single implementation of an RDBMS, such as IBM DB2 UDB, residing on a number of different servers or different implementations of an RDBMS on each server.

In any case, a distributed transaction involves coordination among the various resource managers.

XA is often associated with the Java Transaction API (JTA) interface, which supports complex transactions. It also provides support for decoupling transactions from internal (usually transparent to the application) *connection objects*. A connection object represents a connection to a data source in Java Database Connectivity (JDBC). It is through connection objects that *statement objects* are created for processing SQL statements against the database. An application program can have multiple connections at one time. These connection objects can all connect to the same database (RDBMS) or connect to different databases. As JDBC is modeled after the Object Database Connectivity (ODBC) and the X/Open Call Level Interface (CLI) specifications, JTA is modeled after the X/Open Extended Architecture (XA) specification. JTA and JDBC work together to decouple transactions from connection objects.

Though typically not seen by the application, the decoupling of transactions from connection objects allows the system to have a single connection work on multiple transactions concurrently. Conversely, it allows you to have multiple connections work on a single transaction. In OS/400 V5R2 the marketing term used was *Adaptive e-transaction services* to distinguish this support from previous release implementations.

Needless to say, operating systems, database products, and some middleware products have specific software to implement the wide range of XA and JTA capabilities. For example, since IBM WebSphere Application Server products support distributed transaction and database capabilities, XA TM and RM support are included.

Note: i5/OS V5R3 also has XA support enhancements, but they are not covered in this book.

i5/OS V5R4 and XA support enhancements

In V5R4, Distributed Relational Database Architecture (DRDA) support has been extended to include the following functions. Prior to V5R4, this support was disabled over DRDA.

- ▶ Distributed transaction processing (XA/JTA)
- ▶ Server support for profile tokens
- ▶ DB2 Multisystem over TCP/IP

Extending i5/OS XA support to DRDA allows developers to take advantage of distributed transactions through DRDA-based drivers such as the DB2 Universal JDBC driver within DB2 Connect Unlimited Edition for iSeries.

Figure 4-4 depicts the application environments and IBM DB2-based database servers supporting XA over DRDA.

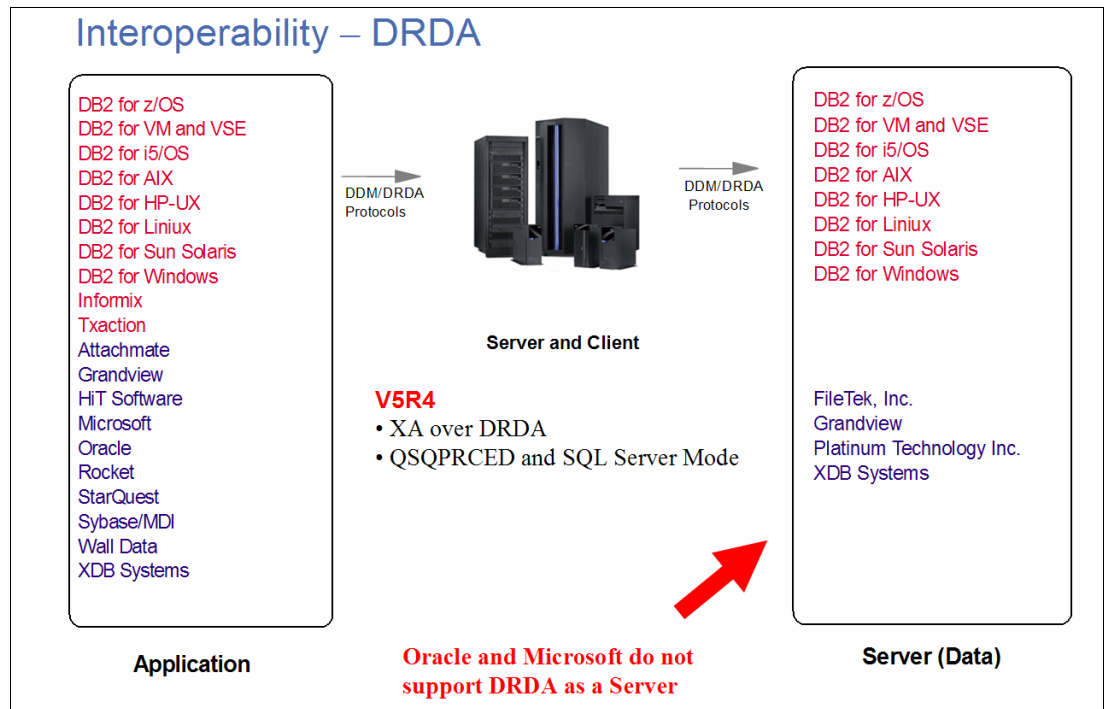


Figure 4-4 XA Interoperability over DRDA

Redesigned SQL Descriptor Area (SQLDA)

The SQL Descriptor Area is used either directly or indirectly by all Dynamic SQL interfaces such as CLI, JDBC., and .NET to manage input and output data. With V5R4, IBM redesigned the internal structures and processing for SQLDA to be more efficient.

JDBC has Version 3 currency and performance improvements, with changes to the Optimization Goal connection attribute. The OLE DB Driver now has System naming and library list support, while the ODBC Driver has an Optimization Goal connection attribute.

This internal redesign has resulted in a nice performance boost for Dynamic SQL applications in V5R4, without any programmer intervention.

We have added:

- ▶ Support for longer column names
- ▶ Faster internal processing

Bigger SQL and table limits

The bigger limits are:

- ▶ 2 MB SQL statement maximum (from 64 K)
- ▶ 1000 tables per query (DML only) (not for views, MQTs, or partitioned tables)
- ▶ 128 byte column names (from 30)
- ▶ 1024 parameters for stored procedures, C, and C++ programs (previous maximum was 90 parameters)
- ▶ 32 K index keys in general (previous maximum was 2000)
- ▶ 32 K index keys for ORDER BY (previous maximum was 10000)

SQL programmers can now be more descriptive with their column names with the maximum size moving from 30 bytes to 128 bytes. The programmers can embed more complicated business logic into their SQL statements by allowing up to 1,000 tables to be referenced and by coding statements with lengths up to 2 MB. The previous statement length maximum was 64 K. This increase in the statement maximum length will be most beneficial to programmers creating SQL procedures, functions, and triggers.

The limit of 256 tables per query continues to apply to views, MQTs, and partitioned tables.

New SET SESSION authorization statement

With V5R4, the new set session authorization statement enables an administrator to create and modify objects under another user's i5/OS user profile. That other profile can continue to own the object.

Characteristics of this support include:

- ▶ New SESSION USER and SYSTEM USER special registers. New special registers allow an application to retrieve the current session user as well as the shared user ID used to set up the connection with the system user register.
- ▶ *ALLOBJ authority required to execute.
- ▶ Note that other settings, such as the SQL Path, may need to be reset.

SET SESSION AUTHORIZATION is also available on V5R3 with DB Group PTF SF99503 level 9 or later.

Note: Use of this statement essentially changes the user profile under which that job was running. For example, a prestart QZDASOINIT job starts running under the default profile QUSER. Specifying a different user profile on the set session authorization statement changes that job to be running under the specified user profile. If you use this statement after running several queries in that job, changing the user ID may require the SQL optimizer to recalculate the current access plan. This could temporarily degrade performance as the new access plan is being built.

4.3.4 Stored procedures

This section describes several enhancements to stored procedures support.

DB2 Expression Evaluator for faster SQL procedural language

Improved performance is the main V5R4 highlight for SQL stored procedures. As you may know, when an SQL procedural object is created, the DB2 UDB engine generates C code with embedded SQL to implement the business logic defined with SQL. You will probably agree that generated code almost never performs as well as code that you write yourself and that has been true for SQL procedures since they were introduced in V4R2.

A significant change was made to the DB2 UDB C code generation engine for the SQL procedural language in V5R4. The improvements on your system will vary, but some tests in the lab showed a 30% performance improvement. All SQL procedures, triggers, and functions must be recreated on V5R4 to benefit from this enhancement.

Easier authority setup

Support for DFTRDBCOL and DYNDFTCOL attributes is useful when porting from other DBMSs. This enhancement is also available on V5R2 and V5R3 with latest Database Group PTF, SF99503, level 9 or later.

Simpler maintenance

The alter procedure statement simplifies the ongoing development and maintenance of SQL procedures. This new statement enables you to quickly change the attributes or logic of an existing procedure and not have to worry about deleting the existing procedure first or resetting the ownership and privileges for the procedure object.

In Example 4-2, the alter is used to change the Increase_Level procedure logic to grant a 10% salary increase instead of the 5% increase in the original procedure.

Example 4-2 Alter used to change Increase_Level procedure logic

```
ALTER PROCEDURE Increase_Level
REPLACE
BEGIN
    DECLARE CurLvl INT;
    SELECT edlevel INTO CurLvl FROM emptbl WHERE empno=Emp#; IF NwLvl >
CurLvl THEN UPDATE emptbl SET edlevel=NwLvl, salary=salary + (salary*0.10)
    WHERE empno=Emp#;END IF
END
```

4.4 On Demand and availability

To improve the availability and On Demand capabilities, V5R4 improves the robustness of the i5/OS cross-reference files that support the DB2 catalog views (SYSTABLES, SYSCSTCOL, SYSTRIGCOL, and so forth). The improved protection of the queue responsible for populating these files should reduce the number of times that recovery efforts are needed in the event that you have an abnormal partition restart (IPL).

In the rare cases in which a recovery of the cross-reference files is still needed, the new Reclaim Database Cross-Reference (RCLDBXREF) command allows for a more granular recovery at a library level, and system-wide recoveries now include a progress and status indicator. In addition, DB2 UDB for iSeries will also automatically rebuild the QSYS2 and SYSIBM catalog views when necessary, eliminating the need for manual intervention.

System Managed Access Path Protection (SMAPP) and access path journaling is now provided for encoded vector indexes (EVIs) and for most access paths with international components for Unicode (ICU) sort sequence tables.

Journaling (or logging) is a key component of any business solution that needs to be highly available. The i5/OS journaling capabilities have been enhanced greatly over the last couple of releases, and V5R4 is no exception. The automatic journaling for SQL-created database is improved with:

- ▶ Support for SEQUENCE object
- ▶ Automatic SQL journaling with QDFTJRN data area after CrtDupObj and Restore commands:
 - Data Area layout
 - Bytes 1–10: Journal Library
 - Bytes 11–20: Journal Name
 - Bytes 21–30: Object Type (*ALL,*FILE,*NONE, *DTAARA, *DTAQ)
 - Bytes 31–40: Option (*ALLOPR,*CREATE, *MOVE, *RESTORE)

Journal management should also be simpler with the enhancements to the Change Journal (CHGJRN) and Work with Journal (WRKJRN) system commands and the new maximum object limit.

Journal minimal data was a feature introduced a couple of releases ago to reduce the size of journal objects on the system by only writing the data that actually changed to the journal. The new Field Boundary option (*FLDBDY) still minimizes the amount of data written to the journal, but it does it at a field (column) level instead of a byte level. This allows the minimized journal entries to be viewable for audit purposes.

Single journal support is extended to 10,000,000 objects.

There was an extension to the QDFTJRN support. This allows you to indicate that tables, data areas, and data queues created, moved, or restored into the library should automatically start journaling.

Figure 4-5 shows a QDFTJRN format extension.

Offset	Field	Format	Description
1	Library name	Char(10)	Name of the library containing the journal
11	Journal name	Char(10)	
21	Repeat the set of Object type and Option as needed		
	Object type	Char(10)	*FILE *DTAARA *DTAQ *ALL *NONE
	Option	Char(10)	*CREATE *MOVE *RESTORE *ALLOPR

Figure 4-5 QDFTJRN format extension

The following are important i5/OS red tips (technotes) — journaling tips that you may find beneficial to your operating environment:

- ▶ Journal Standby Mode on IBM i5/OS: When It Makes Sense to Use
- ▶ Journal Caching: Understanding the Risk of Data Loss
- ▶ Journaling: How to View and More Easily Audit Minimized Journal Entries on the IBM System i Platform
- ▶ The Journal Recovery Count: Making It Count on IBM i5/OS
- ▶ Journaling:– Journal Receiver Diet Tip 1: Eliminating Open and Close Journal Entries
- ▶ Journaling at object creation on DB2 for iSeries
- ▶ Journaling: How many journals should I configure?
- ▶ Journaling: *RMVINTENT: The preferred fork in the road for heavy journal traffic
- ▶ Journaling at object creation on DB2 for iSeries
- ▶ Journaling: How can it contribute to Disk Usage Skew?
- ▶ Soft Commit: Worth a Try on IBM i5/OS V5R4
- ▶ Remote Journal on i5/OS: Are You Selecting the Right Type?

The direct link to ITSO red tips or technotes is:

<http://www.redbooks.ibm.com/redbooks.nsf/tips/>

4.5 Performance and performance management enhancements

A number of changes in these areas improve your response times and controls.

4.5.1 Performance enhancements

In this section we discuss performance enhancements.

SQL Query Engine (SQE) enhancements

The modernized query optimizer introduced in V5R2, known as SQE, has several key enhancements with V54.

Support for LIKE, LOB columns, SUBSTR, and sensitive cursors

SQE is a key IBM i5/OS component for continuously improving SQL performance over large sets of data or complex SQL statements. V5R4 SQL workloads will benefit from SQE's new ability to process SQL requests having a LIKE clause or processing LOB columns, the addition of substring (SUBSTR), and sensitive cursors. Cursor sensitivity refers to a cursor setting that defines whether the result set returned by a cursor will include recent changes to the underlying table.

Enhanced partitioned table optimization

A partitioned table created under i5/OS is a database file with multiple members. A partition is the equivalent of a database file member. Therefore, most of the CL commands that are used for members are also valid for each partition of a partitioned table.

Table partitioning enables you to have more data in a table than can be stored in a single member (default for table creation), while enabling the table to appear as a single object for data manipulation operations, such as queries, inserts, updates, and deletes. The partitions inherit the design characteristics of the table on which they are based, including the column names and types, constraints, and triggers.

Table partitioning allows you to have much more data in your tables. Without partitioning, there is a maximum of 4,294,967,288 rows in a table, or a maximum size of 1.7 terabytes. A partitioned table, however, can have many partitions, with each partition able to have the maximum table size.

You must have the DB2 Multisystem feature (option 27 of 5722-SS1) installed under i5/OS to take advantage of partitioned tables support. DB2 Multisystem provides two ways to partition your data:

- ▶ You can create a distributed table to distribute your data across several systems or logical partitions running i5/OS.
- ▶ You can create a partitioned table to partition your data into several members in the same database table on one system or partition.

In both cases, you access the table as if it were not partitioned at all. There are, however, important differences between DB2 Multisystem and partitioning. For more information about partitioned tables, refer to the iSeries Information center at:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp>

Search with *partitioned table*.

Partitioning can also enhance the performance, recoverability, and manageability of your database. Each partition can be saved, restored, exported from, imported to, dropped, or reorganized independently of the other partitions. Additionally, partitioning allows for quickly deleting sets of records grouped in a partition, rather than processing individual rows of a non-partitioned table. Dropping a partition provides significantly better performance than deleting the same rows from a non-partitioned table.

V5R4 partitioned table processing has been enhanced.

Maintained temporary indexes

A significant SQE enhancement in V5R4 is the automatic creation of indexes by the SQE query optimizer to improve performance, with a minimum of additional work by the user. This created index is referred to as a maintained temporary index (MTI) because it is maintained by i5/OS for use by subsequent queries to improve their performance. This MTI may also be referred to in marketing documentation and an *autonomic index*, but should properly be referred to as an MTI.

The SQE optimizer sits in the background and looks for frequently executed queries that could benefit from an index being created. If such a query is identified, then an MTI is created by DB2 UDB. This MTI is created as a temporary index that is maintained by DB2 UDB as the underlying table data is changed similar to a permanent SQL index. The MTI, however, will disappear once the partition is shut down or restarted.

You might want to use one of the other SQL-based interfaces to find the MTIs during run time and save them instead of making the system go through the same autonomic processing after each partition start. You can see all active indexes, including these MTIs, using the **iSeries Navigator Database** → **Schemas** → **Tables** and take one of the options for tables. See also “Index Evaluator” on page 104 for more information.

The maintained temporary indexes function can be disabled with QAQQINI options `CACHE_RESULTS` and `ALLOW_TEMPORARY_INDEXES` set to `*NONE` and `*ONLY_REQUIRED`.

An MTI is not used for statistics and can disappear if all the *owning* plans are removed from the plan cache. MTIs can be viewed with the index evaluator function.

For more information about this support refer to IBM Redbooks Technote *Index Advisor surfaces Maintained Temporary Indexes on DB2 for i5/OS*, TIPS0624.

Faster XML processing

Enhancements in the operating system and in the XML Extenders code have shown between 50% and 90% improvement in performance on some XML-based workloads. Figure 4-6 shows the performance times in seconds.

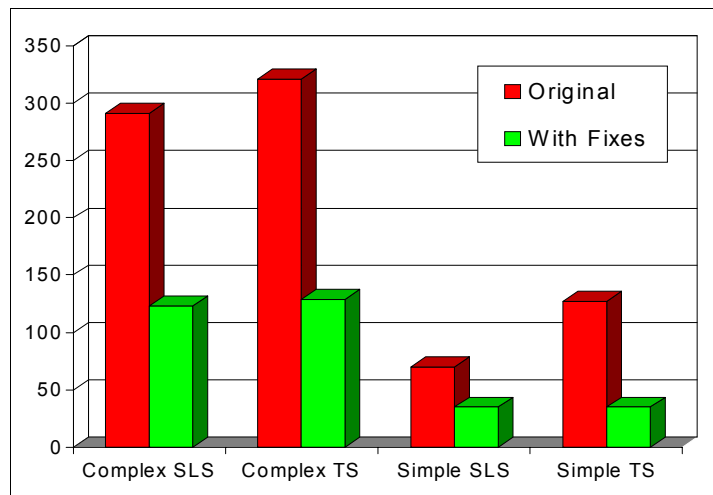


Figure 4-6 Performance times in seconds

In addition, the functionality enhancements allow XML Extenders to be used with more complex XML documents. These advancements include:

- ▶ Ability to generate XML documents with element names being shared.
- ▶ Ability to shred documents containing wrapper elements.
- ▶ Increased scalability.
- ▶ Maximum size XML Document 100 MB (up from 15 MB).
- ▶ Extenders can now insert up to 512,000 rows into a table.

Transaction Manager enhancements

You can improve the performance of database transactions with soft commit. Soft commit limits the number of times that the system writes journal entries associated with a transaction to disk.

Soft commit can improve transaction performance, but it might cause one or more transactions to be lost in the event of a system or partition abnormal termination. Traditional commitment control on DB2 Universal Database for iSeries ensures transaction durability, which means that when a transaction has been committed, the transaction persists on the system, which means that the changes have been formally recorded on disk storage. Soft commit does not provide this durability, although it still ensures the atomicity of the transaction. In other words, the system guarantees a commit boundary, but one or more complete transactions might be lost in the event of a system failure.

To use soft commit, both for a particular job or across the system, specify *NO on the QIBM_TN_COMMIT_DURABLE environment variable. You can change this variable with the Add Environment Variable (ADDENVVAR) command.

For more information about this support refer to IBM Redbooks Technote *Soft Commit: Worth a try on IBM DB2 for i5/OS*, TIPS0623.

Parallel rollback unlock

iSeries customers that make heavy use of commitment control will also be able to recover faster with the parallel rollback option and the ability to customize commitment control disk behavior to get the best performance for their business.

4.5.2 Enhanced performance tooling

In this section we discuss enhanced performance tooling.

Controlling DB2 SMP parallelism via SQL

The CURRENT DEGREE special register specifies the degree of I/O or Symmetric Multi Processing (SMP) parallelism for the execution of queries, index creates, index rebuilds, index maintenance, and reorganizes. CURRENT DEGREE affects static and dynamic SQL statements. The data type of the register is CHAR(5).

Use of parallel processing and the number of tasks used is determined based on the number of processors available in the system, this job's share of the amount of active memory available in the pool in which the job is run, and whether the expected elapsed time for the operation is limited by CPU processing or I/O resources. The database manager chooses an implementation that minimizes elapsed time based on the job's share of the memory in the pool.

The initial value of CURRENT DEGREE is determined by the current degree in effect from the CHGQRYA CL command, the PARALLEL_DEGREE parameter in the current query options file (QAQQINI), or the QQRYDEGREE system value.

VOLATILE table support

Volatile tables are not a new SQL object type. Volatile is just an attribute that can be specified when creating a normal SQL table. This new attribute indicates to the DB2 UDB query optimizer that the number of rows at runtime can vary widely and tell the optimizer that it should use an access method that performs relatively well, regardless of the table size. Many times this will bias the query optimizer towards performing an index scan over this table when

it is referenced on a query. Work tables that are repeatedly populated and then cleared during a batch process may be good candidates for this new keyword.

New filters on STRDBMON command

Advanced database monitor filtering is one of the key attributes that will be discussed later on with the graphical DB2 On Demand Performance Center that is part of iSeries Navigator. The Start Database Monitor (STRDBMON) system command is the technology behind the graphical toolset, so new parameters had to be added to this underlying support.

These filters include:

- ▶ JOB wildcarding (QZDAS*)
- ▶ RUNTHLD, FTRFILE, FTRUSER, FTRINTNETA

Refer to *SQL Performance Diagnosis on IBM DB2 Universal Database for iSeries*, SG24-6654.

4.5.3 DB2 performance analysis with iSeries Navigator

There are many performance-related enhancements and new tools for a database with V5R4 iSeries Navigator interface. Figure 4-7 shows the iSeries Navigator menu.

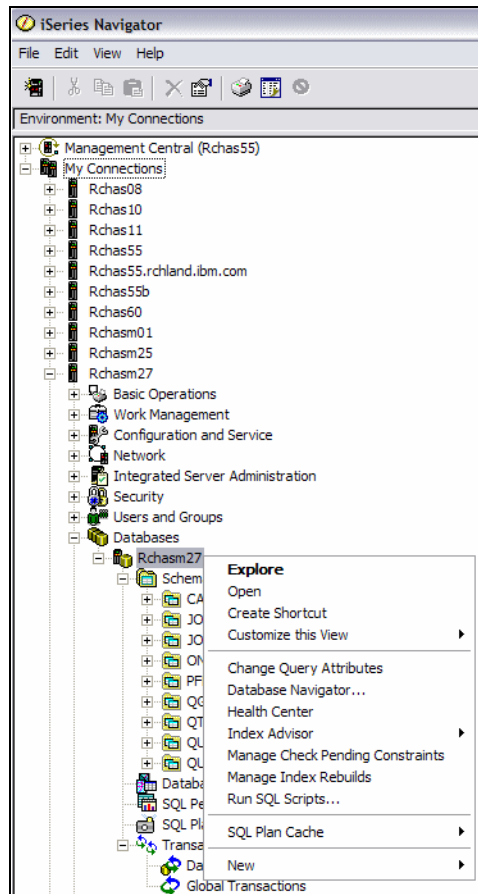


Figure 4-7 iSeries navigator menu

Tools to control resource-intensive queries

In this section we discuss tools to control resource-intensive queries.

Predictive temporary storage governor

DB2 UDB for iSeries has featured a predictive query governor based on execution time since V3R1 with CHGQRYA (Change Query Attributes) command. If the query optimizer estimated that a query would take longer than the defined time limit, an administrator could prevent the a long-running query from taking over the system for an extended period of time. This is a great need when you have users executing *ad hoc* queries.

Another issue caused by *ad hoc* queries results if the database is not properly tuned with the correct indexes. In these cases, the query optimizer may have to create temporary objects such as hash tables to help in the execution of the query. If your server has a lot of memory and temporary storage available, this might not result in a negative performance impact. However, if there are several *ad hoc* queries with this personality running concurrently. It will not be long before the help desk phone is ringing off the hook with users complaining about response time issues.

To help control this, in V5R4 we have added a temporary storage governor via the STORAGE_LIMIT QAQQINI option or the query storage limit parameter on the CHGQRYA command to provide another control for ad hoc queries hogging system resources. Just like the time limit, the query optimizer will estimate how much space will be consumed by the temporary objects it needs to use during the query execution. If the estimated storage exceeds the limit defined for a user, then you can prevent the query from running.

Governor exit point for easier administration

To better automate workflow when the governor storage or time limit is exceed, a new exit point is available in V5R4 for the query governor. In V5R4, the Query Governor Exit Program with Query Governor Exit Point (QIBM_QQQ_QUERY_GOVR) can be used to have a program called every time the governor detects a limit being exceeded. This program could then notify an administrator or keep track of the users that are hitting the limits for future consultation.

The Query Governor can be enabled by setting the QUERY_TIME_LIMIT or STORAGE_LIMIT options in the QAQQINI query options file. It can also be enabled by using the query processing time limit (QRYTIMLMT) or query temporary storage limit (QRYSTGLMT) parameters on the Change Query Attributes (CHGQRYA) CL command. The temporary storage limit feature is new in V5R4.

Tools to manage indexes and MQTs

A materialized query table (MQT) is a database table that provides pre-computed result sets for specified queries, which speeds access to selected data from the larger source tables against which it runs — formally known as automatic summary tables.

Traditionally, automatic summary tables are objects that:

- ▶ Summarize data.
- ▶ Pre-compute query results.
- ▶ Maintain query results that can be used by the optimizer to improve query performance.
- ▶ Are useful in data mining applications.

MQTs are another feature that only can be used by SQE. Just like Instead of Triggers, MQTs were introduced after the GA of V5R3 via PTF and the capabilities have been enhanced in V5R4. Similar to autonomic indexes, they can be used by SQE to automatically improve query performance. Example 4-3 shows a MQT example.

Example 4-3 MQT example

```
CREATE TABLE Example_MQT AS
    (SELECT Geography, Region, Year, Month,
        SUM(Revenue) AS Total_Revenue, SUM(Quantity) AS Total_Quantity,
        COUNT(*) AS Rows_per_Group
    FROM Example_Table
    GROUP BY Geography, Region, Year, Month)
DATA INITIALLY IMMEDIATE
REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER
```

When an SQL query is run, the SQE query optimizer checks to see if there are any MQTs on the system that either match the query being run or partially match. When the query optimizer does find a match or partial match, it will rewrite the query and just use the contents of the MQT as a shortcut to speed up the delivery of the query results. MQTs are most beneficial in data warehousing and business intelligence environments since the data does not change as frequently as a transactional system.

Two entries in the QAQQINI properties file allow the management of the utilization of these tables by the optimizer:

- ▶ `MATERIALIZED_QUERY_TABLE_REFRESH_AGE`, which can be set to 0 (which equals the `*DEFAULT` setting), `*DEFAULT`, `*ANY`, or a timestamp duration.
- ▶ `MATERIALIZED_QUERY_TABLE_USAGE`, which can be set to `*NONE` (which equals the `*DEFAULT` setting), `*ALL`, or `*USER`.

Optimizer could use the MQT instead of fully executing the query, as shown in Example 4-4.

Example 4-4 Optimizer using MQT

```
SELECT Geography, Year,
    SUM(Revenue) AS Total_Revenue, SUM(Quantity) AS Total_Quantity,
    FROM Example_Table WHERE Year IN (2004, 2005)
    GROUP BY Geography, Year;
```

Users retain responsibility for maintaining data currency and activating optimizer awareness.

Reference the white paper on this subject at:

http://www-03.ibm.com/servers/enable/site/education/abstracts/438a_abs.html

Figure 4-8 shows the create table command.

```
CREATE TABLE mqt1 ( c1 , c2 , c3 )
AS (
SELECT itemkey, storekey, COUNT(*) AS cnt
FROM fact
GROUP BY itemkey, storekey )
DATA INITIALLY DEFERRED
REFRESH DEFERRED
MAINTAINED BY USER
ENABLE QUERY OPTIMIZATION

SELECT itemdim.name, storedim.name,
COUNT(*) AS numsold, SUM(itemdim.price) AS sales
FROM itemdim, storedim, fact
WHERE itemdim.itemkey = fact.itemkey AND
      storedim.storekey = fact.storekey AND
      storedim.storekey = fact.storekey
GROUP BY itemdim.name, storedim.name
```

Figure 4-8 Create table command

Figure 4-9 shows the enhanced table menu options.




Figure 4-9 Enhanced table menu options

Index Evaluator

This is also commonly referred to as *show indexes*, which has been enhanced with V5R4.

Before V5R3, it was difficult to determine which of your indexes were unnecessary. With the changes made in V5R3 and V5R4, you can now easily find all indexes and retrieve statistics on index usage. To take advantage of this new feature, you must have the program temporary fix (PTF) SF99503 Version 4 applied. This feature requires that your iSeries Access for

Windows is at V5R3 with iSeries Access PTF number SI15176 installed or V5R4 iSeries Access for Windows.

1. To access this through the iSeries Navigator, navigate to **Database** → **Schemas** → **Tables**.
2. Right-click your table and select **Show Indexes** (see Figure 4-9 on page 104). Indexes include:
 - Last Query Use (V5R3): States the timestamp when the index was last used to retrieve data for a query.
 - Last Query Statistic Use (V5R3): States the timestamp when the index was last used to provide statistical information.
 - Query Use Count (V5R3): Lists the number of instances the index was used in a query.
 - Query Statistics Use (V5R3): Lists the number of instances the index was used for statistical information.
 - Last Used Date (V5R4): The century and date this index was last used.
 - Days Used Count (V5R4): The number of days the index was used. If the index does not have a last-used date, the count is 0.
 - Date Reset Days Used Count (V5R4): The date that the days used count was last reset. You can reset the days used by the Change Object Description (CHGOBJD) command.

Figure 4-10 shows an example. In the top window we select **Show Indexes** for table **ORDERS**. The lower windows show three windows of index information in **1**, **2**, and **3**.

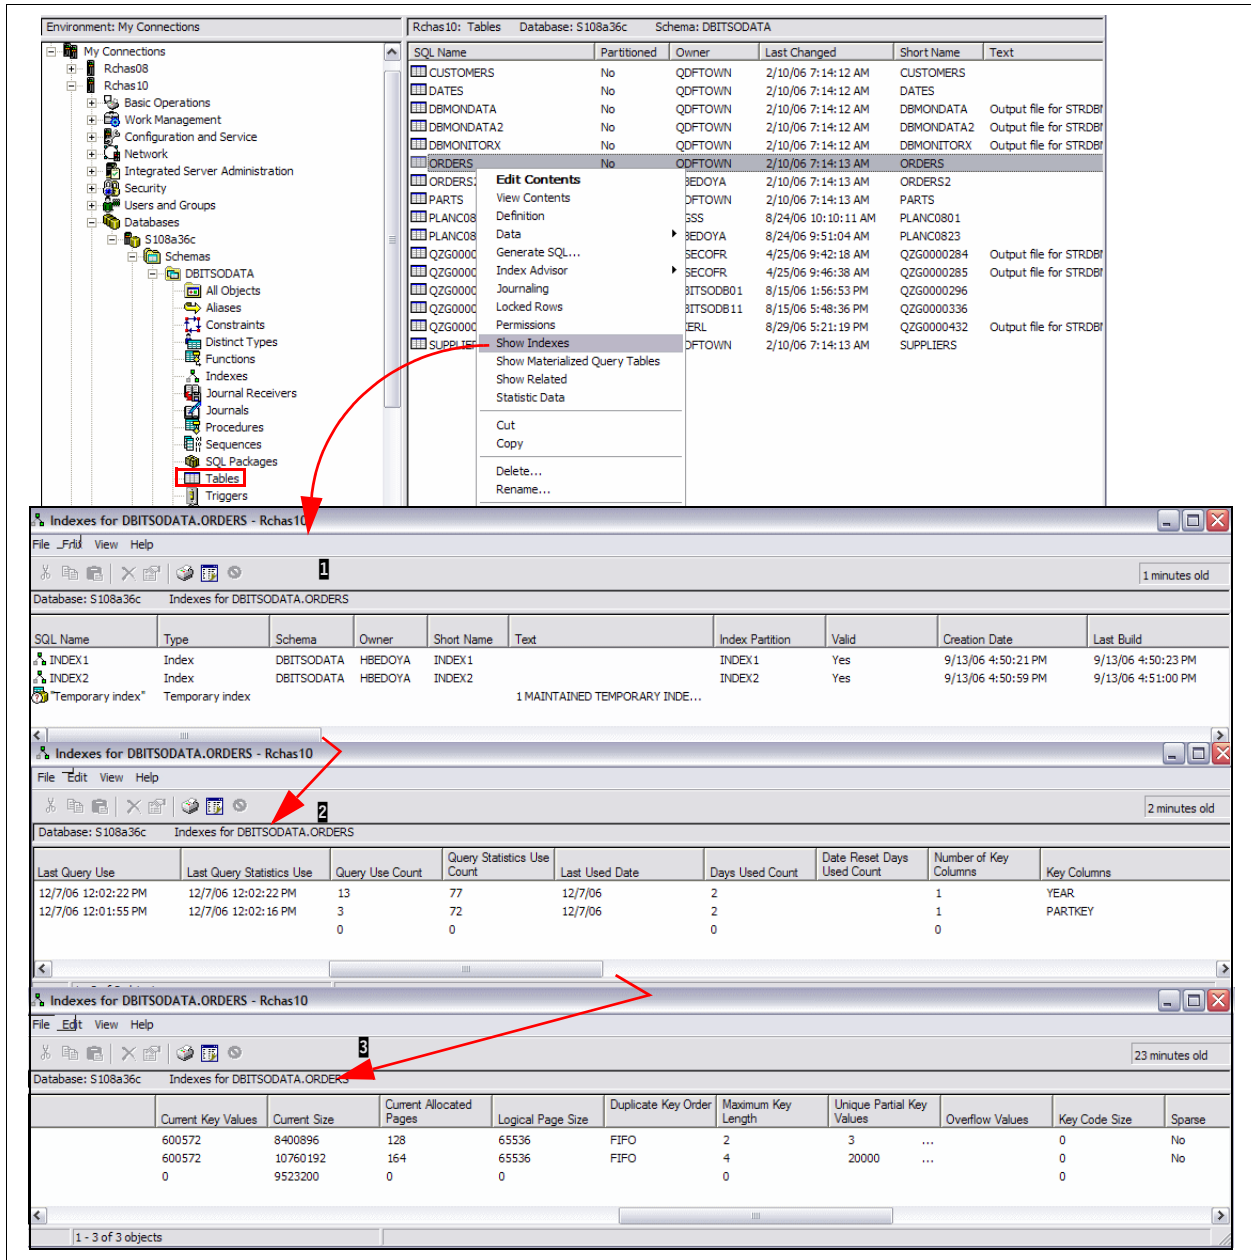


Figure 4-10 Index Evaluator example

As you can see by the position of the scroll left and right bar, more information exists to the right of the last window (B) we show here.

Materialized Query Table Evaluator (new)

This is also referred to as Show Materialized Query Tables.

You can now easily find all Materialized Query Tables (MQTs) and retrieve statistics on MQT usage as a result of iSeries Navigator and i5/OS functionality.

1. Similar to the Index Evaluator interface, to access this through the iSeries Navigator, navigate to **Database** → **Schemas** → **Tables**.

2. Right-click your table and select **Show Materialized Query Tables**.

In addition to all existing attributes of an MQT, two new fields have been added to the V5R4 iSeries Navigator. Those new columns show you statistics on MQT usage in a query:

- ▶ **Last Query Use:** This parameter lists the timestamp when the MQT was last used by the optimizer to replace user-specified tables in a query.
- ▶ **Query Use Count:** This parameter lists the number of instances the MQT was used by the optimizer to replace user-specified tables in a query.

MQT tables must be optimized just like non-MQT tables. Indexes should be created over the MQT columns that are used for selection, join, and grouping as appropriate. Column statistics are collected for MQT tables.

You can find out more information with iSeries Information Center (by selecting **Database** → **Performance and query optimization** → **Processing queries** → **Materialized query table optimization**).

Manage index rebuilds (new)

Through V5R3, you can view a list of access paths that are rebuilding and either hold the access path rebuild or change the priority of a rebuild with EDTRBDAP CL command.

Starting with V5R4, you can now manage the rebuild of your indexes using iSeries Navigator:

1. To display access paths to rebuild with the iSeries Navigator, expand **Databases**.
2. Right-click the database that you want to work with and select **Manage index rebuilds**.

With this function, you can get more information about why an index is being rebuilt and why it was invalidated. (The Edit Rebuild of Access Paths (EDTRBDAP) CL command does not show this information.)

- **Rebuild reason:** Reason why this access path needs to be rebuilt.
- **Rebuild reason Type:** Reason type why this access path needs to be rebuilt.
- **Invalidate reason:** Reason why this access path was invalidated.
- **Invalidation reason Type:** Reason type why this access path was invalidated.

Index Advisor (new)

The database component of iSeries Navigator is the major driver of DB2 for i5/OS simplification in V5R4 with the addition of new tools and streamlining of existing tools. This new combination of tools is known as the DB2 On Demand Performance Center.

The system-wide Index Advisor is one of the new tools in V5R4. Index advice from the DB2 UDB query optimizer is not a new capability in V5R4. Feedback on suggested indexes in past releases could be found in optimizer debug messages within a joblog or a database monitor collection. However, this feedback was only available if someone manually started a database monitor or turned on debug messaging. Even if this was done, an analyst had to deal with the complexities of extracting the advised index details of various joblogs and database monitor files.

The new Index Advisor eliminates these hurdles with a click of mouse. No user intervention such as starting a database monitor is needed since the query optimizer automatically logs index advice for all queries to a repository in V5R4.

Right-click your database name in the iSeries Navigator tree and select the **Index Advisor** task, and you are given the output displayed Figure 4-11, viewable at database, schema, or table level.

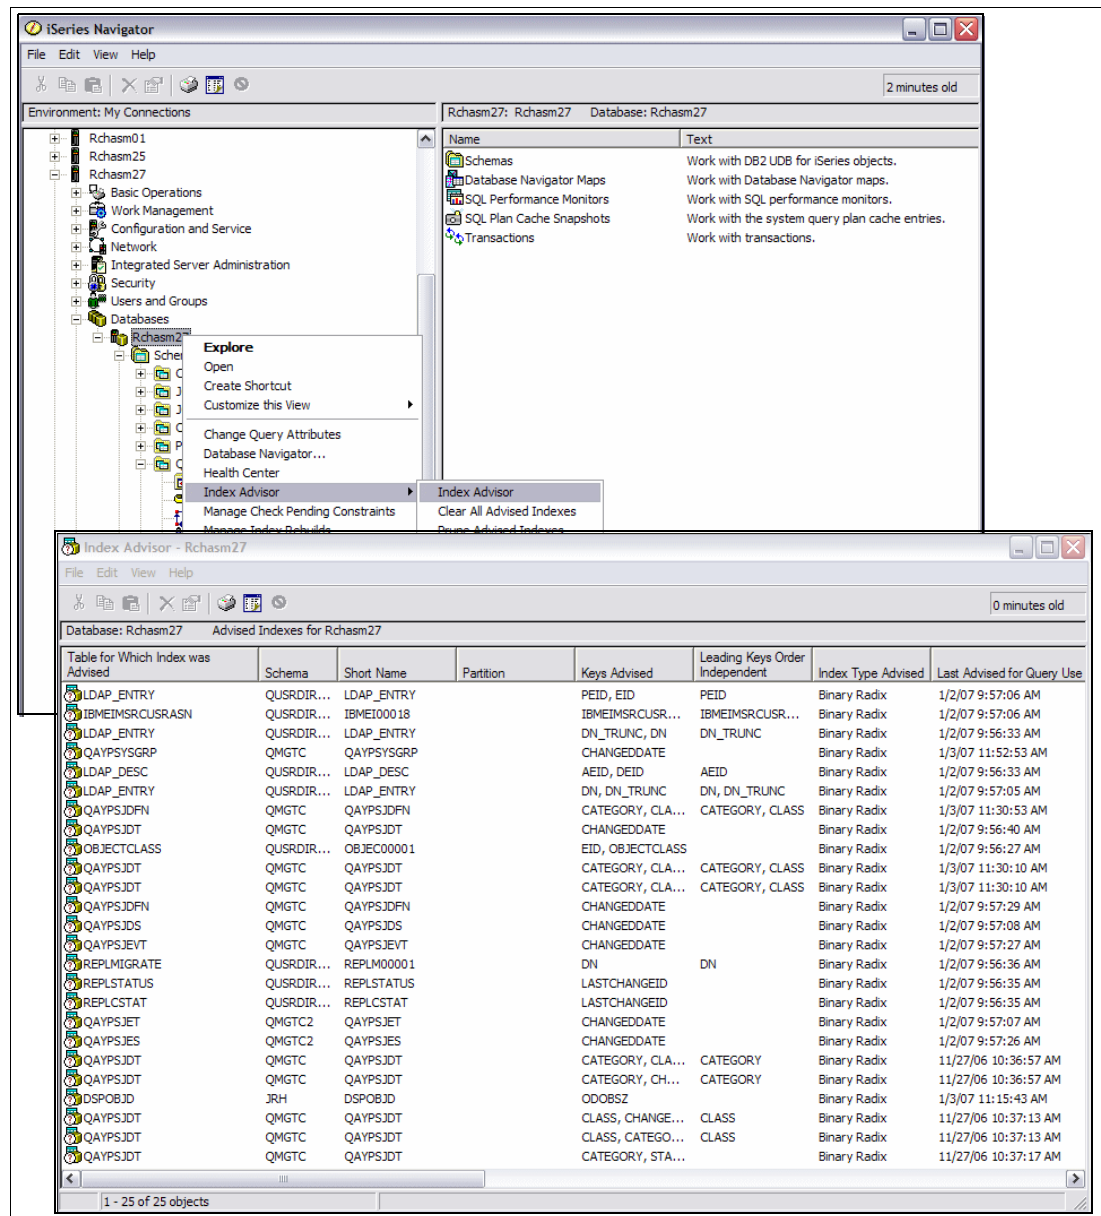


Figure 4-11 Index Advisor

Even better is the fact that index advice provided in V5R4 by the SQL Query Engine (SQE) query optimizer is more intelligent and complete. The index advice given in past releases only focused on filtering criteria of a query, requiring manual intervention from an analyst.

While SQE still does not advise indexes for IN predicates or search conditions that contain an OR, it does include join, grouping, and ordering criteria. Look closely at the Index Advisor output and you will also see that the optimizer provides advice on the type of index to create. Some queries may benefit from a traditional radix index while others would benefit from an encoded vector index structure. Again, index advice can be done at database, schema, and table level. iSeries Navigator displays information found in the QSYS2/SYSIXADV system table.

Note: This functionality is also available for V5R3 with the latest DB Group PTF, SF99503, level 9 or later.

4.5.4 Tools to monitor and analyze SQL performance

In this section we discuss tools to monitor and analyze SQL performance.

Database performance monitor

Database monitors have existed on DB2 UDB for iSeries for many releases, but the usability of these monitors takes a big step forward in V5R4 with the addition of both pre-filters and post-filters. These filters are very similar to the filters that were available with the SQL Plan Cache viewer. Database monitor filters are important for several reasons.

The filters can minimize the overhead and disk space consumed by a database monitor collection. As a case in point, a customer recently ran a database monitor collection on all jobs on their system for 15 minutes. An SQL-based ERP application that was running on the system at the time resulted in almost 3 million rows of detailed monitor data.

Applying one of the filters available in V5R4, such as only collecting monitor data for a certain user, or only for SQL statements referencing a specified table, could have significantly reduced the amount of data that has to be collected by DB2 UDB. These filters are made available when you start a new SQL Performance Monitor collection via the graphical iSeries Navigator interface or the STRDBMON CL command.

Option for pre-filtering

You can start this monitor by right-clicking **SQL Performance Monitors** under the Database of the iSeries Navigator and selecting **New** → **SQL Performance Monitor**. On the monitor wizard, select **Detailed**, which is the default option.

Starting with V5R4, when you create a detailed monitor, you can filter the information that you want to capture. Key filter criteria include:

- ▶ **Minimum estimated query runtime:** Select this to include queries that exceed a specified amount of time. Select a number and then a unit of time.
- ▶ **Job name:** Select this to filter by a specific job name. Specify a job name in the field. You can specify the entire ID or use a wildcard. For example, QZDAS* will find all jobs where the name starts with QZDAS.
- ▶ **Job user:** Select this to filter by a job user. Specify a user ID in the field. You can specify the entire ID or use a wildcard. For example, QUSER* will find all user IDs where the JOB name starts with QUSER.
- ▶ **Current user:** Select this to filter by the current user of the job. Specify a user ID in the field. You can specify the entire ID or use a wildcard. For example, AS03* will find all users where the job's current user starts with AS0301.
- ▶ **Internet address:** Select this to filter by Internet access. The format must be xxx.xxx.xxx.xxx, for example, 5.5.199.199.
- ▶ **Only queries that access a list of tables:** Select this to filter by only queries that use certain tables. Click **Browse** to select tables to include. To remove a table from the list, select the table and click **Remove**. A maximum of ten table names can be specified.

- ▶ Activity to monitor: Select to collect monitor output for user-generated queries or for both user-generated and system-generated queries.

Option for post filtering

You can view SQL statements that are included in a detailed monitor. Right-click any detailed monitor in the SQL performance monitor window and select **Show statements**. The filtering options provide a way to focus in on a particular area of interest, including:

- ▶ Minimum runtime for the longest execution: Filter those queries with at least one long individual query instance runtime.
- ▶ Queries run after this date and time: Filters those queries that have been run recently.
- ▶ Queries that use or reference these objects: Filters those queries that reference or use the tables or indexes specified.
- ▶ SQL statement contains: This is a wildcard search capability on the SQL text itself. It is useful for finding particular types of queries. For example, queries with a FETCH FIRST clause can be found by specifying fetch. The search is case insensitive for ease of use.

Analyze monitor enhancements

The filters discussed in the prior section eliminate the number of rows in a database monitor collection, but there are still hundreds of fields in a single database monitor row and many monitor row types that need to be sorted through to identify and fix performance problems. This maze of data made it a challenge to quickly look at a database monitor collection and determine whether there were any issues, unless you were a DB2 performance expert.

A dashboard summary was added in V5R4 to address this issue. On the first analysis of a database monitor collection, DB2 for iSeries will automatically summarize some of the key performance indicators in the collection and then present the results in the summary window depicted at the top of this chart. This summary will allow an analyst to quickly scroll through some high-level DB2 performance indicators to determine whether more detailed analysis is needed.

Figure 4-12 shows a dashboard and drill-through analysis.

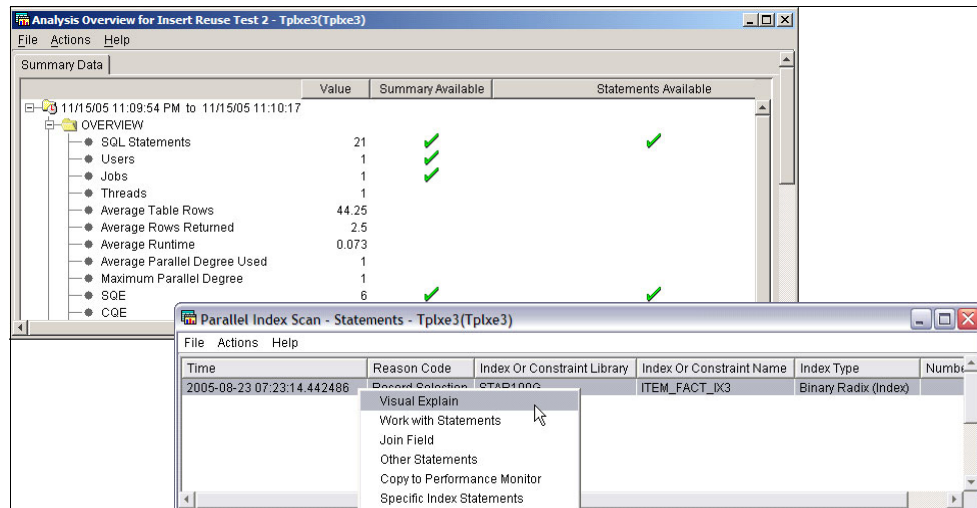


Figure 4-12 Dashboard and drill-through analysis

V5R4 also addresses another issue with the database monitor analysis reports by providing drill-through analysis. The reports could be used to find an SQL statement with performance issues. However, if more detailed research of this statement was needed with a tool such as

Visual Explain, the user had to exit out of the tool and then manually copy that SQL statement for Visual Explain analysis. As you can see with the bottom figure in Figure 4-12 on page 110, an analyst can easily right-click to drill into more detail with Visual Explain. This will be a great enhancement for improving the efficiency of your database analysts and administrators.

- ▶ To analyze your monitor data, right-click a monitor and select **Analyze**. The monitor does not need to be ended in order to view this information. On the Analysis Overview dialog, you can view new overview information or choose one of the following categories:
 - How much work was requested?
 - What options were provided to the optimizer?
 - What implementations did the optimizer use?
 - What types of SQL statements were requested?
 - Miscellaneous information.
 - I/O information.
 - From the Actions menu, you can choose predefined summary reports.
- ▶ In addition, when a green check is displayed under the Summary column, you can select that row and click **Summary** to view information about that row type.
- ▶ Click **Help** for more information about the summary report.
- ▶ To view information organized by statements, click **Statements**.

Compare monitors: new function in V5R4

Often times the performance of an application can change over time as the server configuration changes, new versions are installed, or thousands of records are added. Sometimes changes have an unexpected negative impact on performance.

With the new monitor comparison utility, users can utilize database monitor collections to help determine what change is causing performance issues. This technique requires that a customer collects and saves database monitor collections when a critical DB2 workload or report is performing and running well on their server.

A best practice is to compare snapshots:

- ▶ When performance is acceptable, create a plan cache snapshot.
- ▶ If you experience performance degradation, create a new plan cache snapshot.
- ▶ Run the compare function to identify statements that are running slower than before.
- ▶ Identify these statements and do additional problem determination.

The following is a way to perform the comparison of monitor output:

1. To compare data sets in different monitors, go to **iSeries Navigator** → **system name** → **SQL performance monitors**. Right-click a monitor in the right pane and select **Compare**.
2. On the Compare dialog, you can specify information about the data sets that you want to compare. When you click **Compare**, both monitors are scanned for matching statements. Any matches found will be displayed side-by-side for comparison of key attributes of each implementation.
3. On the Comparison output dialog, you view statements that are included in the monitor by clicking **Show Statements**. You can also run Visual Explain by selecting a statement and clicking **Visual Explain**.

It is important to note that the comparison function could be a long-running process with large database monitor collections, so plan accordingly.

SQL Plan Cache

The SQL Plan Cache is maintained by the SQE optimizer. It is an internal repository used by SQE to store the access plans and associated statistics for SQL statements that are currently running or recently have been executed on your server. The big benefit with this Plan Cache analysis toolset is that for the first time you can perform detailed SQL performance analysis without the overhead of running a database monitor.

New in V5R4 is the capability to *see* its contents.

A detailed database monitor collection can use up disk space in a hurry and noticeably slow your disk response times. Gone are the days when a user calls about poor performance and your only response for an SQL-based application was starting a database monitor and having the user rerun the report or application. Now, when a performance issue is raised on V5R4, you can analyze the contents of the SQL Plan Cache to see if a long-running SQL statement is the cause of the performance problems.

Using the SQL Plan Cache tool is as simple as right-clicking the **SQL Plan Cache Snapshots** component and selecting the **SQL Plan Cache → Show Statements** task. This action will produce the screen shown in Figure 4-13.

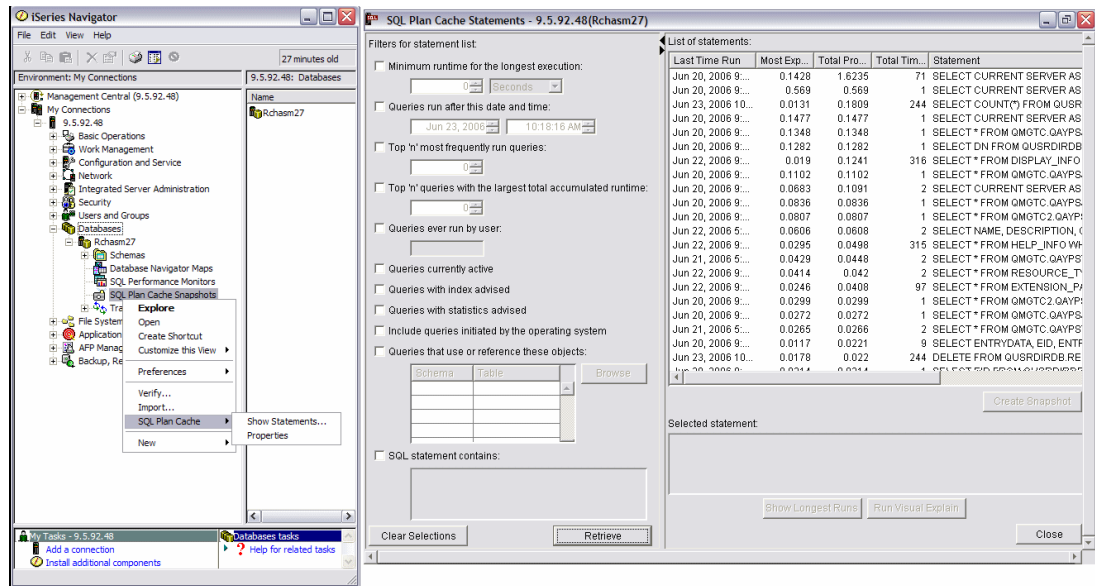


Figure 4-13 SQL Plan Cache

One feature of this Plan Cache viewer not to miss is the statement filters that are available on the left-hand side of the dialog. These filters allow you to intelligently subset the SQL statements that you want to analyze, as opposed to scrolling through a random list of thousands of SQL statements. For instance, you can pick the top 10 statements that are executed most often or take the most time. You can also choose to only analyze the statements currently running or the SQL statements run by a specific user. Furthermore, once you have selected an SQL statement to analyze you can seamlessly launch Visual Explain for that statement from the same window.

You will need to remember that the SQL Plan Cache is fixed in size, so some SQL statements may get purged out of the cache if they have not been recently executed. In these situations, the database monitor tool will have to be used.

If you want to archive the contents of the cache into a permanent object, you can create an SQL Plan Cache Snapshot™. Effectively, this operation takes the contents of the SQL plan

cache and converts it into a database monitor output file. Once the snapshot is created, you can run reports to analyze the contents just like a database monitor collection.

Statistics requests enhancements

Use the Statistics Requests dialog to control the system-wide statistics collection processes. Both user-created and system-generated statistics collections are available. With this dialog, you can:

- ▶ Stop statistics from being collected for a table.
- ▶ Display statistics request details and view the error data for a request.
- ▶ Request that statistics for a table be collected immediately.
- ▶ View and change the background collection status for all background requests.
- ▶ Refresh the list of statistics requests.

4.5.5 Tools to show current activity

In this section we discuss tools to show current activity.

Statistics requests (enhanced)

Use the Statistics Requests dialog to control the system-wide statistics collection processes. Both user-created and system-generated statistics collections are available. With this dialog you can:

- ▶ Stop statistics from being collected for a table.
- ▶ Display statistics request details and view the error data for a request.
- ▶ Request that statistics for a table be collected immediately.
- ▶ View and change the background collection status for all background requests.
- ▶ Refresh the list of statistics requests.

Current SQL for a job (enhanced)

The Current SQL dialog displays the active jobs on your system. You can then select a job and display the SQL currently being run, if any.

Simply select a job from the list and click **SQL Statements**. If the job contains SQL statements, the statement is shown in the SQL statement for select job display. When displayed, you can select the statement and:

- ▶ Click **Edit SQL** to edit the SQL used for the job.
- ▶ Click **Visual Explain** to have the statement displayed in Visual Explain.
- ▶ Click **Job Log** to view the job log.
- ▶ Click **End Job** to end the job.

In V5R4, the information related with SQL cursor open has been added. For active jobs list, not all columns are displayed. To display different columns for the active jobs, click **Columns**.

4.5.6 Tools to show current database limits and sizes

DB2 for i5/OS V5R4 provides an overview of the database objects within your i5/OS partition databases, using the Health Center function for your specific database.

Health Center

The Health Center allows you to monitor and perform a checkup on how many different database objects you have to assist you in managing them. You can also see associated values so you can compare them to different DB2 UDB limits such as the maximum number of rows allowed in a table or maximum size of an index object.

Additionally, you have the ability to change the thresholds of when to mark a DB2 object as being close to a limit and the ability to save these checkups to a permanent file to enable historical comparisons. This feature is launched by right-clicking your named database and selecting the **Health Center** task.

Information is presented under three sets (tabs) of information:

- ▶ Overview
- ▶ Size Limits
- ▶ Design Limits

In Figure 4-14 we show excerpts of two of the sets of information: Overview (1) and Size Limits (2). This is a good way to find out what database objects you have and how many of each type.

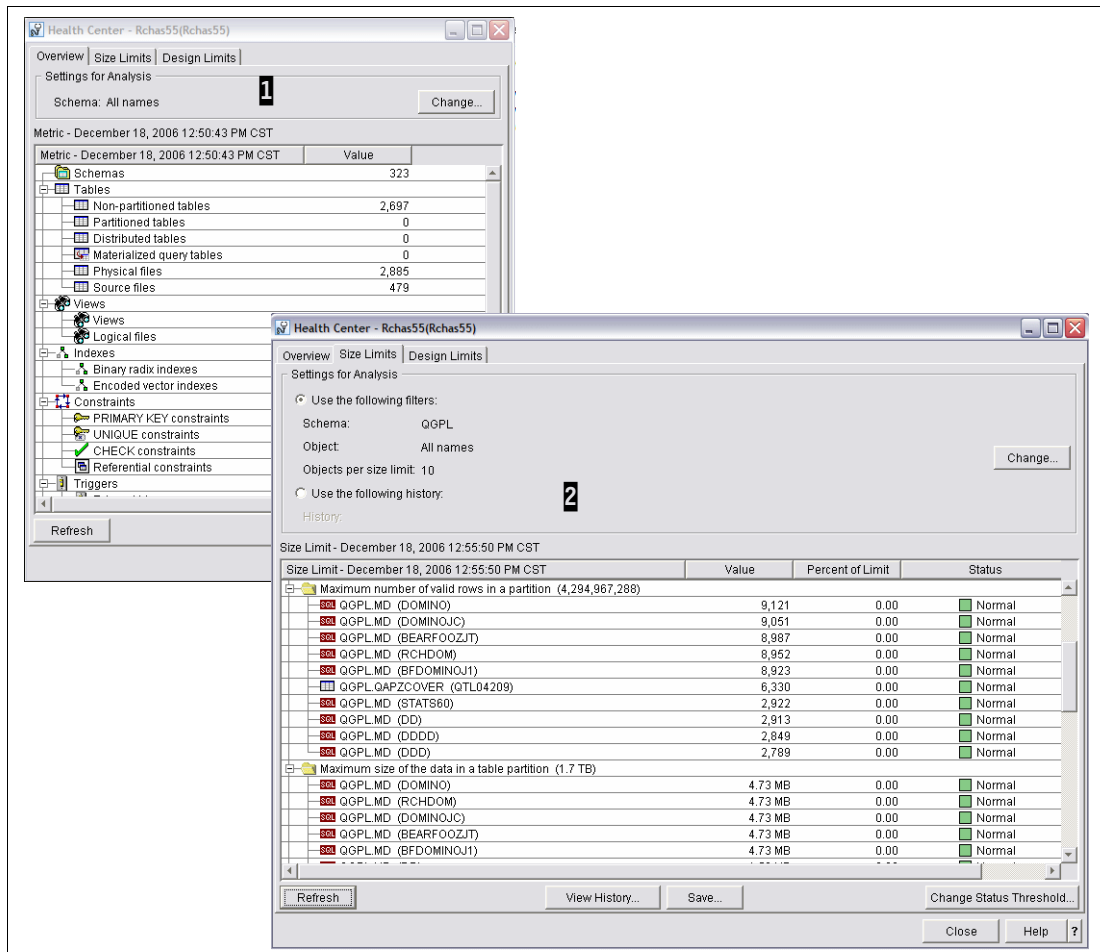


Figure 4-14 Health Check example: overview and size limits

4.5.7 Content Manager

IBM Content Manager software captures, stores, manages, integrates, and delivers all forms of digital content across a company's value chain, from employees to customers to suppliers and partners. It is intended for new customers and not to provide a migration from CM for iSeries V5.3.

DB2 Content Manager Standard Edition V8.3 provides an entry-level content management solution for small and medium-sized businesses. It provides a single, open, comprehensive, and consistent framework to manage, share, reuse, and archive all types of digitized content.

Content types include HTML and XML Web content, document images, electronic office documents, and rich media. Standard Edition's tools include a Windows-based eClient and administrative clients for desktop or browser access and a single GUI interface to manage the system, servers to provide indexing, storage, retrieval, and archiving services to manage all types of digitized content and document routing and workflow options that support a formal, consistent, and repeatable document routing process.

Figure 4-15 depicts Content Manager functions and user interface.

Storage Options & DB2 CM	Disk	Optical	DR550
DB2 Content Manager for iSeries			
V5.3	X	X	X (MBS)
Standard Edition 8.3	X	X (MBS)	X (MBS)
Content Manager OnDemand for iSeries			
Spool File Archive	X	X	
Common Server	X	X	Planned

MBS TECHNOLOGIES

For more information, see <http://www-306.ibm.com/software/data/cm/cmgr/mp/edition-standard.html>

Figure 4-15 Content Manager

DB2 Content Manager Standard Edition is initially available on the iSeries server platform for i5/OS. Although Standard Edition shares a common code base and Windows-based management client with Enterprise Edition, it does not provide the full range of features for larger enterprises, such as Hierarchical Storage Management, migration services (from DB2 CM 5.3), full text search, and federated services (including advanced workflow).

Since Standard Edition does not support HSM to optical or DR550 devices, customers may consider StorageView from MBS Technologies, which enables support for DR550 and any other TSM device to DB2 CM Standard Edition (and CM iSeries v5.3).

4.6 Additional information

See the following for additional information:

- ▶ DB2 UDB for iSeries home page
<http://www-03.ibm.com/servers/eserver/iseries/db2/>
Select the **News and announcements** link.
- ▶ USENET
<http://www.news2mail.com/comp/sys/ibm/as400/misc.html>
<http://www.news2mail.com/comp/databases/ibm-db2.html>
- ▶ iSeries Network SQL and DB2 Forum
<http://www.iseriesnetwork.com/Forums/main.cfm?CFApp=59>
- ▶ Education Resources - Classroom and Online
<http://www.iseries.ibm.com/db2/gettingstarted.html>
<http://ibm.com/servers/enable/site/education/ibo/view.html?oc#db2>
- ▶ DB2 UDB for iSeries publications
 - White papers
<http://www-03.ibm.com/servers/enable/site/education/ibo/view.html?wp>
 - Online manuals
<http://www-03.ibm.com/servers/eserver/iseries/db2/books.html>
 - Porting help
<http://www-03.ibm.com/servers/enable/site/db2/porting.html>
 - DB2 UDB for iSeries IBM Redbooks
<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=db2>
 - *Stored Procedures, Triggers, and User-Defined Functions on DB2 Universal Database for iSeries*, SG24-6503
<http://www.redbooks.ibm.com/abstracts/sg246503.html?open>
 - *Preparing for and Tuning the SQL Query Engine on DB2 for i5/OS*, SG24-6598
<http://www.redbooks.ibm.com/abstracts/sg246598.html?open>
 - *Modernizing IBM eServer iSeries Application Data Access - A Roadmap Cornerstone*, SG24-6393
<http://www.redbooks.ibm.com/abstracts/sg246393.html?open>



i5/OS integration of IBM System x-base servers

This chapter discusses the i5/OS V5R4 System x-based server integration enhancements in general and specifically provides an overview of the V5R4 support on IBM System i5 models for attaching a wider array of the IBM System x family of systems, including BladeCenter models. This iSCSI attachment is via the new in V5R4 iSCSI support.

In this book, unless a specific technology is called out, we use the term *System x* to represent i5/OS integrated server support of xSeries, System x, and BladeCenter technologies.

5.1 IBM System x integration enhancements

The System i platform is designed to combine operating systems, storage, security, middleware, and database into one single integrated system. This means that instead of supporting and maintaining multiple IBM xSeries and System x models (and with i5/OS V5R4, BladeCenter models), you can integrate the operations, networking, and server management into one or more partitions on a System i model.

V5R4 enhancements build on top of the long history of success stories with this integration support from previous i5/OS releases. V5R4 provides a moderate set of enhancements for all previously supported xSeries and System x models using the previously available Integrated xSeries Server (IXS) hardware, and the Integrated xSeries Adapter (IXA) over an HSL loop. V5R4 support is now expanded to include integrated support over new Internet Small Computer System Interface (iSCSI) hardware and Ethernet cabling and switches.

iSCSI support is available only on the POWER5-based System i5 models 520, 550, 570, and 595, which must be running i5/OS V5R4.

The following sections highlight general V5R4 enhancements and then focus on the new iSCSI support. Note that the new iSCSI support configuration offers a wider number and range of processor capacities of integrated servers than is available through IXS and IXA hardware configurations. iSCSI offers the advantages in expanded connectivity and more flexibility in spreading virtual disk activity over multiple connections. This can improve performance throughput in heavy virtual disk I/O environments.

The iSCSI configuration is, in general, more complex to set up than setting up one or a small number of integrated servers using IXS or IXA configurations. This chapter provides an overview of iSCSI configuration requirements. You should review documentation we reference in order to make the right choice of using IXS, IXA, iSCSI, or combinations of these options for your production environment with best performance in mind.

Note: Since this an overview document, you are referred to the following documentation and Web sites that contain important planning, configuration, and performance details and considerations when using i5/OS-base integrated server support over IXS, IXA, and iSCSI attachment to an i5/OS partition:

- ▶ The primary Web site for IBM System i integration of xSeries, System x, and BladeCenter models running supported releases of Windows operating systems and Linux releases, at:

<http://www.ibm.com/systems/i/bladecenter>

Refer to this Web site when you want to know the current supported IBM xSeries, System x, and BladeCenter systems, switches supporting iSCSI, and supported software release levels, and the latest PTFs for this integration of servers running Windows operating systems and Linux distributions.

By February 2007 a group PTF will be available for IBM System x-based integration on the System i platform. It is SF99348 and titled "System i integration with BladeCenter and System x." A direct link to PTFs for this support is:

<http://www.ibm.com/systems/i/bladecenter/ptfs.html>

There is also a white paper describing how the combination of existing System i technologies and service offerings can provide additional value to clients by automating a high-availability environment for an integrated server or AIX 5L environment hosted in an Independent Auxiliary Storage Pool (IASP). Two systems are used with the IASP on both systems residing in an IBM System Storage™ DS6000™. Data replication of the IASP is provided by Metro Mirror (formerly Synchronous PPRC, Peer-to-Peer Remote Copy), an IBM System Storage Remote Mirror and Copy (RMC) solution. Copy Services for System i provides i5/OS command-line management of Metro Mirror and HA switch-hitter automation.

The example solution demonstrated in the paper illustrates two of the core facets of HA: making business data available on a backup site and restarting the AIX 5L partition or integrated server on the backup site.

The paper title is *High Availability Automation for iSCSI Integrated Server and AIX 5L with Virtual Storage on System i*.

- ▶ V5R4 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/index.jsp>

Expand the **Integrated operating environment** link.

- ▶ The Performance Capabilities Reference manual (PDF), dated May 2006 or later, at:

<http://www.ibm.com/eserver/iseriess/perfmgmt>

Select the **Resource Library** link. This document has a chapter devoted to IXS, IXA, and iSCSI performance considerations.

- ▶ IBM Redbook *Implementing Integrated Windows Server through iSCSI to System i5*, SG24-72300

This book includes all V5R4 i5/OS level Integrated Server Support capabilities. This includes the general enhancements discussed in this chapter and significant content on the new iSCSI support. You can search the IBM Redbooks Web site, using this manual number to access this book:

<http://www.ibm.com/redbooks>

5.1.1 Support for integrated Intel servers

i5/OS V5R4 offers changes in the support for the integrated Intel servers (IXS).

Changes in V5R4 include licensed program product packaging of Linux, the replacement of the Windows integration server licensed program product, new extending disk storage space and Windows 2004 Volume Shadow Copy support, and withdrawal of support for some xSeries technology models.

5.1.2 Licensed program product changes

Licensed program (LPP) 5722-WSV iSeries integration for Windows Server® is replaced in V5R4 by Option 29 of i5/OS LPP 5722-SS1 (Integrated Server Support).

In i5/OS V5R4 support for Linux is now available as a licensed product (Licensed Program Product (LPP) 5722-LSV, Integration for Linux on xSeries).

For i5/OS V5R4, a new i5/OS infrastructure for Linux on integrated servers is provided. In V5R2 and V5R3, the i5/OS infrastructure for Windows is used for Linux servers.

For i5/OS V5R2 and V5R3, PTFs are available to enable running Red Hat Linux on integrated servers. IBM has also made PTFs available that enable running SUSE Linux on System x models that are attached via the IXA.

For the latest list of PTFs that are required for all integrated server support on either V5R4 or earlier releases, refer to:

<http://www.ibm.com/systems/i/bladecenter/ptfs.html>

Attention: A one-time migration of Linux servers that were created on V5R2 or V5R3 is required after upgrading i5/OS to V5R4. See the Migrating Linux servers from V5R2/V5R3 to V5R4 Web page for details:

<http://www.ibm.com/systems/i/bladecenter/linux/v5r4migration.html>

5.1.3 Support withdrawal

Support for the 200 MHz and 333 MHz IBM Integrated PC Server for AS/400 (IPCS) and IBM Integrated Netfinity® Server for AS/400 (INS) is withdrawn. The withdrawn IPCS and INS hardware resource types are 6617 and 2850 with feature codes 2854, 2857, 2865, 2866, 6617, and 6618. Since IPCS and INS servers were the only integrated server types that provided host LAN support (sharing LAN adapters between i5/OS and Windows), the host LAN function has also been withdrawn.

5.1.4 Support for Windows 2003 Volume Shadow Copy

In i5/OS V5R4 support for Windows Server 2003 Volume Shadow Copy Service is added. This file level backup utility is enhanced to be a shadow-copy requestor that requests a snapshot of your Windows volume for backup purposes. This allows data for applications that support the Volume Shadow Copy Service to be backed up without stopping the application, which improves application availability.

5.1.5 Expand disk drive support

In i5/OS V5R4 support has been added to enable you to expand a virtual disk drive (network server storage space) without copying the disk drive. To expand a disk drive, follow these steps:

1. Open iSeries Navigator.
2. Expand **Integrated Server Administration** → **All Virtual Disks**.
3. Select a disk drive from the list available.
4. Right-click the disk drive and select **Properties** or click the appropriate icon on the iSeries Navigator toolbar.
5. Click the **Capacity** tab of the disk drive property sheet.
6. Specify the increased disk size in the New capacity field. See the online help for details on valid disk sizes associated with a particular file system format. The extended portion of the disk will be unpartitioned free space.
7. Click **OK**.

If the disk is linked to an active server, a confirmation panel is shown to indicate that the disk drive will be temporarily unavailable to the server while the disk is being expanded. Click **Change** on the confirmation panel to confirm that this is acceptable, or click **Cancel** on the confirmation panel to cancel the disk expansion operation.

Note: The disk cannot be linked to an active server while it is being expanded. If the server supports dynamic unlinking of disk drives, then the above procedure can be performed while the server is active. In this case, the disk is dynamically unlinked, then expanded and then dynamically relinked to the server again. Therefore, the disk is temporarily unavailable to the active server while the disk is being expanded.

The Windows operating system DISKPART command-line utility can be used to expand an existing partition in order to utilize any additional free space.

Note: DISKPART is available by default on Windows Server 2003. It can also be downloaded from the Microsoft Web site. Also refer to the Microsoft Knowledge base articles for DISKPART for details and limitations. The Microsoft Web site can be found at: <http://www.microsoft.com>

You can also use the i5/OS Integrated Server Support CL command Change Network Storage Space (CHGNWSSTG) to expand the disk storage space. Note that when using CHGNWSSTG to expand the disk, the disk cannot be linked to an active server. CHGNWSSTG will not automatically unlink and relink the disk if the server is active.

5.2 System x and BladeCenter support via iSCSI

The Integrated Server Support (option 29 of i5/OS V5R4) that replaces the Windows operating system integration known in previous releases as iSeries Integration for Windows Server (LPP 5722-WSV) includes new V5R4 support for integrating System x and BladeCenter models via the Internet Small Computer System Interface (iSCSI) protocol and its associated hardware. iSCSi works over Ethernet cabling and appropriate TCP/IP support switches.

iSCSI support includes IXS and IXA capabilities, such as:

- ▶ All disks are virtual disks, enabling storage consolidation
- ▶ Centralized server management
- ▶ i5/OS tape, optical, and virtual Ethernet devices

Additionally, iSCSI attachment offers the following benefits:

- ▶ An extensive scalability range, from connecting many System x-based servers through one iSCSI Host Bus Adapter (HBA) for a lower cost connectivity, to allowing up to a number of iSCSI HBAs per individual System x-based servers for scalable bandwidth.

Each iSCSI HBA can be shared by up to eight integrated servers. In configurations with low to moderate virtual disk I/O rates this can result in a lower cost solution compared to an IXS or IXA option per integrated server.

Each integrated server can access multiple iSCSI HBAs. This enables multiple iSCSI and virtual Ethernet paths between i5/OS partition and the integrated servers. This can sustain higher virtual disk I/O rates.

For specific numbers of servers attachable see the System i integration with BladeCenter and System x Web site at:

<http://www.ibm.com/systems/i/bladecenter/>

- ▶ Capability to specify multiple virtual disk paths that can improve performance of the Windows applications in heavy disk I/O environments.
- ▶ Support extended to a specific set of BladeCenter servers.
- ▶ Lower incremental cost per integrated server.
- ▶ Utilization of TCP/IP skills to install and configure.
- ▶ Leveraging industry standards with iSCSI.
- ▶ Increased flexibility options for redundancy (backing up is easier) and performance.
- ▶ Coexistence with Integrated xSeries Servers and xSeries systems attached via Integrated xSeries Adapters, which can also be used as hot spare components.

While iSCSI offers many benefits, you may find that IXS and IXA configurations are sufficient for your work environment. You may continue to use this support and, where appropriate, selectively use or even mix with new iSCSI configurations. In consolidation scenarios with moderate performance and *cabling connection topology* requirements, you will probably find these *consolidations* quicker to get up and running.

We again remind you to read the additional IXS, IXA, and iSCSI documentation referenced in this chapter to decide which configuration option is best for your environment. Definitely review the Performance Capabilities Reference manual, dated August 2006 or later. The chapter devoted to IXS, IXA, and iSCSI has significantly enhanced performance guideline information.

At the time this book was published iSCSI is supported on Windows Server 2003 with Service Pack 1 (SP1) or Windows Server 2003 Release 2. Over time the range of supported Windows operating systems and releases will be expanded.

Support for specific releases and distributions of 64-bit x86 versions of Linux has been tested and verified on selected System x and BladeCenter models attached via an iSCSI attachment. This includes SLES9 (SuSE LINUX Enterprise Server 9) running on:

- ▶ xSeries and System x servers: x236 (8841), x336 (8837), x346 (8840), x366 (8863), x460 (8872), x3850 (8863), x3950 (8872)
- ▶ BladeCenter chassis model s8867, 8850
- ▶ Blades HS20, LS20, HS21, LS21, LS41

To get the latest information about which Windows and Linux software release and which System x and BladeCenter model feature numbers are supported, go to the Web site we referenced earlier for System i support and specific Linux and Windows Web pages at:

- ▶ <http://www.ibm.com/systems/i/bladecenter/>
- ▶ <http://www.ibm.com/systems/i/bladecenter/linux/>
- ▶ <http://www.ibm.com/systems/i/bladecenter/windows/>

5.2.1 iSCSI introduction

This is a short summary of iSCSI and its protocol as implemented under V5R4 i5/OS.

iSCSI is Internet Small Computer System Interface (SCSI), an Internet Protocol (IP)-based storage networking standard for linking data storage facilities, developed by the Internet Engineering Task Force (IETF). By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. In the IT industry the iSCSI protocol is expected to help extend the capabilities of the storage area network (SAN) technologies. Because of the commonness of IP networks, iSCSI can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet, and can enable location-independent data storage and retrieval.

iSCSI is an industry standard protocol for exchanging SCSI commands within IP packets between a *storage initiator* and a *storage target*. By using IP packets the exchange of commands can be over intranets, LANs and WANs, and supported switches.

iSCSI takes advantage of IP packet header support of larger frame sizes (up to 9 K *jumbo frames*) and faster re-assembly of packets. This is required to make *storage I/O* as fast as possible over an IP-based protocol. Thus, the rules for packet processing are different between iSCSI and standard IP connections, which implies considerations on which *network devices/switches* are used in between the System i HBA and the System x HBA.

Important: Most, but not all, network devices support jumbo frames. System i experience with iSCSI has shown that jumbo frames may not be the right choice in a specific network configuration. iSCSI supports whatever frame size works best in your environment.

iSCSI support requires a hardware feature, an iSCSI Host Bus Adapter (HBA), on both the System i5 model and the System x-based integrated server. A single iSCSI HBA target on the i5/OS system can be addressed by 1–8 System x model iSCSI HBA initiators. Depending on I/O requirements, multiple System i5 HBAs may be required. Each System x iSCSI HBA initiator can address 1–4 System i5 HBA targets.

iSCSI hardware supported on System i models uses industry standard adapters manufactured, at this time, by the company QLogic. System i iSCSI support, including i5/OS-based virtualization capabilities, enables a single point of control to deploy and manage this technology.

Figure 5-1 depicts a simplified view of the connection between an IBM BladeCenter HBA and an IBM System i5 model HBA. Each gold colored *halo* represents the HBA. A more detailed depiction of the iSCSI hardware and i5/OS configuration objects appears later in this chapter.




Figure 5-1 iSCSI connections to System i5

5.2.2 iSCSI installation requirements

The following are the hardware and software installation requirements for the System i5, System x, or BladeCenter systems and Ethernet switch.

Note: In this chapter we use the following terms that are used in the user installation documentation and interfaces:

- ▶ iSCSI target: i5/OS partition receiving the iSCSI requests for data
- ▶ iSCSI initiator: IBM System x-based product making the iSCSI requests for data

For the latest information about xSeries, System x, BladeCenter models, and network switches supported, go to:

<http://www.ibm.com/systems/i/bladecenter>

iSeries System i5 (iSCSI target)

For the iSeries System i5:

- ▶ System i 520, 550, 570, or 595
- ▶ iSCSI Host Bus Adapter (HBA) for System i5 models (up to 1 Gbps)
 - Feature 5783 (CCIN #573B) - Copper
 - Feature 5784 (CCIN #573C) - Fiber

Either feature adapter runs without an IOP. The adapter card is plugged into a 32-bit or 64-bit 3.3v slot (64-bit recommended) attached to a POWER5-based or later system.

A single HBA can attach up to eight System x or BladeCenter servers, depending on customer network requirements and environment and performance considerations. For example, if significant virtual disk activity is anticipated, a lower number of System x or blade servers can reasonably be supported. See the May 2006 or later Performance Capabilities Reference manual for iSCSI performance considerations. This can be found at:

<http://www.ibm.com/eserver/series/perfmgmt>

Select **Resource Library**.

- ▶ i5/OS V5R4 or later is required, with:
 - Integrated Server Support (5722-SS1 opt. 29)
 - Digital Certificate Manager (5722-SS1 opt. 34)
 - IBM HTTP Server for iSeries (5722-DG1)
 - TCP/IP Connectivity Utilities for i5/OS V5R4 (5722-TC1)
 - IBM iSeries Access for Windows (5722-XE1)
 - IBM Director 5.10 (Option of Virtualization Engine (5733-DR1)) Server

All of these i5/OS-based products are no charge and can be included in an order for i5/OS V5R4 install media.

IBM Director 5.1 and i5/OS iSCSI support

The IBM Director 5.1 server component running on i5/OS is used for the *remote server discovery* and off/on control of the iSCSI initiator system. Its functions are invoked by the i5/OS Integrated Server Support to:

- ▶ Discover the server on the external network using the i5/OS iSCSI support remote server object and service processor object. These objects are discussed later in this chapter.
- ▶ The remote server's service processor (SP) should be connected to the intranet to take full advantage of the SP's capabilities. This is especially important since IBM recommends that the HBA connections be made over the iSCSI network, where this is isolated (physically secured) from the external network. Virtual LAN (VLAN) can be configured and used if the same network switch is involved.

The key here is that there are two networks involved: one is the internal intranet iSCSI network and one is the external world Internet network. That is, the SP is located in the corporate intranet to be accessed using the i5/OS LAN adapter.

- ▶ Perform power control functions. This includes turning on (starting) the integrated remote server or performing an operating system shutdown for the corresponding network server description object *vary configuration* functions. These are invoked through Integrated Server Support i5/OS commands or iSeries Navigator interfaces.
- ▶ Retrieval of power on/off status.
- ▶ Configuration of the remote server where possible. Some remote server functions can be configured from IBM Director through the remote server's service processor. Note that the IBM Director server can also perform this same configuration function outside of the iSCSI environment.

You have minimum IBM Director for i5/OS Server setup steps to perform that are dependent upon your environment. You also must specify IP configuration network addresses of your i5/OS connections and your integrated servers to enable use of the IBM Director server with this iSCSI support. Once you have it setup and the IBM Director server is started under i5/OS, you do not directly use the IBM Director interfaces. As stated earlier, you use the Integrated Server support command or iSeries Navigator interfaces.

It is worth noting that with IBM Director Server under i5/OS installed and up and running you can easily start exploring IBM Director's primary functions (available in any supported network), which include managing the software and hardware, and resource utilizations for all the Windows, Linux, and AIX 5L workstations in your IBM System i5-based network.

Only the IBM Director for i5/OS Server at Release 5.10 or later is required to set up and manage iSCSI connections. The i5/OS iSCSI support recognizes that IBM Director server has been installed and uses its functions with no requirement for an IBM Director Console or IBM Director Agent installed anywhere in this network.

Having IBM Director Server installed within i5/OS and also one or more workstations with IBM Director Console and Agent components installed also enables you to perform a wide array of management and monitoring of i5/OS functions. In Chapter 17, “i5/OS and IBM Director” on page 455, we provide overview information about IBM Director 5.10 capabilities. These independent of SCSI capabilities run on either V5R3 or V5R4 i5/OS.

xSeries and System x (iSCSI initiator) installation requirements

The requirements are:

- ▶ xSeries system supported:
 - x236, x336, x346, x366, x460, 3850, 3950
 - RSAll Slimline Adapter 2 required on some models
 - Open bay (diskless server)
- ▶ Each xSeries requires an iSCSI HBAs
 - 32/64 bit 3.3v PCI-X slot
 - P/N 30R5201 - Copper
 - P/N 30R5501 - Fiber
- ▶ Windows Server 2003 with SP1 or Release 2
- ▶ SUSE Linux Enterprise Server 9
- ▶ Gigabit Ethernet switch and cables
 - Ethernet cables (CAT6 or CAT 5e - 100m max distance)
 - Fiber cables (500m - 50/125, 220m - 62.5/125 max distance)

BladeCenter (iSCSI initiator) installation requirements

The requirements are:

- ▶ Requires iSCSI I/O expansion card support (HBA)
- ▶ Blade servers
 - Selected disk-less Intel processor-based blades
 - With iSCSI I/O expansion card (part number 32R1923)
- ▶ 1 Gigabit Ethernet switches or Pass-thru Modules for iSCSI connection
Bays 3 and 4: Bay 4 is optional. It can be used only when more than one target is available.
- ▶ 1 Gigabit Ethernet switch or Pass-thru Module for external network connection
Bays 1 and 2
- ▶ Windows Server 2003 with SP1 or Release 2
- ▶ SUSE Linux Enterprise Server 9
- ▶ Ethernet cables

BladeCenter components for System i5 integration

The components are:

- ▶ BladeCenter chassis (8677xxx, 8852xxx)
- ▶ Blades (1-14) (HS20, LS20, HS21, LS21, LS41)
- ▶ Each with iSCSI HBA (initiator) (part number 32R1923)
- ▶ Power Supply Modules
 - Two are standard (Blades 1–6)

- Two are offered as option (Blades 7–14)
- ▶ BladeCenter Management Modules
 - Full remote video redirection
 - Out-of-band/lights out systems management
 - Concurrent Serial connectivity
 - One standard, second offered as option for redundancy
- ▶ Switches or Pass through Modules for Bays 1 and 2
Exposes Blade Ethernet ports to clients and servers
- ▶ Switches for Bays 3 and 4
I/O expansion card connection for Blades to System i5 iSCSI HBA (target)
- ▶ Flatpanel Monitor
- ▶ Rack and Power Distribution Units (PDUs)
 - For Local Console Manager and Monitor
 - For BladeCenter Chassis

Ethernet switch minimum installation requirements

The requirements are:

- ▶ 1 gigabit layer 2 switch, preferably dedicated to the iSCSI network and isolated from other networks.
- ▶ Copper or fiber optic ports consistent with your iSCSI HBAs.
- ▶ Any copper cables and premises wiring involved should be category 5e or category 6.
- ▶ Fiber optic connections may be preferred in environments with high electrical noise or longer cable length requirements. Fiber optic connections may cost considerably more than copper.
- ▶ Enough ports to accommodate all iSCSI HBAs and possibly the service processor connections.
- ▶ Switch ports should be access ports rather than trunk ports. Access ports are for device connections. Trunk ports are for connections to other switches. You may need to configure access ports using switch features such as Spanning Tree Portfast or STP Fast Start.

5.2.3 iSCSI-attached xSeries or BladeCenter installation

When planning to set up a System x-based iSCSI network with i5/OS, you must read the README FIRST document on this topic, to guide you through the process. Print or get this document at:

<http://www.ibm.com/systems/i/bladecenter/iscsi/readme/>

You also need to use the *System i integration with BladeCenter and System x iSCSI Network Planning Guide* to correctly plan and configure your iSCSI network. This document can be downloaded from the System i integration with BladeCenter and System x Web site listed in several places throughout this book:

<http://www.ibm.com/systems/i/bladecenter/>

Important areas you must carefully plan for and then perform related functions include:

- ▶ The naming conventions you will use for your server-related i5/OS objects. Meaningful names can help you manage the i5/OS objects.

- ▶ Your IP host naming and IP addressing schemes used within the iSCSI network. In a simple to medium iSCSI network, the Planning Guide provides guidance in using naming conventions and IP network parameters. The planned IBM Redbook, SG24-7230, is to provide more complex network setup information.
- ▶ Downloading and installing the necessary System x-based firmware updates. The readme document tells you when to do this and refers you to documentation that tells you how to do this.
- ▶ Downloading and installing additional device drivers you determine to be necessary. For example, drivers for:
 - Embedded or on-board Ethernet controllers
 - IBM options for external LAN adapters
 - Any other IBM options for external devices such as FAX

Note that when you download the firmware updates, you can usually use the same Web site to download the latest device drivers for your system as well. In particular, you should download drivers for:

- IBM Options for external LAN adapters
- Any other IBM Options for external devices you plan on using, such as FAX
- Tape devices when not using i5/OS tape devices

The actual device drivers that you should download vary based on the hardware that is installed in your system, and some of these may not apply to you. Use the documentation that came with your system to determine what hardware and optional adapters are installed in your system and what drivers are appropriate for them.

Using the numbers shown in Figure 5-2 on page 129, we provide a general overview of the hardware configuration objects when configuring iSCSI to connect i5/OS to one or more supported xSeries, System x, or BladeCenter servers. Match the numbers listed below with the corresponding numbers in the figure. The associated software objects for a configuration are shown later in this chapter.

1. You need V5R4 i5/OS and IBM Director 5.1 in a partition (or entire system as a single partition, as shown in the figure) on a System i5 model 520, 550, 570, or 595.
2. The i5/OS console (twinax 5250, HMC 5250 console, Operations Console (Direct attach or LAN), or a new Thin Console) from which you connect to the i5/OS partition using iSeries Navigator or the character-based 5250 i5/OS command interface. This console is shown to make clear the distinction between it and the Windows console.
3. Depending on the type of the physical network you want to use, copper or fiber iSCSI HBAs are available. This iSCSI adapter (HBA) serves as the target device and connects to an Ethernet network using a standard Ethernet cable.
4. An integrated System x-based server does not have its own hard disk drive. i5/OS emulates hard disk space for it to use from the System i5 hard disk drive. These drives and other i5/OS storage devices are accessed through the iSCSI HBA as virtual devices.
5. The iSCSI HBA network cables are connected to a standard 1 Gigabit Ethernet switch.
6. An additional iSCSI HBA is required in the System x-based server. This adapter provides the connection to the iSCSI HBA under i5/OS. This adapter can be viewed from the System x-based server as the storage adapter, where the virtual disks are found across the network.
7. The i5/OS partition must have a network card (LAN adapter) with one or more LAN ports (lines). A LAN line connection to the external network System x-based service processor (SP) is required. This is shown to the switch box at **5** and service processor at **8**. IBM Director uses this connection to discover and manage the remote System x-based server.

Note that depending on the BladeCenter model used, a component called the Management Module replaces the Service Processor. In the context of this chapter use of the term *service processor* includes a Management Module.

- Remember when using an i5/OS partition that both ports of a 2-port LAN adapter are assigned to the partition. You may use one or both ports (lines) for the IBM Director SCSI-based functions. In this figure we depict only one port. On a 2-port LAN adapter the other port can be used by i5/OS to connect to any other network. On a System i5 a 2-port LAN adapter could be the embedded adapter within the processor enclosure (520, 550, 570), or a 5706 or 5707 2-port LAN adapter in an I/O tower or I/O drawer.
- The System x-based server's service processor allows the i5/OS partition (IBM Director server) to discover and manage the remote system. The service processor may be a Remote Supervisor Adapter (RSA II), a Baseboard Management Controller (BMC), or a Management Module of an IBM BladeCenter.

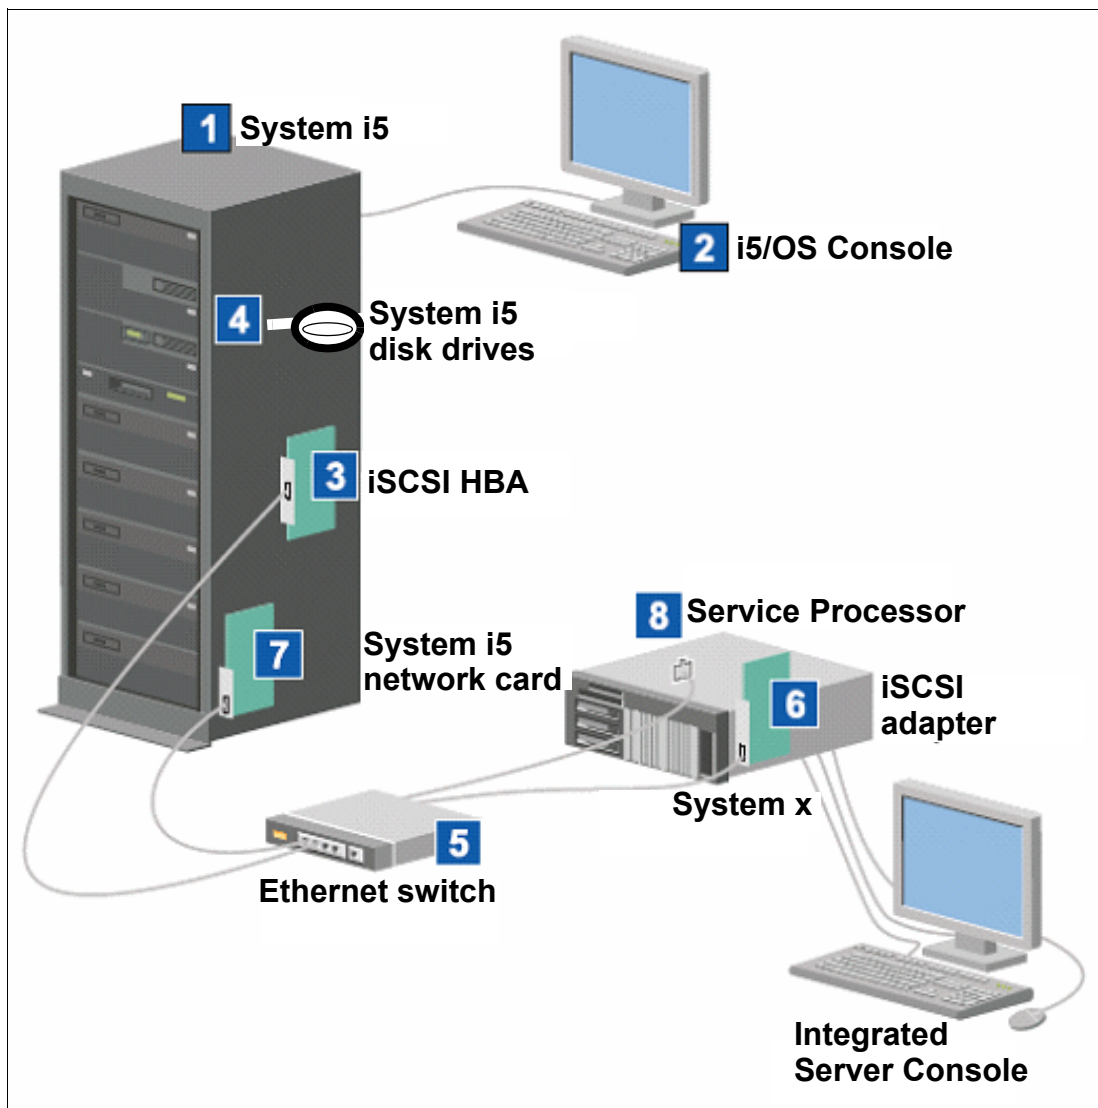


Figure 5-2 iSCSI installation diagram

Completing the hardware view of the iSCSI configuration, we now give you a brief introduction to the new V5R4 iSCSI-based commands, command parameter values, and new objects.

Before discussing the new objects, we first remind you, for iSCSI, that you still must have IXA and IXS support-based objects for:

- ▶ A Network Server Description (type *NWSD) representing the server. For iSCSI support there are additional parameters that include identifying the following:
 - Network Server Host Adapter (type *NWSH) device
 - New Remote System object
Note that the new Service Processor object is not directly referenced by the NWSD. It is indirectly referenced through the Remote System object.
 - New Connection Security object
 - Any multiple path disk I/O options
- ▶ Two or more Network Server Storage Space objects (type *NWSSTG) representing virtual disks.
- ▶ Virtual Ethernet objects representing a point-to-point connection for internal communications and the optional virtual LAN connection for Windows communications with the i5/OS partition.

The new iSCSI support objects are:

- ▶ Network Server Host Adapter (type *NWSH) device: This represents a physical iSCSI Host Bus Adapter in the i5/OS partition.
- ▶ Service Processor object (new NWSCFG object type, with subtype of *SRVPRC): This object is used in Ethernet communication between the i5/OS partition and service processor on the xSeries or System x server or a BladeCenter Management Module (which manages the BladeCenter configured processors).
- ▶ Remote System object (new NWSCFG object type, with subtype of *RMTSYS): This object is used when communicating over the iSCSI network associated with the HBA adapter and its associated Network Server Host Adapter (type *NWSH) device.
- ▶ Connection Security object (new NWSCFG object type, with subtype of *CNNSEC): This object has to exist. However, its functions to increase iSCSI network connection security are planned as a future enhancement.

The Network Server Description object is used to tie all of these components together. You still need at least a virtual C: disk drive where the Microsoft Windows operating system code is stored and a virtual D: disk drive where the installation image is stored.

Refer to Figure 5-3 and Figure 5-4 when reading the topics that follow them. The topics provide additional information about these new command parameters, values, and objects.

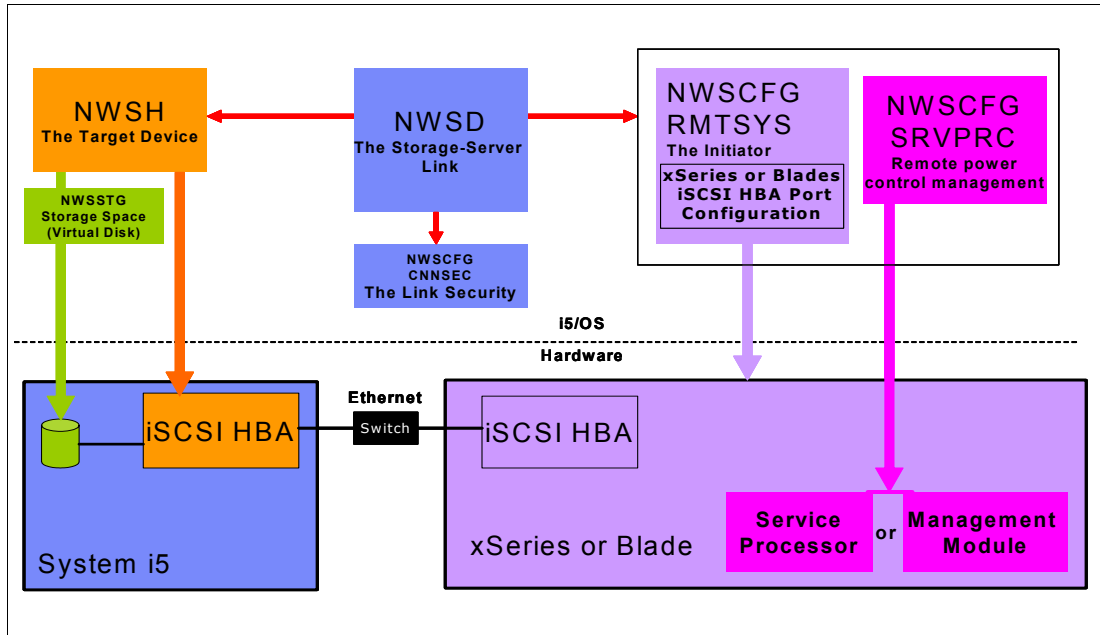


Figure 5-3 V5R4 iSCSI support configuration objects: view 1 of 2

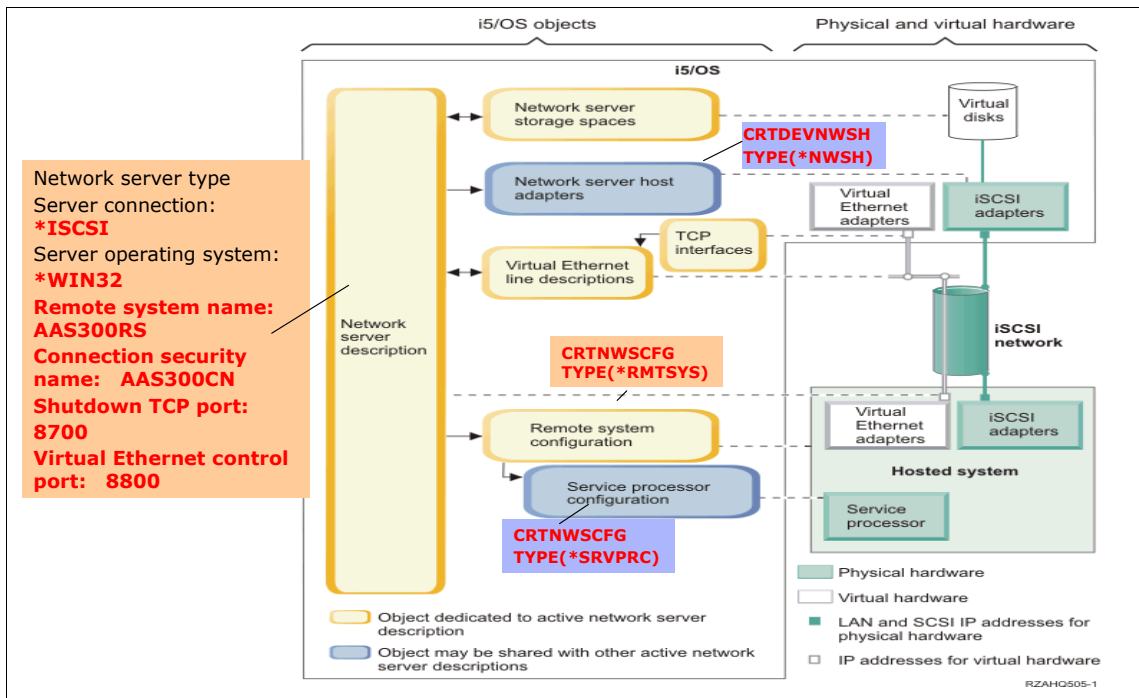


Figure 5-4 V5R4 iSCSI support configuration objects: view 2 of 2

The following topics provide additional information about the new iSCSI support objects.

Network server host adapters (CRTDEVNWSH command)

The network server host adapter (*NWSH) device description object shown in these figures represents the iSCSI host bus adapter (HBA) that is used by the System i5 side of the iSCSI connection:

- ▶ It identifies the iSeries hardware resource name (for example, LIN33) for the iSCSI HBA.
- ▶ It defines how communications errors are logged and communications recovery processing options.
- ▶ It defines the Internet addresses, ports, and so on, for the SCSI and LAN interfaces on the iSCSI HBA.

Remote system configuration (CRTNWSCFG TYPE(*RMTSYS) command)

A remote system network server configuration object, type *RMTSYS, which represents the iSCSI attached System x, xSeries, or IBM BladeCenter server:

- ▶ It identifies the server hardware by serial number and type and model.
- ▶ It contains configuration information for the iSCSI host bus adapters (HBAs) that are used by the xSeries or IBM BladeCenter server.
- ▶ It contains values required to boot the server (such as specifying which iSCSI adapter to boot from).
- ▶ It contains a reference to the service processor NWSCFG object (see below) that is used to control the System x-based server.
- ▶ The remote system configuration can optionally contain values used to secure the server boot process.
- ▶ The remote system configuration object for an integrated server is referenced via a parameter in the NWSD.

Service processor configuration (CRTNWSCFG TYPE(*SRVPRC)command)

A service processor network server configuration object represents the System x, xSeries service processor, or the IBM BladeCenter management module:

- ▶ It identifies the service processor or management module hardware by serial number and type and model.
- ▶ It defines how to find the service processor or management module on the Ethernet network using an Internet address or host name.
- ▶ The service processor object can optionally contain values used to secure the i5/OS to service processor communications.

Connection security network configuration (CRTNWSCFG TYPE(*CNNSEC))

The connection security network server configuration (NWSCFG subtype CNNSEC) object referenced in the figure defines the connection security name parameter. It must be created, but currently does not provide additional security attributes for the iSCSI connection.

Additional security attributes are planned for a future release.

5.2.4 iSeries Navigator support for iSCSI connections

i5/OS V5R4 added support in iSeries Navigator to be able to configure an iSCSI connection. In V5R3 iSeries Navigator has a Windows Administration folder to support Windows

integrated servers. In V5R4 this folder has been renamed to Integrated Server Administration to indicate that support is not only for Windows but also for Linux integrated servers.

Under the Integrated Server Administration folder, there is a new iSCSI Connections folder that supports iSCSI connection management and configuration, as shown in Figure 5-5.

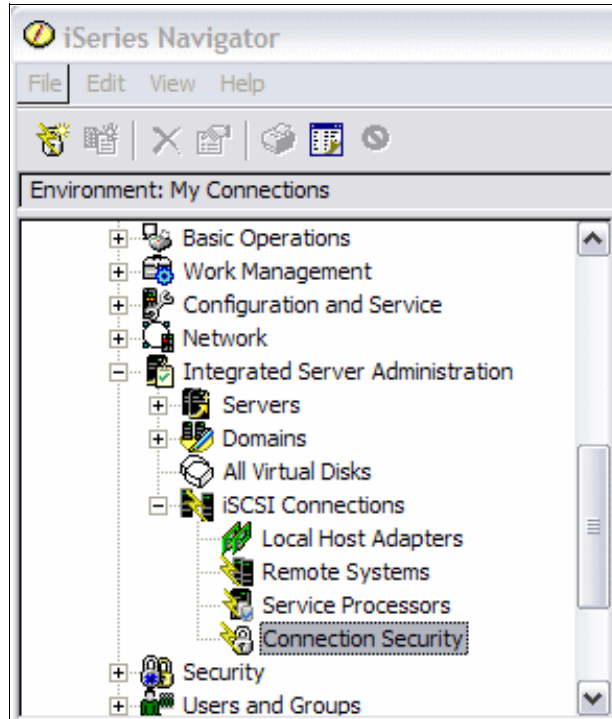


Figure 5-5 iSCSI Connection subfolder in iSeries Navigator V5R4



i5/OS-based security

Security is integrated into the base i5/OS.

This chapter summarizes V5R4 security enhancements, including topics on system integrity, intrusion detection, cryptographic support, antivirus program call updates, Enterprise Identity Mapping (EIM), and auditing feature changes.

We also summarize the capabilities of the Licensed Program Offering (LPO) Portable Utilities i5/OS. This product is based on open source capabilities focused on offering secure application alternatives to standard Telnet and FTP applications. This product became available during V5R3 and provides another example of i5/OS openness to industry-validated open interfaces.

6.1 V5R4 i5/OS security overview

Starting with the first release of OS/400 through the latest release of i5/OS, major improvements have been made to ensure that data and programs on the system are protected.

Today many virus programs have been spread across the Internet, with the sole intention of destroying or stealing information from your system. These programs may have the power to alter your data, render your system unusable (denial of service), or gather confidential data from your system and pass this information over the Web.

i5/OS V5R4 includes a wide range of security and compliance enhancements. For example, new hardware storage protection helps prevent against rogue or altered programs from directly accessing system objects, such as database records. Objects, including database and file system containers, are accessible only to IBM Licensed Internal Code, helping protect against rogue system state programs. Certain vendors proved i5/OS applications that have used *system state* to provide useful functions. Those vendors have been contacted and different interfaces continue to enable those functions. No changes are required for any other valid application programs.

New intrusion detection capabilities enable an administrator simply to configure intrusion event monitoring with entries created in the audit journal for intrusion events, such as scanning for open TCP/IP ports. A policy file allows the administrator to configure intrusion event monitoring and create audit journal entries created for intrusion events such as malformed packets, SYN floods, restricted IP options and protocols, traffic regulation events, and port scans.

Additional auditing and compliance enhancements include the ability to audit special authority violations. So, for example, an audit journal entry would be recorded if a user is authorized to CRTUSRPRF create user profile command without the corresponding *SECADM security administrator authority. Also, now users will not be able to display their audit attributes. This function will be limited to system administrators. Note that these capabilities add to the extensive enhancements to the audit journal in V5R3, with much more granular options to save only the security data required.

This chapter discusses the security enhancements made available to the i5/OS V5R4 since its announcement in January of 2006. This chapter includes topics on system integrity, intrusion detection, cryptographic support, antivirus updates, Enterprise Identity Mapping (EIM), and some auditing feature changes.

Other topics on security are detailed in the *IBM Systems - iSeries Security Reference V5R4*, SC41-5302-09. You can download this manual from:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r4/topic/rzamv/sc415302.pdf>

6.2 System integrity

i5/OS V5R4 is the first release to have Enhanced Hardware Storage Protection (Enhanced HSP). This new feature is a stronger HSP *force field* applied to the System i objects so that only the Licensed internal code (LIC) has access to them. This means any program, malicious in nature or otherwise, attempting an unsupported access to these objects will definitely fail.

Enhanced Hardware Storage Protection

Objects that are created in V5R4 already have this Enhanced HSP attribute. Each time a RISC instruction is executed, the hardware via the LIC determines whether the page frame is hardware storage protected. This process provides the control and type of access to that object. As such, hardware storage protection allows a user domain object to be created so it can be read from a user state program but cannot be modified by the user state program.

A simple way to look at this is to assume that a process running under i5/OS is granted addressability (with an identified type of access) to a data object when it can obtain a valid system pointer. This validation is done by checking the hardware tag bits, a hardware controlled process. Even when an attempt is made to modify a hardware tagged pointer, the tag bit is cleared and the data object becomes unaddressable. The only allowed modification can be performed by an authorized machine interface (MI) instruction. This methodology prevents any type of change made to an object on the system by any type of program or other unsupported access that cannot make use of an authorized MI instruction.

To help with restricting the existence of user domain objects on the system, the QALWUSRDMN system value can be used. The QALWUSRDMN system value specifies a list of libraries that allow user domain objects. These libraries on the system may then contain the user domain user objects *USRSPC (user space), *USRIDX (user index), and *USRQ (user queue).

The internal representation of SQL statements in *PGM, *SRVPGM, *MODULE, and *SQLPKG objects changed in V5R4. *PGM, *SRVPGM, and *SQLPKG objects that were created prior to V5R4 will be converted automatically to the new format on first use. You can also initiate object conversions using the Start Object Conversion (STROBJCVN) command. *MODULE objects that were created prior to V5R4 cannot be permanently converted, and a conversion will be performed each time the *MODULE is bound into a program or service program. The amount of time to convert a single object is small, less than 160 milliseconds of CPU time on a model 890 system for a program approaching the system limit on number and size of SQL statements. On a model 840 system, it is less than 350 milliseconds. The conversion time is proportional to the number of SQL statements in the object and to the cumulative size of the SQL statements. If large numbers of objects need to be converted, the conversion time could be noticeable.

For objects created on earlier releases, they are realigned with the new HSP security information upon first touch to ensure maximum performance on V5R4. This conversion is done in background SLIC tasks upon first use, but will have some performance impact (impact varies, dependent on particular objects) on the system.

Key objects being converted include data queues, user queues, and database objects.

To enhance the runtime performance of journaled data queue (*DTAQ object) synchronization and to increase the debug capabilities of all queues (*DTAQ and *USRQ objects), the first time a queue is touched after V5R4 is installed, an internal conversion takes place. This conversion might take a noticeable amount of time for large journaled data queues, but should be relatively quick for other data queues. To avoid delays in a production environment, you might want to touch all queues immediately following the installation. One way that you can do this is to run the Display Library (DSPLIB) command, and then select either Option 5 or 8 (to display attributes) for each *DTAQ or *USRQ object. You can also use commands, such as Display Library (DSPLIB) or Display Object Description (DSPOBJD), within a submitted batch job.

Database files that were created on a release earlier than V5R4 are also converted to eliminate a performance degradation that might occur due to Enhanced Hardware Storage Protection support. Because database files (*FILE objects) consist of several machine

interface (MI) objects, the time to convert them could be significant when an application that accesses many files is started for the first time after V5R4 has been installed.

Also, because the MI objects will have their header size increased to 4 K during the conversion, each MI object associated with a database file will grow in size by 4 K (4096 bytes). To spread the MI object conversion process out over a period of time, the database file conversion support is designed to be done when the database MI objects are first accessed or used after V5R4 is installed. Similar for data queues, the user can be proactive on this by setting up a time to do this first access via display commands before running a performance-critical application for the first time on V5R4.

If you would like to evaluate the impact that the database file conversions might have on one or more of your applications and estimate how much time you will need during your upgrade window to take the necessary steps to eliminate the startup performance degradation for your more critical applications, then run the QDBFIMIC program that is included in PTF SI18716 (if you are upgrading from V5R2) or PTF SI18615 (if you are upgrading from V5R3). See the PTF cover letter for detailed usage instructions and what happens if a database file object is locked too long while QDBFIMIC is running.

To run QDBFIMIC to process all file objects in library MyLibrary, use the command:

```
CALL QSYS/QDBFIMIC (MyLibrary)
```

To process a file MyFile in library MyLibrary use the command:

```
CALL QSYS/QDBFIMIC (MyLibrary MyFile)
```

Running the commands above will result in a message at the bottom of the screen, which, when displayed, will show something like in Figure 6-1. The information in the message CPF9898 is useful in estimating the amount of time to convert the database objects of your system after V5R4 is installed. To estimate this conversion time do the following:

- ▶ Take the number of objects needing conversion as listed in message CPF9898.
- ▶ Take the CPW of your V5R4 system.
- ▶ Extrapolate to an estimated conversion time by using the above two values and the V5R4 *memo to users* examples showing a large number of objects being converted and conversion runtimes on three CPU models (830, 825, and 570). You can download the PDF copy of the V5R4 Memo to Users from:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzaq9/rzaq9.pdf>

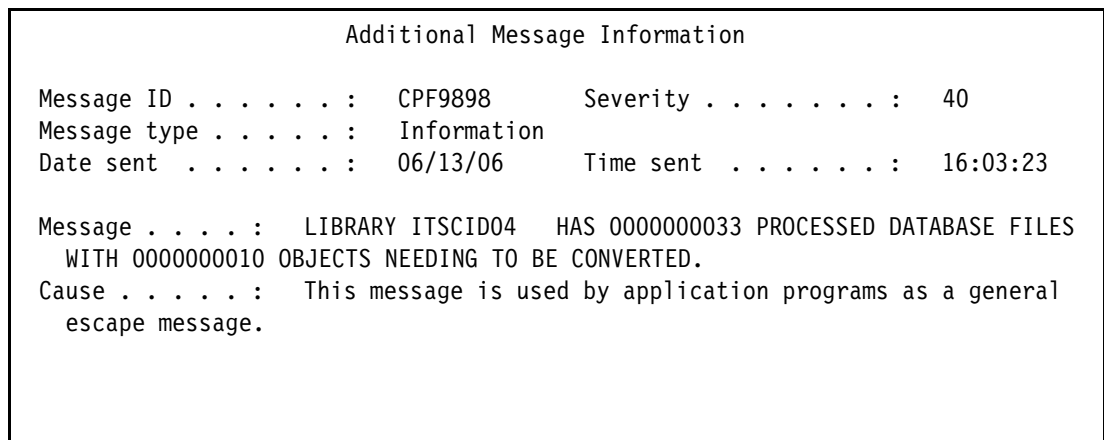


Figure 6-1 CPF9898 message details

The impact in doing the conversion is dependent on the number of objects needing conversion, the CPW, and the amount of storage pool allocated to that conversion job. Most customers upgrading to V5R4 have reported very little impact due to the conversion process. However, if you have a large number of files (and thus potentially a large number of objects to be converted), we strongly recommend running this program and examining its output.

6.3 Intrusion detection

One of the security enhancements of the i5/OS V5R4 is that the V5R4 TCP/IP stack comes with intrusion detection. Intrusion is being used here as gaining unauthorized access to your system with the sole purpose of stealing information, denial of service (harming the business), or to gain unauthorized use of data for whatever purpose.

If managing the server requires you to gather information about unauthorized access attempts and attacks coming in over the TCP/IP network, then this is called intrusion detection. A specialized administrator in security could analyze the auditing records that intrusion detection provides to secure the network of the server.

Intrusion includes activities such as information theft and denial of service attacks. The objective of an intrusion may be to acquire information that a person is not authorized to have (information theft). The objective may be to cause a business harm by rendering a network, system, or application unusable (denial of service), or it may be to gain unauthorized use of a system as a means for further intrusions elsewhere. Most intrusions follow a pattern of information gathering, attempted access, and then destructive attacks. Some attacks can be detected and neutralized by the target system, whereas some may not. Most of the attacks also make use of spoofed packets, which are not easily traceable to their true origin. Most of the time the identity of the attacker is hidden or an attempt is made to hide the attacker's identity. For these reasons, a vital part of intrusion detection is gathering information, and detecting access attempts and attack behaviors.

In V5R4 the user has been given the capability to create an intrusion detection policy that audits suspicious intrusion events that come in through the TCP/IP network. The i5/OS intrusion detection functionality is designed to look for the following:

- ▶ Denial of service attacks and unauthorized gathering of information
- ▶ Port scans and network sniffing
- ▶ Malformed packets
- ▶ Internet protocol (IP) fragments
- ▶ Restricted IP options and protocols
- ▶ Internet Control Message Protocol (ICMP) redirect messages
- ▶ Perpetual echo attacks on User Datagram Protocol (UDP) port 7 (the echo)

Intrusion detection is implemented by using a policy file. This policy file is used by the quality of service server to instruct the TCP/IP stack as to what conditions require what actions.

The process flow shown in Figure 6-2 summarizes how intrusion detection is implemented within the system.

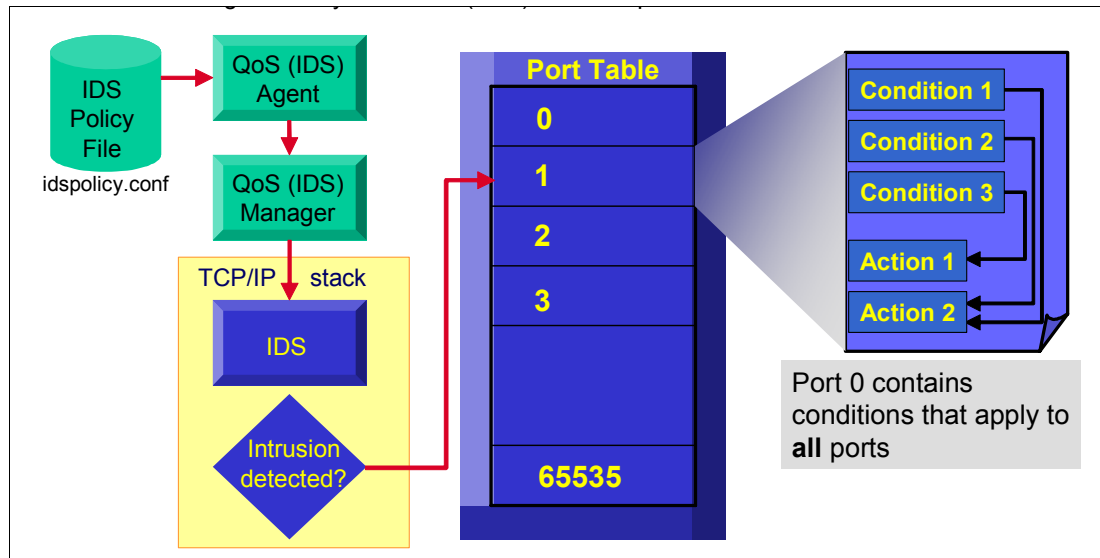


Figure 6-2 Implementation flow of intrusion detection

To implement intrusion detection follow these steps:

1. Check for proper authority on file idspolicy.conf in the following directory:
 /QIBM/ProdData/OS400/QOS
 The file is delivered with the operating system.
2. Edit the idspolicy.conf file to detect specific types of intrusions.
3. Start the QoS server. IDS uses the Quality of Service (QoS) service to enable intrusion detection services. QoS in the TCP/IP attributes must be enabled. To enable the Quality of Service function use the command:
 CHGTCPA IPQOSEN(*YES)

The QoS policy agent reads the intrusion detection policy in the idspolicy.conf file.

The QoS policy agent sends a message with machine instructions to the QoS manager.

The QoS manager interprets the machine instructions and sends them to the intrusion detection system inside the TCP/IP stack. The TCP/IP stack manages outbound traffic and inbound traffic in the network, and routes requests to other computers in the network.

The intrusion detection system creates the policies in the port table. The port table entries represent ports 0 through 65,535. When port 0 contains conditions and related actions, they apply to all ports. In the example on Figure 6-2, there are three conditions and two actions defined for port 1. When one of the conditions 1 or 2 are met, action 2 will be performed. When condition 3 is met, action 1 is performed. You can have multiple conditions defined that point to the same action.

When the TCP/IP stack detects an intrusion, it looks for matching conditions in the port table and executes the specific action, for example, creating an IM auditing record or system statistics.

The system creates an IM audit record that describes the type of intrusion event.

Your system administrator analyzes the IM audit record to determine which security actions to take, such as closing off the port from where the intrusion originated and contacting network security specialists to pursue action against the originator.

For more information about how to start and use audit journals and for finding more security information such as setting the appropriate security system values, use the *iSeries Security Reference manual*, SC41-5302.

See also Redpaper *IBM i5/OS Intrusion Detection System*, REDP-4226. This provides additional information and examples of setting up V5R4 intrusion detection, which is sometimes referred to as the Intrusion Detection System capability.

6.4 Cryptographic support

In V5R3, the Kerberos server was included as part of the optional, no cost, 5722-AC3 product. In V5R4, the 5722-AC3 product is no longer available. The Kerberos server is now shipped in the Network Authentication Enablement (5722-NAE) product. If V5R4 is installed over V5R3, and the 5722-AC3 product is currently installed, then the 5722-NAE product is automatically installed to ensure that the Kerberos server that was part of the 5722-AC3 product remains installed.

If you are using Kerberos before moving into i5/OS V5R4 and you need to know if the licensed program 5722-NAE is already installed on your system, you can check this using Option 10 in the GO LICPGM.

If V5R4 is installed over V5R2, and the 5722-AC3 product is currently installed, then the 5722-NAE product is not automatically installed, since the Kerberos server was not part of 5722-AC3 in V5R2.

Note that the licensed program product (LPP) 5722-NAE is part of the 5722-SS1 media set in V5R4.

Cryptographic support is enabled via Option 35 of 5722-SS1, Common Cryptographic Architecture Cryptographic Service Provider (CCA CSP).

The Client Encryption (5722-CE3) product will no longer be shipped with V5R4. The client encryption function, Secure Sockets Layer (SSL), is now included in the base of the V5R4 5722-XE1 iSeries Access for Windows product. The licensed program 5722-CE3 can still be used for supporting clients that use supported iSeries Access for Windows versions older than V5R4.

6.4.1 Cryptographic support for AS/400 (5722-CR1)

IBM plans to drop support of IBM Cryptographic Support for AS/400 (5722-CR1) in the release after V5R4. This license program product provides cryptographic function APIs similar to the functions available with the Common Cryptographic Architecture (CCA) APIs that are used with a cryptographic coprocessor, such as the 4758-023 or 4764. The V5R4 iSeries Information Center provides migration scenarios for moving applications that currently use 5722-CR1 to an environment that exploits the Cryptographic Services APIs. As outlined in the Information Center, in some cases it might be a better choice to move to a solution that utilizes a cryptographic coprocessor. Customers and ISVs who maintain applications that use the 5722-CR1 product should plan for migrating their applications to an alternative environment.

6.5 Enterprise identity mapping

Enterprise Identity Mapping (EIM) for iSeries allows administrators and application developers to managing multiple user registries across their enterprise. Most network enterprises are challenged by multiple user registries, which require each person, user, or entity within the enterprise to have an identity in each registry in each system. The need for multiple user registries quickly grows into a large administrative problem that affects users, administrators, and application developers. Enterprise Identity Mapping (EIM) makes this easier for you to manage multiple user registries and user identities in your enterprise.

EIM allows you to create a system of identity mappings, called associations, between the various user identities in various user registries for a person in your enterprise. EIM also provides a set of APIs that can be used across platforms to help develop applications that can use the identity mappings that you create to look up the relationships between user identities.

EIM itself does not provide authentication. You can use EIM in conjunction with network authentication service, the i5/OS implementation of Kerberos, to provide a single sign-on environment.

You can configure and manage EIM through iSeries Navigator, the iSeries graphical user interface, as shown in Figure 6-3 on page 143. The iSeries server uses EIM to enable i5/OS interfaces to authenticate users by means of network authentication service. Applications, as well as i5/OS, can accept Kerberos tickets and use EIM to find the user profile that represents the same person as the Kerberos ticket represents.

In releases prior to V5R4, there is only a users' registry definition per system. However, as these user registries grow across the enterprise, this can become a challenge in administration.

Logically grouping the registry definitions enables the reduction of the amount of work that must be performed to configure EIM mapping. With V5R4 it is possible to manage a group registry definition similar to the way an individual registry definition can be managed.

All members of the group registry definition typically contain at least one common user identity to which a target or source association would be created. By grouping members together it is now possible to create only one association, rather than multiple associations, to the group registry definition and user identity.

This is possible if the following conditions are met:

- ▶ All of the members (individual registry definitions) of the group registry definition must be defined in the EIM domain before they can be added to a group registry definition.
- ▶ A registry definition can be a member of more than one group, but specifying an individual user registry as a member of multiple group registry definitions should be avoided.
- ▶ All of the members (individual registry definitions) of the group registry definition must have the same case sensitivity.

By using group registries, you can avoid identical user definitions as a member of multiple group definitions.

In Figure 6-3 we show a portion of an iSeries Navigator window used to add a group registry.

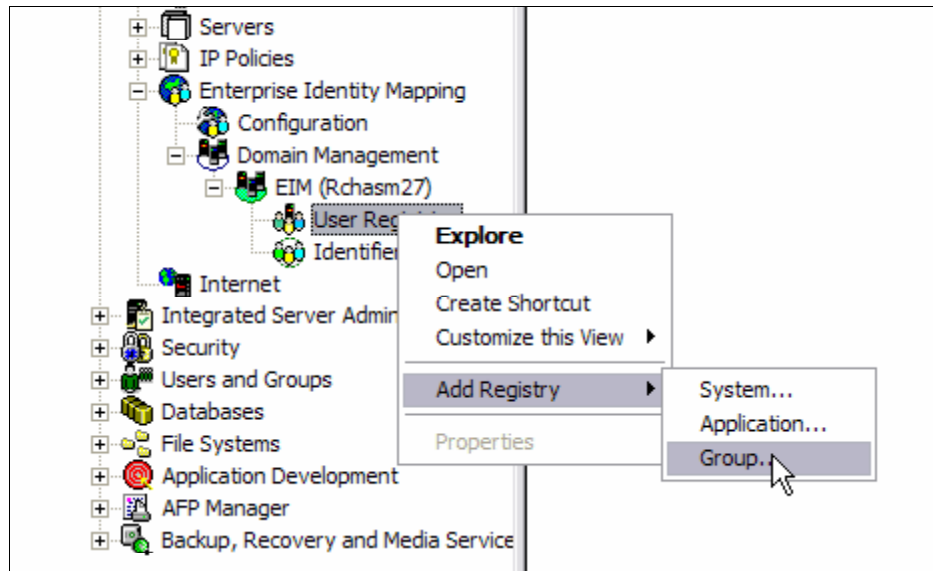


Figure 6-3 Adding group registry using iSeries Navigator

To add a group registry:

1. Expand **Network** and select **Enterprise Identity Mapping**.
2. Click **Domain Management**.
3. Right-click **User Registry** and click **Add Registry**.
4. Fill in the required information under the General tab. Figure 6-4 shows an example.

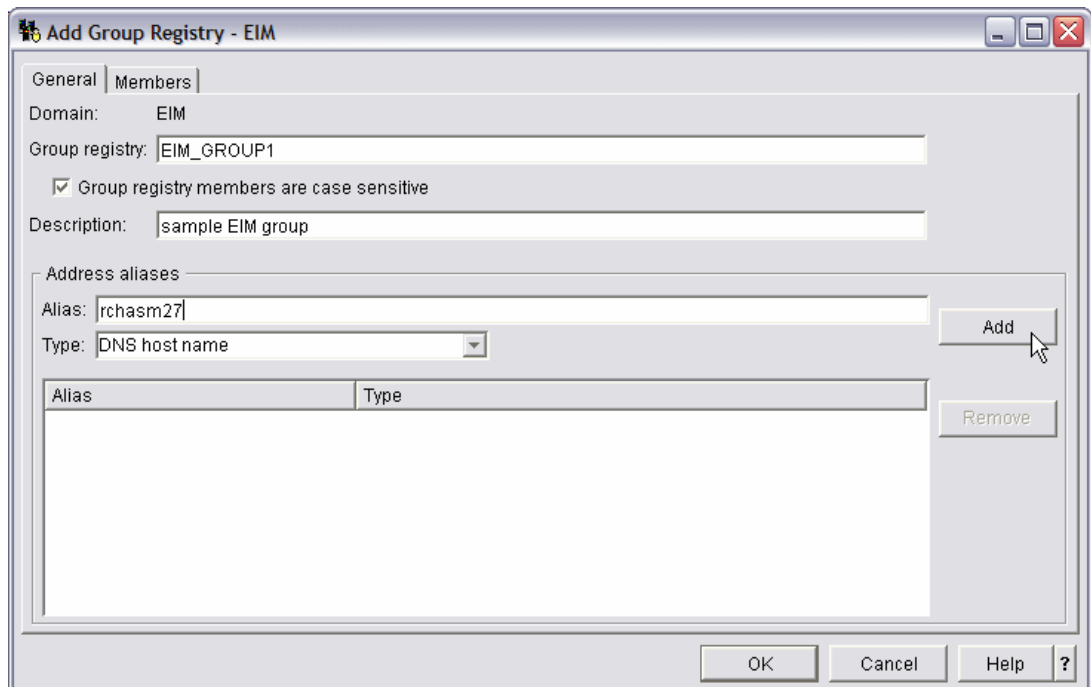


Figure 6-4 Adding a group registry in EIM

5. Ensure that you have an existing user registry to add to your new group registry.

6. Click **Members** tab and provide the necessary information, like who will be the members of this group registry.

For a more detailed discussion on i5/OS V5R4 EIM support, please visit:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/index.jsp?topic=/rzalv/rzalvmst.htm>

6.6 Improved antivirus scanning

Every day around the world there are probably hundreds if not thousands of systems being affected by viruses. These rogue programs can bring a network or a server down and render it unusable, causing businesses to lose huge amounts of revenue.

Although there are no known viruses that run on a System i, viruses can still be transferred to and from other servers via the Integrated File System (IFS). Because of this the i5/OS operating system has been designed to enable third-party antivirus software to scan the Integrated File System.

In i5/OS V5R3 and now in V5R4, IBM has introduced user exit enablement for enhanced virus scanning of the Integrated File System (IFS) objects. It also enabled third-party vendors to develop antivirus scanning software that plugs into i5/OS. Software like StandGuard Antivirus for iSeries from Bytware and Raz-lee's Security's iSecurity Suite, which includes their antivirus software, are examples of antivirus software available for iSeries. There are other vendors that may also provide this support.

It is also possible now with the use of the system value QSCANFCTL that a user exit program can start a scan on an object once that object is opened, closed, saved, or restored.

6.7 Portable Utilities for i5/OS

Though Portable Utilities i5/OS, 5733-SC1, became available initially with i5/OS V5R3, we include an overview of its capabilities here to ensure that you are aware of these capabilities, especially for interaction with non-i5/OS operating systems.

TCP/IP connectivity applications such as Telnet and FTP transmit data and passwords over the network in plain text. This means that the data and passwords can be intercepted and read by other users on the network. You can secure some of these transmissions with various encryption mechanisms, including digital certificates and open source implementations. 5733-SC1 enables open source implementation of secure alternatives to Telnet and FTP under i5/OS.

The Secure SHell (SSH) protocol suite is a software solution that provides secure alternatives for Telnet and FTP. SSH verifies that the authenticity of both the client and the server and all of the data (including user IDs and passwords) is encrypted as it travels on the network. This encryption is done transparently to the end user. With this support you can also include, typically within your FTP data, scripts that can perform functions on the target system, such as installing software.

OpenSSH is the open source implementation of the SSH protocol suite. OpenSSH is widely available for use on many other operating systems including Linux, AIX, and z/OS. The OpenSSH development team maintains Web pages at:

<http://www.openssh.org/>

The 5733-SC1 Licensed Program Offering (LPO) contains the OpenSSH, OpenSSL, and zlib open source packages, which were moved to i5/OS using the i5/OS PASE runtime environment. 5733-SC1 requires i5/OS V5R3 or later and also requires that i5/OS Option 33 (i5/OS PASE - Portable Solutions Application Environment) be installed.

The OpenSSH portion of 5733-SC1 consists of these utilities:

- ▶ ssh: This is a secure Telnet replacement that allows an i5/OS user to connect as a client to a server running the sshd daemon. An ssh client can also be used to connect to the Hardware Management Console (HMC) on the IBM System i5 models.
- ▶ sftp: This is a secure FTP replacement. As with all implementations of sftp on other platforms, sftp can only transfer data in binary format. Note that sftp does not provide the enhanced functions available in the i5/OS FTP utility when transferring files in the QSYS.LIB file system, nor does it provide the CCSID data conversion options available in the i5/OS FTP utility.
- ▶ scp: This is a secure file copy program. This is essentially an alternative to sftp for copying a single file in the Integrated File System (IFS).
- ▶ ssh-keygen: This is a public/private key generation and management tool. SSH allows users to authenticate using these public and private keys as an alternative to using their operating system sign-on password.
- ▶ ssh-agent: This is an authentication agent that can store private keys. ssh-agent allows a user to load their public/private key pass phrase into memory to avoid needing to retype the pass phrase each time an SSH connection is started
- ▶ sshd: This is a daemon that handles incoming ssh connections. The sshd daemon utility allows users to connect to i5/OS via an ssh client. In contrast to connecting to i5/OS via Telnet and being presented with a 5250 screen interface, users that connect via ssh to an i5/OS partition running the sshd daemon will be presented with a character interface and an i5/OS PASE command line.

For specifics on the OpenSSH utilities see the reference manual pages at:

<http://www.openssh.org/manual.html>

The OpenSSL and zlib components of 5733-SC1 have the following characteristics:

- ▶ OpenSSH requires the OpenSSL and zlib open source libraries in order to function. A user can write i5/OS PASE applications that call OpenSSL or zlib APIs.
- ▶ OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL). Refer to the OpenSSL Web pages for more information:
<http://www.openssl.org/>
- ▶ zlib is a set of data compression APIs. For more information refer to the zlib Web page at:
<http://www.zlib.org/>



i5/OS-based backup and recovery

This chapter covers enhancements to backup and recovery subjects, including i5/OS and V5R4 Backup and Recovery Media Services (BRMS).

Journaling enhancements are covered in several other areas, including:

- ▶ 4.4, “On Demand and availability” on page 96.
- ▶ 6.3, “Intrusion detection” on page 139
- ▶ “Optional journaling for SQL tables” on page 303

7.1 i5/OS: saving and restoring data

There is an old saying, *better safe than sorry* — a good backup routine is the basis of all disaster recovery plans. Recognizing that one size does not fit all, we continue to improve the tools that are available to you. We cover these enhancements in some depth, including saving Spool files, parallel saves, and Backup and Media Recovery Services (BRMS).

7.2 Spooled file support

Save/restore of spooled files is now supported in i5/OS. In this section, we summarize the save and restore enhancements for spool files on output queues. Refer to Chapter 15, “i5/OS printing and printer output enhancements” on page 377, for more details on these save and restore functions, which are included in both i5/OS commands and BRMS interfaces.

The new function preserves spooled file data, identity, and attributes:

- ▶ Spooled file name
- ▶ Spooled file number
- ▶ Creation date/time
- ▶ Fully qualified job name
- ▶ Job system name

There are some attributes that may not be preserved:

- ▶ System 36 identifier
- ▶ Output queue name/library (restored to different output queue/library)
- ▶ ASP (output queue on different ASP)
- ▶ New expiration date (restore occurs after expiration date)
- ▶ New save/restore attributes of spooled file

We have added a new SPLFDATA parameter on save/restore commands:

- ▶ SPLFDATA(*NONE) is the default on saves.
- ▶ SPLFDATA(*ALL) is supported.
- ▶ SPLFDATA(*NEW) is the default on restores.
- ▶ Restore spooled files that do not exist.
- ▶ SPLFDATA(*NONE) is supported on restores.

Save and restore is of entire OUTQs, as there is no support to save/restore individual spooled files. Menu options allow you to optionally save and restore *all* spooled files, no matter the queue.

You can display the save/restore media to see a list of the spooled files with the DSPTAP, DSPOPT, and DSPSAVF commands. This is similar to displaying members in a database file.

We have been able to realize a performance improvement, and this method is now faster than existing methods to save/restore spooled files. There are no requirements to copy the spool file to a database file, or for Document Library Objects (DLOs). Files are saved and restored directly, as is any other object.

BRMS uses the new support, and while there are no user interface changes, you do have the option of filtering the files you save in BRMS. We also see the performance improvement in BPRMS.

7.3 Dynamic device allocation

Dynamic device allocation allows you to allocate tape devices as needed for parallel saves. There have been enhancements to the Media Definition Object.

Prior to V5R4, all tape devices for the save had to be allocated at the start. Enhancements have been made to parallel saves by having dynamic device allocation. For example, on a *save while active* you could only allocate one tape drive initially and not require more until you reach the *save while active* checkpoint. Enhancements include allowing tape drives to be used for other jobs during checkpoint processing. Also, you can begin using additional tape drives as they become available during the save.

This allows you to specify the minimum number of resources required at write time. On SAVLIB LIB(*ALLUSR) at each library boundary we check with each potential physical unit to see if it is ready to write to.

BRMS manages device allocation automatically.

Enhancements to Media Definition (MEDDFN) Object include changes to the QsrCreateMediaDefinition API (used internally by BRMS) for device allocation:

- ▶ Default value 0 - all tape devices allocated at start of save
- ▶ 1 - one tape device allocated at start of save and additional tape devices allocated when ready to write data
- ▶ 2 - minimum number of resources specified allocated at start of save and additional tape devices allocated when ready to write data

The restrictions include:

- ▶ All save operations will continue to initially allocate at least one tape.
- ▶ Any save operation that does not use a media definition will allocate its device at the start of save.
- ▶ Devices will not be dynamically deallocated.
- ▶ Attempts to dynamically allocate devices will be limited to these points in time:
 - After save-while-active checkpoint
 - When the library data is ready to be written to an available device, that is, just prior to writing the data

If multiple libraries are being saved, will check for available devices before writing each library.

7.4 IFS parallel save/restore support

There is now a single SAV or RST command when using multiple tape drives. This can reduce the Save/Restore window for large Integrated File System (IFS) objects, such as large network storage spaces and Domino databases. Performance decreases with small objects, so we advise against using this to save a collection of millions of objects each less than 1 MB in size.

You create the media definition with the QsrCreateMediaDefinition/QSRCRTMD API. You specify the Media Definition Object in the Device (DEV) parameter. You can use up to 32 tape devices in parallel.

You can save the entire IFS to multiple tape devices once the media definition X is created.
SAV DEV('/QSYS.LIB/Y.LIB/X.meddfn') OBJ('//*') ('QSYS.LIB' *OMIT) ('/QDLS' *OMIT))

This is supported with the QsrSave and QsrRestore APIs.

The end-user interface of BRMS simplifies this. There is no need to create a Media Definition. You simply specify minimum and maximum parallel resources on the BRMS Control Group.

7.4.1 Automate parent directory creation

There is a new Create Parent Directories (CRTPRNDIR) parameter. This will optionally create parent directories of objects if they do not exist. For example, if object /a/b/c/file1 is being restored, then directories /a, /a/b, and /a/b/c must exist. This parameter only applies to root (/), QOpenSys, and user-defined file systems, and will be ignored for all other file systems.

*NO is the default, meaning that parent directories are not created. Diagnostic message CPD375B is sent and the object is not restored.

If you specify *YES, the parent directories are created with *EXCLUDE public authority. These objects are owned by the user profile specified on the Parent Directory Owner (PRNDIROWN) parameter. If a value is specified for this parameter, *YES must be specified for the CRTPRNDIR parameter.

The values you may set in PRNDIROWN are *PARENT and name. *PARENT means that the owner of a parent directory being created by the restore will be the same as the owner of the directory it is being created into. Name allows you to specify the name of a user profile to be the owner of any parent directories that are created by the restore.

7.4.2 IFS restore performance improvements

Internal code changes were made to improve performance. This includes utilization of multiple threads. The System Level Internal Code (SLIC) now has an optimization that reduces disk operations. You should refer to the V5R4 Performance Capabilities Reference, available online at:

<http://www-03.ibm.com/servers/eserver/series/perfmgmt/resource.html>

7.5 New SAVSYSINF and RSTSYSINF commands

The Save System Information (SAVSYSINF) command performs a partial save of data saved by the Save System (SAVSYS) command. This is a cumulative save since the last SAVSYS. SAVSYS requires the system to be in a restricted state. This is not required for SAVSYSINF.

The save includes:

- ▶ System reply list entries
- ▶ Certain system values
- ▶ Some system values are not saved
- ▶ Not required for system recovery
- ▶ Service attributes
- ▶ Network attributes
- ▶ Environment variables
- ▶ PTFs for 5722SS1 and 5722999 applied (TEMP or PERM) since last SAVSYS

- ▶ All group PTFs
- ▶ All PTF cover letters
- ▶ Loaded PTFs copied into *SERVICE

The SAVSYSINF command saves the following object types in QSYS:

- ▶ *DTAARA
- ▶ *JOBBD
- ▶ *JRN
- ▶ *CLS
- ▶ *IGCTBL
- ▶ *MSGF
- ▶ *JOBQ
- ▶ *JRNRCV
- ▶ *MSGQ 1
- ▶ *EDTD
- ▶ *SBSD
- ▶ *TBL
- ▶ *CMD 1

One object changed since the last SAVSYS.

The SAVSYSINF command does not save the following:

- ▶ Licensed Internal Code
- ▶ QSYS library
- ▶ System values that cannot be changed or saved, or are related to date and time or password level
- ▶ Configuration objects (use SAVCFG)
- ▶ Security Data (use SAVSECDTA)

Loaded PTFs are only saved if copied into the *SERVICE area. This enables SAVSYSINF to find the save files of the PTFs. The Copy PTFs (CPYPTF) service attribute specifies whether to copy PTF save files into *SERVICE when PTFs are loaded from a tape or optical device. Use the CHGSRVA CPYPTF(*YES) command to change the service attribute on your system to copy PTF save files when loading PTFs from media.

7.5.1 SAVSYSINF and RSTSYSINF considerations

Base SAVSYS is still required. SAVSYSINF is not a replacement for SAVSYS. *ALLOBJ or *SAVSYS Special Authority is required for this command. It cannot be used for restoring to another system, or for system upgrades or migrations.

The command is intended for customers who cannot bring the system to the restricted state required for SAVSYS, or who cannot take downtime for SAVSYS.

PTF save files must remain on the system until the next SAVSYS. You do not want to run the Delete Program Temporary FIX (DLTPTF) command except right before or after SAVSYS. The default when installing PTFs is not to restore save files. You can either change the system default CPYPTF(*YES) on the CHGSERVA command, or use it in the INSPTF command consistently. This can increase the time and complexity of saves.

We recommend that a save of the entire system including a SAVSYS be done in restricted state. This can be accomplished by performing a Go Save Option 21, a combination of an Option 22 and 23, or by using the equivalent functions within BRMS.

Recall that, starting with V5R3, you can submit a BRMS function to the batch queue of the QCTL subsystem.

7.5.2 SAVSYSINF backup strategy

After the base SAVSYS, perform these commands to save system changes:

- ▶ SAVLIB LIB(*IBM)
- ▶ SAV OBJ('/QIBM/ProdData')('/QOpenSys/QIBM/ProdData') UPDHST(*YES)
- ▶ SAVSYSINF

Use these commands to save user data:

- ▶ SAVESECDTA
- ▶ SAVCFG
- ▶ SAVLIB LIB(*ALLUSR)
- ▶ SAVDLO DLO(*ALL) FLR(*ANY)
- ▶ SAV OBJ('//*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT) UPHST(*YES)

An example SAVSYSINF command to save file is:

```
SAVSYSINF DEV(*SAVF) SAVF(QGPL/SAVF) CLEAR(*ALL) OUPUT(*OUTFILE)
OUTFILE(QGPL/OUTPUT)
```

This command example saves the system information to the save file named SAVF in library QGPL. The save file will be cleared automatically. Information about what was saved will be written to the first member of the file name OUTPUT in library QGPL. The file and member will be created if they do not exist.

7.6 Changes to SAV commands

This section highlights the new and changed parameters on the group of i5/OS save commands.

7.6.1 SAVLIB, SAVOBJ, and SAVCHGOBJ

We have made changes to the SAVLIB, SAVOBJ, and SAVCHGOBJ commands. You can now use the user space (*USRSPC) value on LIB, OMITLIB, and OMITOBJ parameters. You can specify a user space/library name on the CMDUSRSPC parameter.

A user space can now have up to 32767 specified values for each parameter. They are built in the same format (key values) as QSRSAVO Save Object API. The previous option to specify 300 objects or generic object values is still allowed.

7.6.2 Save data queue contents

The new QDTA parameter on SAVLIB, SAVOBJ, and SAVCHGOBJ commands allows for:

- ▶ *NONE (default)
 - Only description of data queue saved.
- ▶ *DTAQ
 - Description and data queue contents of standard data queue saved.
 - Only description of Distributed Data Management (DDM) data queue saved.

This is supported on the QSRSAVO API.

7.6.3 Parallel library save enhancement

Prior to V5R4, in order to save multiple libraries in parallel format, the user or application needed to create a data area named QSRPARFMT in library QTEMP. The save format can now be specified in the media definition. Again, in practical use you use BRMS to do this.

QSRCRTMD, QsrCreateMediaDefinition API, allows the addition of save format:

- ▶ -2 Value (default)
 - One library specified, saved in parallel
 - Multiple libraries specified, saved serially

Note: If DTAARA QSRPARFMT exists, all libraries are saved in parallel.
- ▶ -1 Value
 - All libraries saved serially
- ▶ 0 Value
 - All libraries saved in parallel

SAVLIB (*ALLUSR, *IBM, *NONSYS) is still saved serially.

7.7 Backup and Recovery Media Services

i5/OS's strategic backup management product, BRMS (5722-BR1), consolidates all of your backup tasks. It manages your media, automates your backups, and simplifies recovery. It provides detailed reporting on what was saved and what was not saved. It produces detailed instructions on the recovery process.

BRMS supports V5R4 i5/OS virtual tape support and, with the availability of an RPQ, supports The IBM Virtualization Engine TS7510 virtual tape hardware solution.

Figure 7-1 is a representation of the base, network, and hierarchical storage components of BRMS.

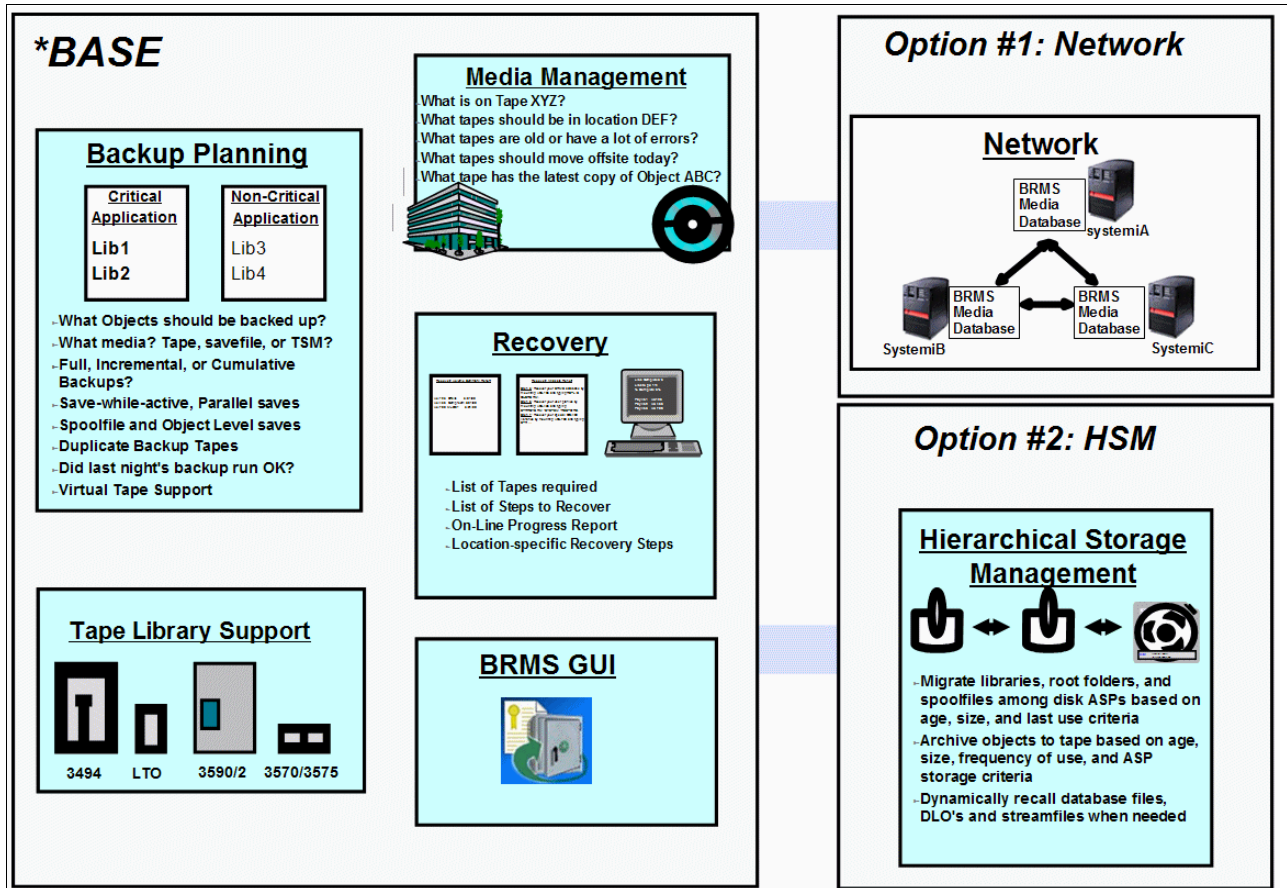


Figure 7-1 BRMS product overview

The *BASE option is always required. It allows you to save and restore all functions to tape media, and save files, Tivoli® TSM servers, and virtual tapes.

The *NETWORK function is optional and separately priced. By placing multiple iSeries systems in a BRMS network, you can share BRMS policies, media information, and storage locations across the network group. This allows you to manage backups across all of your iSeries systems in a consistent manner. It also optimizes the use of your media.

The Hierarchical Storage manager (*HSM) function is optional and separately priced. Hierarchical Storage Management (HSM) automatically and transparently manages customer data across a storage hierarchy. The storage hierarchy can consist of high-performance disk, compressed disk, and tape libraries. Through the BRMS user-defined policies, HSM can migrate or archive and dynamically retrieve infrequently used data or historical data up or down a hierarchy of storage devices.

For an in-depth review of changes and functionality, refer to the Infocenter or the manual *iSeries Backup Recovery and Media Services for iSeries Version 5*, SC41-5345. This can be found online at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/books/sc415345.pdf>

7.7.1 BRMS iSeries Navigator GUI client

The iSeries Navigator has had numerous additions and enhancements in V5R4.



Figure 7-2 BRMS menu in iSeries Navigator

Archive options

Archive options have been added to the Navigator. Archiving provides a method of saving disk space by saving infrequently used objects to media, then removing them from the disk. During the archive operation, the storage on the primary disk is freed, allowing you to store more vital data in that disk space. Archiving data to less expensive forms of media, like tape or save files, provides a more cost-effective way of handling legacy data, such as past financial records. You have the benefit of retaining this data, but without an impact to the storage capacity of your primary disk space. BRMS tracks the information associated with the objects that you archive.

The *create an archive* policy launches a wizard to walk you through the steps required to make a new policy for your date. You can define what you want to archive, under what conditions, and when you want to run it. You have the opportunity to prepare your archive media. Once created, an archive policy can be reused as often as you wish, or it can be modified as needed.

The two other options allow you to monitor these policies. Lists are just that — lists where you have the option of filtering the results to include what you want to see. From the list, you can copy, edit, or delete specific policies.

The *view archives policies* report generates an HTML BRMS Archive Policy Report, which will launch in a browser window. This fully hyperlinked report provides extensive detail regarding the policies, what they save, when they ran, and the condition of the system before and after the archival process.

E-mail policy

In the *view or edit global policies* properties, you can control the networking options for BRMS. By placing multiple systems in a BRMS network, you can share BRMS policies, media information, and storage locations across the network. This allows you to manage saves across all of your systems in a consistent manner.

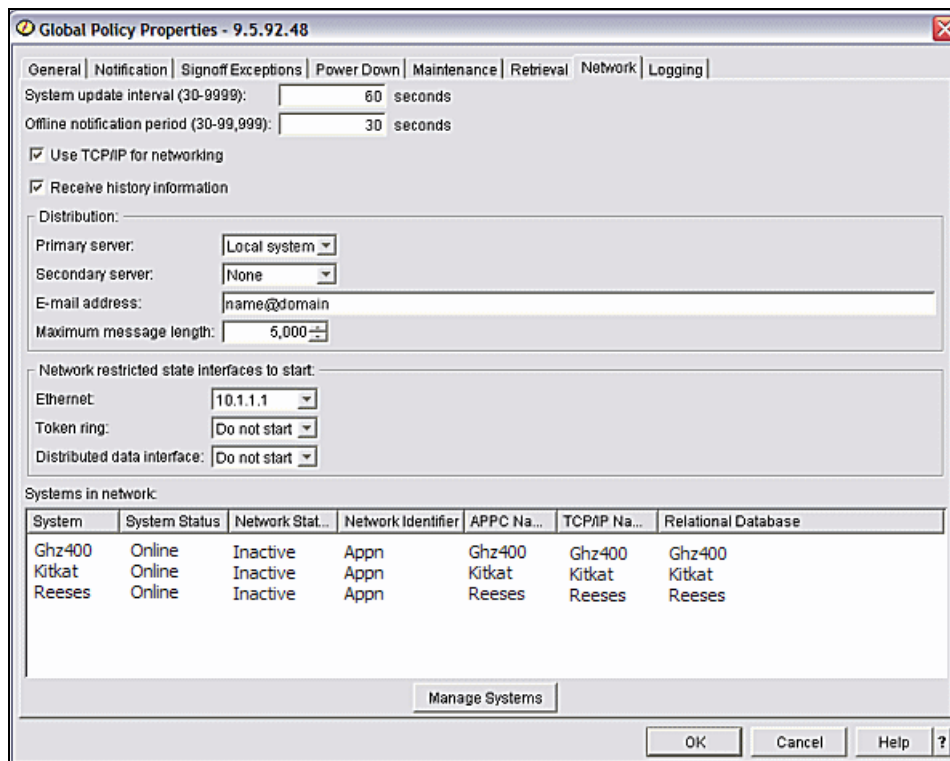


Figure 7-3 Network tab of global policies properties

The Network tab contains two functions. The first addresses the fact that, when in restricted state, it may be necessary for BRMS to communicate with other systems in the network. You can specify which network interfaces should be started when the current system is in restricted state. Only one network interface of each type can be started.

The other part of this page deals with the setup of e-mail distribution of messages sent to the BRMS log. You can select specific messages for distribution on the Logging page.

The system that you select as the primary server will distribute these messages. If the primary server cannot complete the distribution, it sends the message to the secondary server for distribution. For example, you could select a local primary server and a remote secondary server. When the local system is in restricted state, BRMS sends the e-mail through the remote secondary system. If a secondary server is not specified and the primary server is in restricted state, the message will be queued and distributed as soon as the BRMS networking is restarted.

Use the Logging page to specify how messages sent to the BRMS log are managed. You can choose to include or exclude specific messages from the log.

Containers support

You can also manage containers, which are any holders that store multiple media volumes, with Backup, Recovery and Media Services (BRMS).

Containers usually hold one media pool, although you can specify any media pool. Using containers adds complexity to your storage management system and is optional. You can specify that you want to store media in physical containers at specified locations, and use these containers to move volumes in the container according to a move policy.

From this panel you can complete the following tasks:

- ▶ Create a new container.
- ▶ Delete a container.
- ▶ Edit a container.
- ▶ Manage a container pool.

User media library (MLB) support

Under Manage devices, you can now use the New User Media Library dialog to define third-party (non-IBM) media library devices to use with BRMS. To use third-party media libraries, you must enter the commands that enable functions for these libraries. For specific commands for these functions, consult the documentation that is provided with the third-party media library. Note that all user media library commands must be synchronous.

Virtual volume support

While we cover virtual tape in Chapter 8, “i5/OS-based virtual tape” on page 169, there are additions that we mention here that allow you to control this feature.

The BRMS volume information now shows the image catalog of virtual volumes. In the add media wizard, you are prompted for image catalog selection. You can now add virtual volumes to the BRMS inventory.

Save joblogs

We go in-depth regarding saved job logs, and their potential impact on your system in 15.4.2, “Action plan: gaining control over your output queue” on page 387. We would like to mention that BRMS job logs may be a good category of job logs to keep.

*LINK Omit

BRMS uses the *LINK Omit to exclude directories from save operations. The Navigator makes it very easy for you to select these directories. From the Manage Lists panel, select **Edit** and then browse to the directories you wish to include.

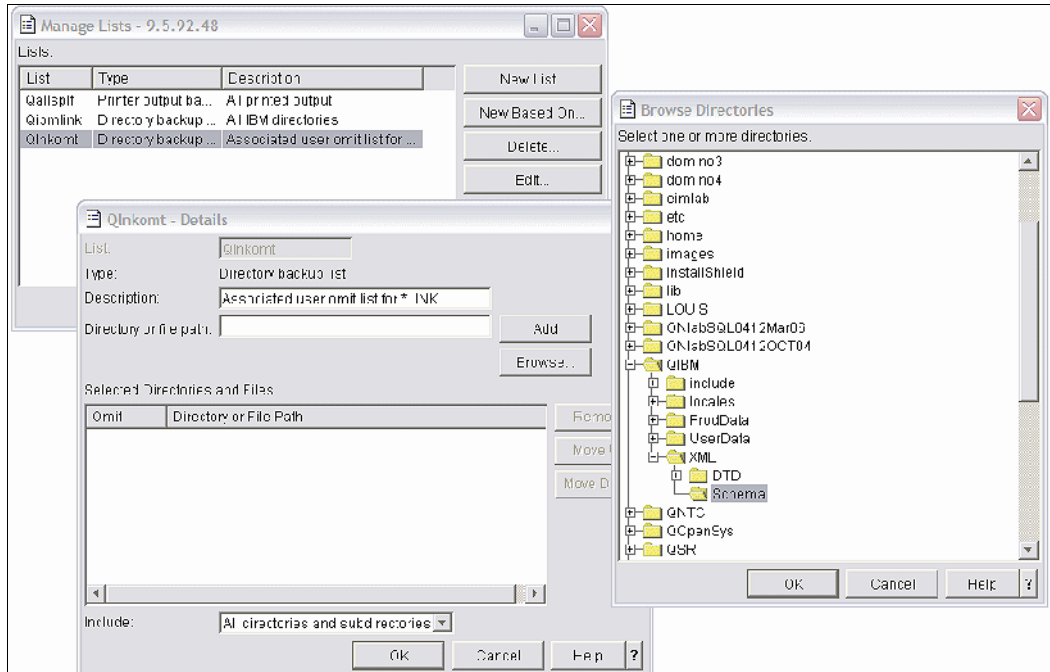


Figure 7-4 *LINK Omit selection

Differences in the BRMS interfaces

Table 7-1 is a matrix of differences between the character-based (5250) and GUI (Navigator) interfaces.

Table 7-1 Interface differences

BRMS plug-in to iSeries Navigator	Character-based interface	Definition
Backup policy	Backup control group and media policy	A backup policy is the combined attributes of a backup control group and the media policy in the character-based interface. Both of these functions control how, when, and where a backup is performed.
Changes-only save	Non-incremental cumulative backup	Saves only those items that have changed since the last full save was performed.
Media pool	Media class	A grouping of media based on similar characteristics, such as density and capacity.
Disk pool	Auxiliary storage pool	A software-defined group of disk storage units on a system.
Disk pool group	Auxiliary storage pool class	A group of disk pools that are organized by similar purpose.
Container pool	Container class	A grouping of similar types of containers that hold similar media volumes.

Compatibility differences

If you are currently using the character-based interface control groups and they do not reference the system policy (*SYSPCY), archive policy (*ARCPCY), or backup policy (*BKUPCY) and do not share media policies among control groups, then using the BRMS plug-in should have no functional impacts to your character-based control groups. If this is not true, read the following information.

The backup and archive policies created by the BRMS plug-in can be viewed and run from the character-based interface. Also, control groups created from the character-based interface can be viewed and edited via the BRMS plug-in backup and archive policy properties. A message will be presented to the user when a character-based created or edited control group is about to be changed by the BRMS plug-in. The user will have the opportunity to cancel the operation or continue. However, changing control groups created on the character-based interface using the BRMS plug-in is not recommended unless you understand the following:

- ▶ The BRMS plug-in backup and archive policies are designed to be independent of each other. Thus, if you make changes to one policy, it will not affect the others. To ensure this independence, control groups updated with the BRMS plug-in will have all references to the character-based interface system policy (*SYSPCY), archive policy (*ARCPCY), and backup policy (*BKUPCY) removed. These references will be resolved and replaced with the actual values when the backup policy is saved by the BRMS plug-in. The archive policy, backup policy, or system policy do not affect backup or archive policies created or edited by the BRMS plug-in. The exceptions to this are the network function, sign-off exceptions, and subsystem to check controls in the system policy, which are used by the BRMS plug-in.
- ▶ The BRMS plug-in hides the media policy from the user (no BRMS plug-in functions exist to show or edit media policies). Media policy attributes are indirectly shown and edited through backup and archive policy properties. As mentioned earlier, the BRMS plug-in backup and archive policies are designed to be independent of each other. To ensure this, the BRMS plug-in will create a new media policy based on the currently used one if the control group being edited contains a media policy that is used by other control groups or is a shipped BRMS media policy. If a new media policy is created, the control group will be changed to use the new one.
- ▶ When you create a new backup or archive policy using the BRMS plug-in, it will use a new media policy that is dynamically created.
- ▶ Changes made to other, non-IBM- supplied policies are reflected in the BRMS plug-in.

7.7.2 Backup enhancements

In this section we discuss backup enhancements.

Control group attributes

Control groups consist of groups of libraries, objects, folders, spooled files, and other types of data that share something in common. Perhaps it is the characteristics of the items, perhaps it is just that you want to group them for backup purposes. While policies determine how data is processed, control groups determine what is saved. Control groups have been enhanced to allow you more flexibility.

Exit program

We have added an option to invoke an exit program. You use this exit to receive information about entries processed in control groups. You may want to use this exit to:

- ▶ Perform additional actions for processed entries.
- ▶ Review the messages generated by processed entries.
- ▶ Retry processed entries.
- ▶ Force termination of subsequent processing.

See *Backup Control Group Backup Item Exit Program in the Backup, Recovery and Media Services book*, SC41-5345, for additional details about required parameters for the exit program.

New missed object policy

This allows BRMS to build a control group with a list of missed objects. This control group can then be re-run again or looked at. This can be used with multiple control groups using the same missed objects control group.

These are also supported in the properties for a backup policy in the GUI client.

Library lists in object lists

Backup object lists can be used as library lists. Library lists can be useful to ensure that the libraries in the list are being saved using the same save while active attributes.

You can create a library list by first creating a backup object list. For each library or generic library name you enter for library, also specify *ALL for object, type, and member. When the list is run as part of a backup, BRMS will start with the first entry in the list and group subsequent entries, which specify *ALL for object, type, and member into a single save command until it reaches an entry where *ALL is not specified for object, type, and member, or until the number of libraries exceeds the limit of the library (LIB) parameter of the Save Library (SAVLIB) command.

Normal backup object list processing occurs for entries where *ALL is not specified for object, type, and member. Library list processing resumes with the next subsequent entry where *ALL is specified for object, type, and member.

Note: Normal backup object list processing occurs when saving to save files or TSM servers since these media types only support saving one library at a time.

Force full backup days

In your backup policy, you can select to force full backup days. This parameter will specify when a full backup should occur relative to the last backup. This option is reflected in both interfaces.

Save system information (SAVSYSINF)

As noted in 7.5, “New SAVSYSINF and RSTSYSINF commands” on page 150, we have added support for this new function. Again, SAVSYSINF is not to be considered a replacement for the SAVSYS command and is not to be used for system upgrades or migrations.

Native spool file save/restore

As noted in 7.2, “Spooled file support” on page 148, we have added the capability of natively backing up and restoring your spooled files. BRMS provides full support of this function, with all the control that you have come to expect from the product.

FLASHCOPY support

BRMS is integrated with SAN for the FlashCopy solution. FlashCopy is an optional function available with IBM TotalStorage products. From an i5/OS viewpoint, FlashCopy copies all storage from a single system (or partition) on a single TotalStorage server to a separate set of disks in a few seconds.

To get the most complete information captured you must quiesce i5/OS activity before copying starts. This would include performing the following:

- ▶ ENDSBS *ALL to get the partition to a restricted state
- ▶ From the console in subsystem QCTL issuing the Clear Storage Pool (CLRPOOL) command for all active storage pools
- ▶ Re-IPLing the system

Note that there is research going on to determine whether the IPL step is required.

Once the FlashCopy has transferred the files, BRMS will back up the copied system as if it was the production system doing its normal backup.

The BRMS Networking Enhancement stops other systems from accessing BRMS data from the production system until BRMS media information is returned.

The iSeries Copy Services Toolkit is available and blends the technologies provided by Independent Auxiliary Storage Pools (IASPs) along with the advanced functions provided by IBM TotalStorage Copy Services

For more Information see IBM Redbook *iSeries and IBM TotalStorage: A Guide to Implementing External Disk on eServer i5*, SG24-7120, found at:

<http://www.redbooks.ibm.com/abstracts/sg247120.html>

For the latest information about i5/OS and TotalStorage products working together, refer to:

<http://www.ibm.com/eserver/series/availability>

Parallel save/restore enhancements including support for IFS

As covered in 7.4, "IFS parallel save/restore support" on page 149, parallel save and restore functions have been enhanced, including support for IFS. BRMS has fully implemented these enhancements.

7.7.3 Recovery and retrieval enhancements

In this section we discuss recovery and retrieval enhancements.

Start recovery using BRMS

Start recovery using BRM (STRRCYBRM) has had multiple enhancements.

BRMS continues to enhance the Recovering the Entire System (QP1ARCY) report. New steps were added along with new attention notices. These help to identify problems in advance, such as libraries not saved, which could affect your recovery.

The *RESUME option was enhanced to allow sharing of saved items across jobs. This allows you to run concurrent recoveries.

The use duplicate media (USEDUPMED) parameter was added. This allows you to do a restore or retrieve from any volume. This is especially helpful when the original volume is in off-site storage and the local copy is a duplicate.

You can create special recovery instructions. If you wish to include these instructions in the recovery report, you would indicate this through the user recovery information (USRRCYINF) parameter.

Restore objects and libraries with BRMS

The RSTOBJBRM/RSTLIBBRM have had important enhancements.

Special value *SAVDATE was added to the save level (SAVLVL) parameter. This is used in conjunction with the new save date (SAVDATE) parameter to specify by date the level of the library to be restored.

With regard to the information covered in 7.2, “Spooled file support” on page 148, the spool file data (SPLFDTA) parameter was added to these commands. This allows restoration of spooled files, which were saved concurrently with saved output queues.

As noted in 7.4.1, “Automate parent directory creation” on page 150, we have added create parent directories (CRTPRNDIR) and parent directories owner (PRNDIROWN). These parameters allow you to better manage the restoration of IFS objects.

7.7.4 Devices/media management enhancements

In this section we discuss devices/media management enhancements.

Native virtual tape support

We cover virtual tape support in Chapter 8, “i5/OS-based virtual tape” on page 169, so we will only discuss it briefly here.

Virtual tape can be used to shorten the backup window. Also, you can eliminate data and equipment checks during critical backup windows. Virtual tape devices can be used for all save/restore, copy, tape utility functions, virtual I/O, and high-level language interfaces. For instance, SAVSYS requires restricted state, with only subsystem QCTL active. Starting with V5R3, SAVSYS can be performed in batch mode in subsystem QCTL, so SAVSYS to a VT device is supported. There is no support for i5/OS install, SAVSTG, or use by SST/DST.

Parallel saves are supported with virtual tape. You just need to add virtual tapes and associated media class into the BRMS inventory and BRMS will use them.

Once the backup is complete, you can resume operations and do duplicate media using the BRMS (DUPMEDBRM) command in the background. Then you can keep virtual tape for restoring files needed to recover during the week and send real tapes to vault.

Considerations

Consider the following:

- ▶ Virtual tape can consume large amounts of disk.
- ▶ You must create virtual tape devices outside of BRMS.
Use the Create Device Tape (CRTDEVTAP) command.
- ▶ Must create image catalog and virtual volumes outside of BRMS. Use:
 - Create image catalog (CRTIMGCLG).
 - Add image catalog history (ADDIMGCLGE).

- iSeries Navigator.
- ▶ Virtual tape acts just like a physical tape device.
 - BRMS automatically performs virtual tape mounts.
- ▶ For disaster recovery, remember:
 - Duplicate to physical media (DUPMEDBRM).
 - Physical tapes become the original save volume.
 - Virtual tapes become duplicated save volumes.
 - Virtual tape volumes are stored as IFS objects.
 - Not saved during normal save
 - No virtual volumes in image catalog after disaster recovery
- ▶ Virtual media volumes are not shared in a BRMS network.
 - Local system only.
 - IASPs allow sharing by switching IASP.

Virtual tape support of the IBM Virtualization Engine TS7510

The IBM Virtualization Engine TS7510 (sometimes referred to as the TS7510 Virtualization Engine) is a member of the IBM Virtualization Engine TS7000 Series of virtual tape libraries intended for the Small and Medium Business (SMB) environment. The TS7510 Virtualization Engine combines hardware and software into an integrated solution designed to provide tape virtualization for a wide variety of operating systems connecting over Fibre Channel physical connections.

The TS7510 is an ideal choice as part of your backup choices when you have an environment where it can be used by several hardware platforms, operating systems, and partitions at the same time.

The TS7510 Virtualization Engine combines IBM server and disk technology and is designed to virtualize or emulate tape libraries, tape drives, and tape media. Real tape resources can then be attached to the TS7510 Virtualization Engine to help address information life-cycle management and business continuance.

Each operating system communicating with the TS7510 thinks it is performing I/O operations and managing media with a tape library and attached tape drive. The Virtualization Engine for Tape Console interface is used to define the virtual tape libraries and associated virtual tape devices.

i5/OS releases V5R2, V5R3, and V5R4 support the TS7510 via an RPQ. To order this RPQ contact your local IBM System Storage representative. An IBM Redbook, number SG24-7510, on i5/OS support of the TS7510 is scheduled to be available January 2007.

We recommend using BRMS to take full advantage of the TS7510 capabilities.

Worm (write once read many) device support

Change media class interfaces now have an option for write once media. BRMS will treat all media in this class as WORM when set to *YES. Media so marked can only be written to once. You can continue to append data until the media is full, then it is done. Some drives will ensure this feature also.

This is important in certain regulatory compliance situations. If you need a clear history of changes, for instance, you can store them on this type of media, and they can be reviewed, but not altered.

```

Change Media Class
Media class . . . . . : VIRTAP
Type choices, press Enter.
Density . . . . . *VRT256K F4 for list
Media capacity . . . . . *DENSITY *DENSITY, Number nnnnn.n
Unit of measure . . . . . 1=KB, 2=MB, 3=GB
Mark for label print . . . . . *NONE *NONE, *MOVE, *WRITE
Label size . . . . . 1 1=6 LPI, 2=8 LPI, 3=9 LPI
Label output queue . . . . . *SYSPCY Name, *SYSPCY, *PRTF
Library . . . . . Name, *LIBL
Shared media . . . . . *NO *YES, *NO
Write once media . . . . . *NO *YES, *NO
Text . . . . . TEST VIRT

More...

F3=Exit F4=Prompt F12=Cancel

```

Figure 7-5 WORM selection

Add cartridge type to displays

We have added a Cartridge Type column to the *work with media library media* using the BRM (WRKMLMBRM) command and *add media library media* to BRM (ADDMLMBRM) command.

```

Work with Media Library Media MYSYS
Media library . . . . . : TAPMLB01
Position to . . . . . Starting characters
Type options, press Enter.
1=Add MLB media 2=Work with media 5=Initialize 6=Change category
7=Eject 8=Mount 9=Demount
Media
---BRM Information---
Opt Volume Type Category Media Class Expired Status
___ 000B84 H *SHARE400 *NONE Available
___ 000C3C G *SHARE400 FMT3570E *YES Available
___ 000D63 G *NOSHARE FMT3570E *YES Available
___ 000FC9 G *NOSHARE FMT3570E *YES Available
___ 000FF8 G *NOSHARE FMT3570E *YES Available
___ 0003BC G *INSERT *NONE Inserted
___ 001019 G *INSERT FMT3570E *YES Inserted
___ 0784AE H *SHARE400 FMT3570 *YES Available

More...

F3=Exit F5=Refresh F9=System command F12=Cancel F13=Repeat

```

Figure 7-6 Media Type column

Tivoli Storage Manager (TSM)

In most cases, BRMS manages the retention and deletion of objects saved to TSM servers. When media retention expires, BRMS deletes the stored data from the TSM server and from the save history located in BRMS.

If the TSM administrator has enabled device retention protection in the TSM management class, then the TSM server restricts data from being deleted. If you are saving objects to a TSM server that uses retention protection, for example, when using the IBM DR450/550, the media retention that you specify in BRMS must match the retention that is defined in the TSM management class. This is required so these objects can be deleted from BRMS history and the TSM server concurrently at expiration.

Version V5R1 or later of the Tivoli Storage Manager Application Programming Interface (TSM API), program product 5733-197, is required to use retention protection.

Before setting this retention protection for TSM devices on BRMS, verify whether the TSM server uses device retention protection with your TSM administrator.

DR450/550 support (data retention)

We have integrated TSM device support for TSM servers that have been pre-configured for data retention protection. The IBM Data Retention 450 (DR450) or DR550 servers are customer setup solutions.

BRMS has been enabled to allow backup or archive operations to be directed to this TSM server type. This function is available by PTF in releases prior to V5R4M0.

New TSM performance report

We have added a TSM Performance Report (Figure 7-7).

BRMRPT	Reports	System: RCHASM27
Select one of the following:		
46. Print migration candidates		
47. Print migration information		
48. Print migration log		
Other		
49. Print system policy		
50. Print calendars		
51. Print centralized media audit		
52. Print BRMS log		
53. Print TSM performance analysis		
Selection or command		Bottom
===>		
F3=Exit	F4=Prompt	F9=Retrieve
F10=Commands	F12=Cancel	F13=Functions

Figure 7-7 Go BRMRPT, note Option 53

DUPMEDBRM several enhancements

Duplicate media using BRM (DUPMEDBRM) has a new smart append function. There are two new volume (VOL) values:

- ▶ *RESUME was added to allow you to resume a pending duplication operation that ended abnormally.

This special value is used in conjunction with the new Resume key (RSMKEY) parameter to identify the specific duplication operation to be resumed.

- ▶ *SCHHST allows duplication of marked saved items in the history information.
This is nice for when appending saves to tape.

```

Specify Value for Parameter VOL
Type choice, press Enter.
Type . . . . . : CHARACTER
From volume identifier . . . . . > _____
:
:
*LIST          :          From volume identifier (VOL) - Help          :
=> *RESUME     : *RESUME                                                           :
=> *SCHHST     : Specify to resume history or media duplication                   :
*SEARCH       : operations which end with errors. Use the Resume key           :
              : (RSMKEY) parameter to specify the duplication                   :
              : operation to be resumed.                                         :
              : Note: You must correct any errors which caused                 :
              : the duplication operation to end before you resume             :
              : the duplication                                                 :
              : *SCHHST                                                         :
              : All saved items that are marked for duplication, and           :
              : meet the specified search values are duplicated.               :
              : Search values are specified by using the following             :
              : parameters:                                                     :
              : o File group (FILEGRP) parameter                               :
              : o File group type (GRPTYPE) parameter                          :
              : o Media class (SCHMEDCLS) parameter                             :
              : o From system (FROMSYS) parameter                               :
              :                                                                 :
              :                                                                 :
              : F2=Extended help   F10=Move to top   F11=Search Index         :
F3=Exit  F5= : F12=Cancel       F20=Enlarge     F24=More keys         :
              :.....:

```

Figure 7-8 Specify value for parameter VOL

The special value *LEAVE was added to the list of special values for the from end option (FROMENDOPT) and to end option (TOENDOPT) parameters for consistency with other commands.

```

                                Duplicate Media using BRM (DUPMEDBRM)
Type choices, : .....
                :          From device end option (FROMENDOPT) - Help          :
                :                                                                 :
From volume i : *LEAVE                                                           :
From device   : The volume does not rewind or unload after the                 :
To device    . : operation ends. It remains at the current position on         :
From sequence : the device.                                                     :
    Starting fi :                                                                 :
    Ending file :                                                                 More... :
To sequence n : F2=Extended help  F10=Move to top  F11=Search Index           :
To media poli : F12=Cancel        F20=Enlarge     F24=More keys              :
Input volume  :                                                                 :
                : .....
From device end option . . . . . FROMENDOPT    *LEAVE  <==
To device end option . . . . . TOENDOPT       *LEAVE
From system . . . . . FROMSYS                 *LCL
                                                More...
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 7-9 *LEAVE value

The Mark history for duplication (MARKHST) parameter was added to allow you to remark duplicated saved items for a subsequent duplication.

The save media information (SAVMEDI) parameter was added to allow you to save the BRMS media information after the duplication operation has successfully completed.

The options for both of these values are *MEDPCY, *NO, and *YES. *MEDPCY is dependant on the entry in the media policy (MEDPCY) parameter. If *NONE is specified in the MEDPCY parameter, you cannot specify *MEDPCY in this parameter.

```

Duplicate Media using BRM (DUPMEDBRM)
Type choices, press Enter.

                                Additional Parameters

To volume identifier . . . . . TOVOL          *MOUNTED
                                + for more values
Retain object detail . . . . . OBJDTL        *NO
Expiration date . . . . . EXPDATE           *MEDPCY
Move policy . . . . . MOVPCY                *MEDPCY
Media class . . . . . MEDCLS                 *MEDPCY
Location . . . . . LOC                      *MEDPCY
Secure volume . . . . . VOLSEC               *MEDPCY
Required volumes . . . . . MINVOL           *MEDPCY
Mark volumes for duplication . . MARKDUP     *MEDPCY
Mark history for duplication . . MARKHST     *MEDPCY <==
Save media information . . . . . SAVMEDINF    *NONE <==

                                                                Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys

```

Figure 7-10 Mark history and save media information

7.7.5 Miscellaneous enhancements

In this section we review miscellaneous enhancements.

Network enhancements

BRMS iSeries Navigator Client Enhancements BRMS continues to move character-based interface functions into the graphical user interface provided by the BRMS iSeries Navigator and to improve the existing BRMS client functions.

We now support available/unavailable status and user control of message sending from the GUI. There is a restricted state network interface as well. We also have library manager support for 3584 tape libraries.

Archive

The start archive using BRM (STRARCBRM) command selects a control group and begins the archive process. You can start the archive immediately or you can schedule it using the system job scheduler. Support of save media information about STRARCBRM has been added.

Saved objects

We added an ASP parameter to work with saved objects (WRKOBJBRM).



i5/OS-based virtual tape

Virtual storage consists of objects that, when used together, imitate tape, CD, DVD, and write-once read-many (WORM) media on your disk units. The imitated media appear to the server to be actual media. In this chapter we discuss enhancement to virtual tape support.

8.1 Virtual tape support

There are some distinct advantages to the use of virtual tape. There are also considerations you need to take into account before you add this function to your backup arsenal. We cover these topics, along with implementation and management details, in this chapter. We would like to note that there is a separate IBM Redbook draft available: *i5/OS V5R4 Virtual Tape: A Guide to Planning and Implementation*, SG24-7164.

8.2 Key advantages

The primary uses for virtual tape are:

- ▶ When you have lots of information to save during the same time period.
- ▶ You do not have a large number of tape devices to do the saves directly.
- ▶ You need extra capabilities not available with a save to a save file.

Virtual tape saves can be faster than saving directly to tape, with performance similar to save files. You will get the best performance in a separate ASP.

Virtual tape is supported on all save/restore commands, virtual I/O, and APIs (except save storage - SAVSTG).

Many of the save file restrictions are supported with virtual tape, as you can see in Table 8-1.

Table 8-1 Save file versus virtual tape

Function	Save file support	Virtual tape support
Number of libraries per save file	One	Multiple
SAVSYS	No	Yes
Parallel saves	No	Yes
1 TB size limitation	Yes	No

Virtual tape eliminates tape device or media errors causing your save to fail.

The ability to then duplicate saves to media with the duplicate tape (DUPTAP) or the duplicate media using BRM (DUPMEDBRM) command allows you to save at your convenience, when the tape devices are available.

There are on-site and off-site storage benefits, as well. You can keep virtual volumes on the system as needed. You can send the duplicated volumes to an off-site storage facility for safety. However, with the virtual tapes on disk, you need not wait until the off-site media can be delivered in order to recover. Nor will you incur fees for emergency service.

There is full support by the Backup Recovery and Media Services (BRMS) product, 5722-BR1, as we discussed in Chapter 7, "i5/OS-based backup and recovery" on page 147.

8.3 Considerations

There are considerations you should take into account before implementing the new virtual tape support in your environment. You must manage objects created when using virtual tape

support. Failure to do so can result in your disk storage becoming significantly consumed with unmanaged files.

There are additional disk requirements. How much is something best worked out before you commit to a strategy including virtual tape.

There are no interface points through dedicated service tools (DSTs) or system service tools (SSTs). This means that functions like dump to media from DST/SST are not supported.

There is no ability to install from SAVSYS virtual volume, as a D-IPL can only be from media. While you cannot restore the operating system from virtual volumes, you can do restore functions, such as restore license program (RTSLICPGM). Remember that, although you cannot install from a save system to a virtual tape image, starting in V5R3 you could perform a SAVSYS in batch mode in subsystem QCTL, when the system is in a restricted state.

While it may be faster, it may also not be faster. This all depends on your tape technology, system configuration, and environment.

Data compaction is not supported, and data compression (DTACPR) support has areas to note. First, you must change DTACPR to *YES in the parameter on save commands, as it defaults to *NO. The compression supported is SNA low data compression, and this can have a significant performance impact.

There is a save-while-active checkpoint restriction, in that, once a checkpoint is reached, the saves cannot be restarted.

8.4 Overview and flow

Figure 8-1 is a good graphical representation to have in mind when learning about i5/OS virtual tape support. The objective is to save to virtual tapes that are really storage areas on the i5/OS partition. No physical tape devices need to be attached during backup (or restore).

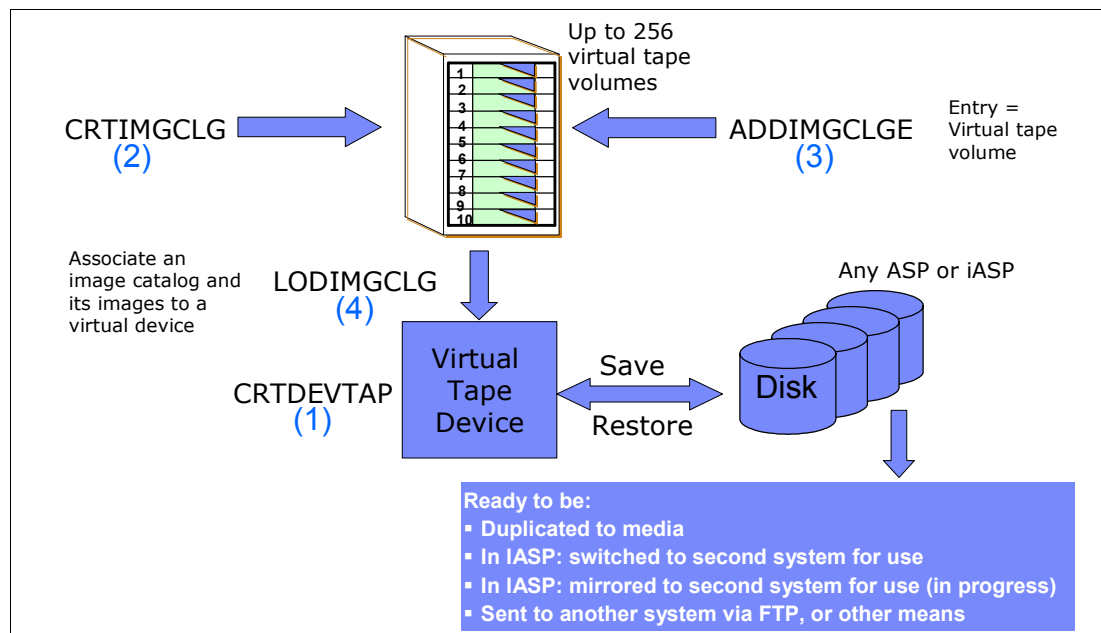


Figure 8-1 Overview and flow

The numbers represent the steps in setting up to use i5/OS virtual tape support:

1. Create a tape device description with a virtual attribute.
2. Create a tape image catalogue.
3. Add image catalog entries (up to 256) to the image catalogue. Entries represent a virtual tape volume.
4. The load/unload image catalog (LODIMGCLG) command is used to associate an image catalog and its images to a virtual device.

After this you can save to the virtual tape or restore from it. Then you can copy the virtual tape image catalog entry data to an actual tape device or other media, such as sending it to a second system.

If the catalog is in an Independent Auxiliary Storage Pool (IASP), the IASP contents can be switched to another system or be mirrored to a second IASP on another system.

As a reminder, you do need to manage this storage on the original system as well as on a second system if the data is switched or duplicated to a second system. If switched or duplicated onto a second system, you do have to manage any function you want to perform. You also have to manage any new *tape catalog image entry* versus previously created entries.

An example of this would be if you do the actual save to media on the second system. You must determine when and how to use the new entry, as well as the space for multiple image entries, and when to discard the data. You may want to perform these functions manually first and then do them programmatically or use the BRMS product to help you.

8.4.1 Implementation of virtual tape

Virtual tape support is included in the base i5/OS V5R4. A virtual tape device configures on the system as a stand-alone tape device, and works as a tape library device, which can automatically mount the virtual tape volumes. See also 8.5.1, “Create virtual tape devices” on page 175.

You can have up to 35 virtual tape devices. Virtual tape volumes support multiple optimum block sizes for compatibility with physical tape devices.

You utilize existing image catalog commands used for virtual optical to control virtual tape. The tape volume management interface is through iSeries Navigator. The user interface can be either i5/OS commands or iSeries Navigator.

BRMS also provides a complete end-user interface along with tape management support.

8.4.2 Virtual volumes stored as IFS stream files

The virtual tape volumes are stored as Type 2 Integrated File System (IFS) stream files. The stream files have a special attribute to prevent:

- ▶ Saving with storage free
- ▶ Scan processing
- ▶ Journaling

The stream files cannot be saved while the image catalog containing them is in a *ready* status. Manual editing of the stream files could cause data corruption.

When using File Transfer Protocol (FTP) to send to other servers or partitions, consider:

- ▶ Bandwidth
- ▶ Management of transmitted files
- ▶ Disaster recovery

The stream files can be stored in Dependent or Independent ASPs. When using an iASP, the iASP can, for example, be created on a remote set of DASD that are Fibre Channel (FC) attached. This type of setup can facilitate or improve a disaster recovery (DR) scenario. This is why we advise that you consider use of Independent IASPs that can be switched to or mirrored to a second system or partition.

8.4.3 Image catalog

Commands have been modified to support both previously available virtual optical and new virtual tape volumes. You can store up to 256 virtual tape volumes. Refer to the IBM Redbook *i5/OS V5R4 Virtual Tape: A Guide to Planning and Implementation*, SG24-7164, for information about managing the size of virtual tape volumes within the disk storage of the i5/OS partition.

*ALLOBJ authority is no longer required.

The loading of the image catalog on a virtual device makes the virtual volume accessible.

8.4.4 Virtual media

The virtual tape I/O manager is reloaded whenever the virtual tape I/O manager is varied on with a reset.

Virtual media can be write protected to avoid overwriting.

A virtual volume can be opened for input by multiple virtual tape devices. The volume status must be LOADED or MOUNTED before it can be used.

Auto-create of a new volume occurs when end of volume (EOV) is reached and either of the following occurs:

- ▶ The volume list has been exhausted and a new volume is specified in the reply to message CPA6798.
- ▶ The specified volume is MOUNTED and the end of the image catalog is reached.

The attributes of an auto created volumes include:

- ▶ Default values *except*:
 - Maximum size: 1 TB
 - Allocate storage (*MIN)
- ▶ Same as previous volume:
 - Density
 - Type (standard label, nonlabeled)
 - Encoding (EBCDIC, ASCII)

New volumes are placed in position 256 (IMGCLGIDX parameter). There is no auto create capability with BRMS.

8.4.5 Virtual tape devices with guest partitions (Linux, AIX)

This is supported using virtual I/O. The virtual tape device description must be configured with no volume unload during vary off. The virtual volume mounting requires that the i5/OS partition mounts the device in advance.

Here are some of the reasons why you may want the virtual tape image on another system or partition. On the system on which you are running a save or a restore operation, there is no hardware to support the media you use for backup. In this case you may decide to back up to virtual tape and then transfer the image to a partition or system that has a suitable physical tape device attached.

For a straightforward restore of the data in the virtual tape image to another partition or server without using tape media there are also different options.

Virtual tape images can be used in an iASP or in a user ASP. When used in a user ASP or (non-switchable) iASP the only benefit is that the user ASP or iASP is separated on another set of DASD than the system ASP. For switchable iASP the image can be used for backup/restore on multiple systems or partitions within the cluster domain, without requiring the image to be transferred to another system. You can simply switch the iASP.

There is also the option to use FTP for simply sending the image file to a remote server or partition.

Virtual tape fulfills a requirement for improved availability and recovery during system backup.

Virtual tape devices can be used for high-speed backup with a much lower risk of having the backup fail due to media or other errors. A permanent media error causes a backup to fail, which implies that there is no complete backup. You have to restart the backup or keep the incomplete backup for that day. You can avoid this situation using virtual tape support.

i5/OS Virtual Tape support is separate and distinct from *hardware virtual tape solutions* available in the marketplace. For example as of March 2006, i5/OS does not formally support the IBM Virtualization Engine TS7510 as a Virtual Tape Library (VTL). A VTL provides High Performance Backup and Recovery (HPBR) by utilizing disk arrays and virtualization software. In simple terms the TS7510 is an xSeries server running Linux attached to an IBM DS4000™ disk subsystem, running tape virtualization software.

Contact your IBM TotalStorage representative to investigate possible i5/OS support at a later date.

8.5 Implementing virtual tape

Now we are going to cover the actual implementation of virtual tape.

8.5.1 Create virtual tape devices

To create a virtual tape device, use the create device description (CRTDEVTAP) command.

```
                                Create Device Desc (Tape) (CRTDEVTAP)

Type choices, press Enter.

Device description . . . . . > NEWVRT01      Name
Device type . . . . . *RSCNAME             *RSCNAME, 2440, 3422, 3430..
Device model . . . . . *RSCNAME           *RSCNAME, 1, 2, 12, A01...
Resource name . . . . . > *VRT             Name, *NONE, *VRT...
```

Figure 8-2 Prompted command CRTDEVTAP DEVD(NEWVRT01) RSCNAME(*VRT)

When you specify *VRT for RSCNAME, the operating system will generate a virtual hardware resource name to use for the newly created device description. There is an important note here. Once the maximum number of virtual hardware resource names have been generated, any additional device descriptions that are created will be assigned the virtual hardware resource name that was last generated.

The resource name will not be removed if a device description is deleted. You may create a new device description for existing resources by specifying the resource by name.

For device type (TYPE) and device model (MODEL), we advise that you use the default *RSCNAME. You can specify the type and model, but that will be ignored. The virtual tape device description will always be created with device type 63B0 and device model 001.

You cannot directly create a virtual tape device from iSeries Navigator.

You can change an existing device description to be a virtual device.

8.5.2 Create the image catalog

To create a tape image catalog in iSeries Navigator, follow these steps:

1. In iSeries Navigator, expand **My Connections**.
2. Expand the iSeries server that you want to work with.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Tape Devices**.

6. Right-click **Tape Image Catalogs**. Select **Create Image Catalog**.

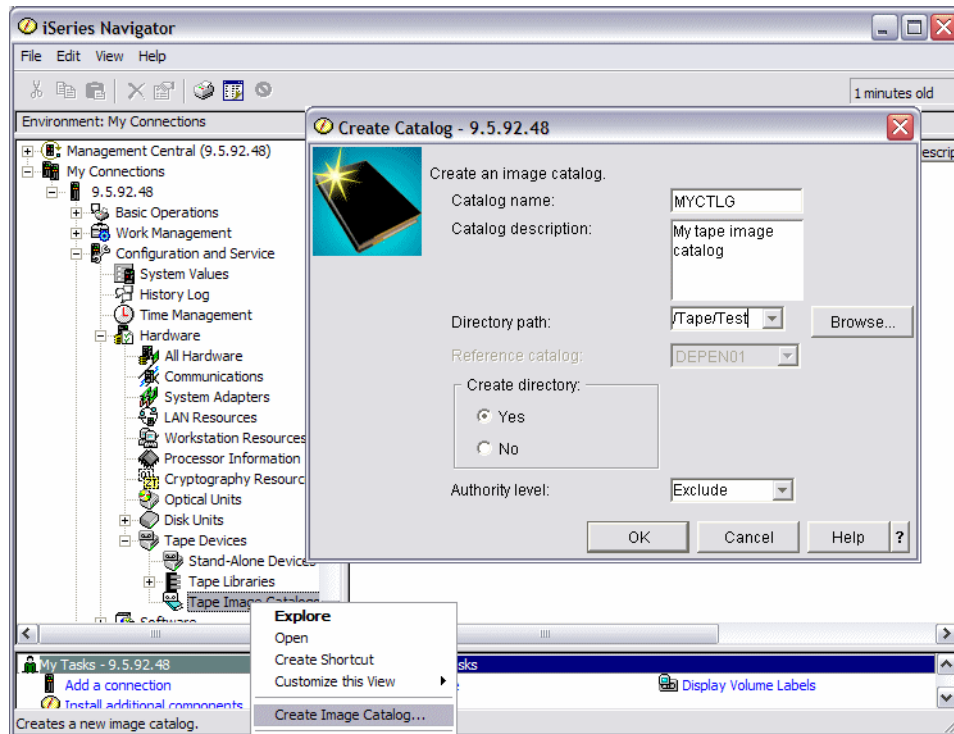


Figure 8-3 Create image catalog

The options are:

- ▶ Catalog name - Specify the name of the new tape image catalog.
- ▶ Catalog description - Type a description of the new tape image catalog (up to 50 characters).
- ▶ Directory path - Specify a directory to associate with the new tape image catalog. Use the Browse button to browse to a list of directories in the IFS file system.
- ▶ Reference catalog - Select the reference catalog from the list of image catalogs on your system. This is the catalog that you want to reference for the new catalog.
- ▶ Create directory - Specify whether a directory should be created if it does not already exist.
- ▶ Authority level - Specify the authority given to users.

To create a tape image catalog you can also use CRTIMGCLG CL command, as shown in Figure 8-4.

```

Create Image Catalog (CRTIMGCLG)

Type choices, press Enter.

Image catalog . . . . . > TAPECLG      Name
Directory . . . . . > '/TapeDir'

Image catalog type . . . . . > *TAP      *OPT, *TAP
Create directory . . . . . *YES        *YES, *NO

Additional Parameters

Text 'description' . . . . . *BLANK

Authority . . . . . *EXCLUDE      Name, *EXCLUDE, *LIBCRTAUT...

```

Figure 8-4 Prompted CRTIMGCLG IMGCLG(TAPECLG) DIR('/TapeDir') TYPE(*TAP)

8.5.3 Create a virtual tape image

In iSeries Navigator, expand **Configuration and Service** → **Hardware** → **Tape Devices** → **Tape Image Catalogs**.

Right-click the Image catalog that you want to add a volume to, and select **Add Volume**.

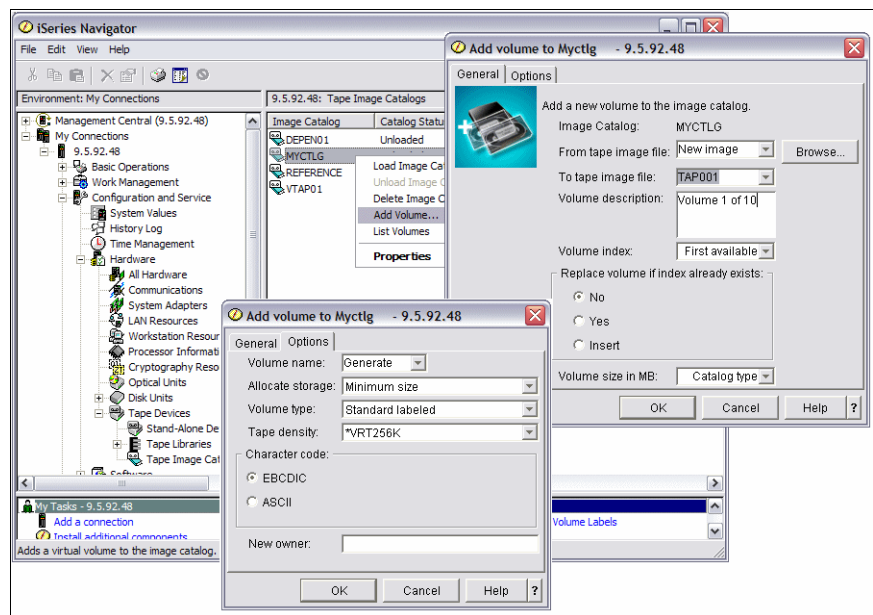


Figure 8-5 Add volume through iSeries Navigator

To tape image file (TOFILE)

This is the name of the file that will be copied to the target directory.

Volume size in MB

Specify the number of megabytes for the new image file. The range is 48 to 1,000,000 MB (one terabyte). The default is *catalog type*. For tape type image catalogs, the size of the image file created will be 1 GB.

Allocate storage parameter has options

When the default *minimum size* (*MIN) is taken, the newly created virtual tape volumes will initially use 4 KB of storage and will allocate additional storage as needed, until the image size is reached. Other important size considerations include:

- ▶ The maximum storage allocation for virtual tape is the greater of 95% or 5 GB free storage. That is, for ASPs smaller than 100 GB, the system can allocate up to 95% of the ASP size. For ASPs larger than 100 GB the system can allocate storage until there is only 5 GB remaining within the ASP.
- ▶ CPF4373 - End of Media on Device. This message is sent when the storage threshold of the ASP containing the virtual tape volume exceeds 95% of the up to 100 GB ASP or exceeds the 5 GB remaining of the larger ASPs, or when the maximum storage allowed threshold of the user profile that owns the virtual tape is exceeded and additional storage cannot be allocated to continue the output operation.
- ▶ When *volume size* (*IMGSIZ) is used the full amount of storage as specified by the image size parameter is allocated at create time.
- ▶ If you try to create a specific new virtual tape volume with the *IMGSIZ parameter and the volume size causes you to exceed the 95% capacity or 5 GB remaining rules for the ASP, you will get following error messages in your job log. With iSeries Navigator, you will get a popup error window that has the same error messages.

CPD4F11: Insufficient storage available in auxiliary storage pool #.

CPFBC28: Image catalog entry not added to image catalog CTLGNAME.

Volume name parameter

The tape volume name is equivalent to a cartridge ID label used by a tape library. It uniquely identifies the volume, even when it is non-labeled format, so that it can be auto-mounted when specified in a volume list. If you select Generate (Default), the volume name will be generated by the system.

Duplicate volume names are not allowed within the same image catalog.

The tape volume name can be changed via INZTAP.

Format of virtual tape images

The density (format) parameter controls the optimum block size that the volume will use so the created volume will be compatible for duplication to your physical tape device.

- ▶ Volumes with a density of *VRT256K will use an optimum block size of 256 KB and will be compatible with 35xx type devices and the newer QIC tape devices.
- ▶ Volumes with a density of *VRT240K will use an optimum block size of 240 KB and will be compatible with VXA and 8 mm technology devices along with the 35xx and newer QIC devices.
- ▶ Volumes with a density of *VRT64K will use an optimum block size of 64 KB and will be compatible with 3490F model 18 track media, VXA and 8 mm technology devices, and with the 35xx and newer QIC devices.
- ▶ Volumes with a density of *VRT32K will not use an optimum block size and will be compatible with all devices.

Volumes written using one of those formats can only be duplicated to tape devices that support the same maximum block size or greater.

You want to pick the largest compatible optimum block size to maximize performance. The Initialize Tape (INZTAP) command can be used to change the density of an existing volume.

Initialize tape (INZTAP) removes any existing data on the tape.

You can also use ADDIMGCLGE CL command to add the virtual tape volume from an existing virtual tape file or create a new virtual tape volume.

8.5.4 Working with virtual devices

As seen in Figure 8-6, to vary on a virtual tape device volume into a tape image catalog using a stand-alone tape device, follow these steps:

1. In iSeries Navigator, expand **My Connections**.
2. Expand the iSeries server that you want to work with.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Tape Devices**.
6. Expand **Stand-Alone Devices**.
7. Right-click the stand-alone device.

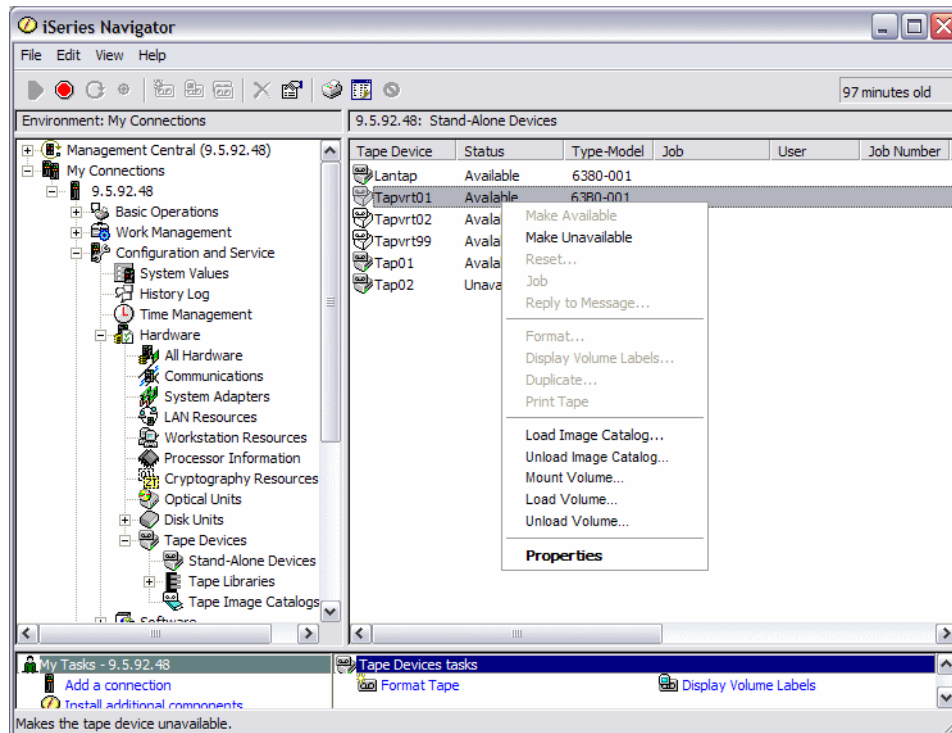


Figure 8-6 Working with virtual devices

You vary on by selecting **Make Available**. Also, you can load, unload, and mount the volume.

You can also use the WRKCFGSTS CL command to list the virtual tape devices, as in this example:

```
WRKCFGSTS CFGTYPE(*DEV) CFGD(*TAP).
```

You can vary on and off the virtual tape devices just like actual tape devices.

8.5.5 Working with image catalogs

To display the properties of the virtual volume, follow these steps:

1. In iSeries Navigator, expand **My Connections**.
2. Expand the iSeries server that you want to work with.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Tape Devices**.
6. Expand **Tape Image Catalogs**.
7. Right-click the tape image catalog where the virtual volume is located. Select **List Volumes**.
8. Right-click the virtual volume that you want to display properties for, and select **Properties**.

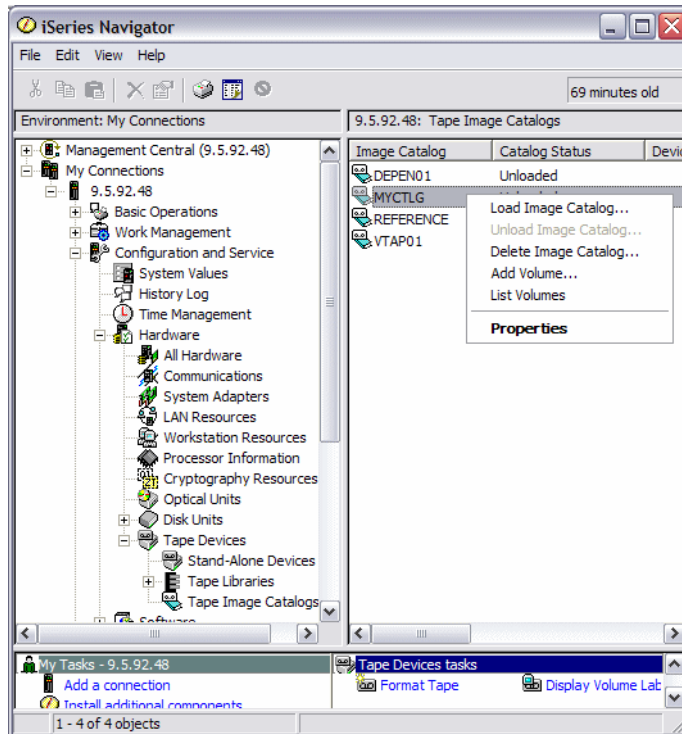


Figure 8-7 List image catalog

You can also use `WRKIMGCLGE CL` command if you know your image catalog name, or you can use the `WRKIMGCLG CL` command.

8.5.6 Working with virtual volumes

In iSeries Navigator, with List Volume window, you can also load, unload, format, and duplicate a tape volume.

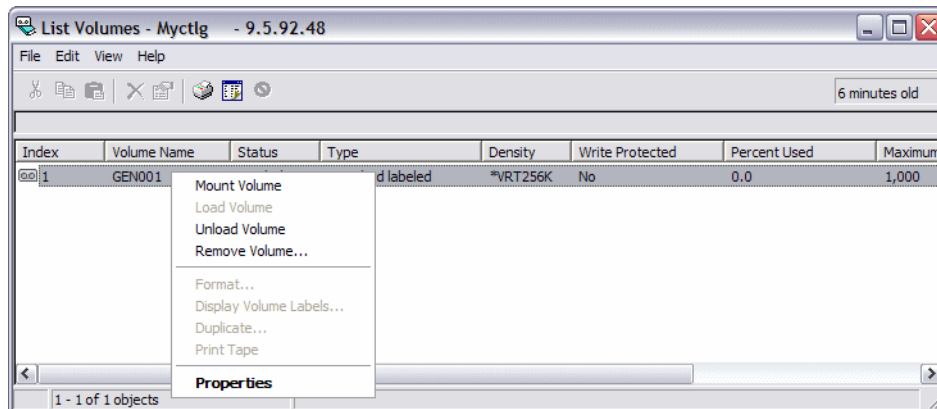


Figure 8-8 List volume window

In the list volume window, you can find out the maximum size allowed for the virtual volume associated with the image catalog entry and the percentage of the virtual tape volume that has been used.

You can also use the CL command to list volumes in your image catalog.

Use Option 12 = Work with entries in the WRKIMGCLG CL command or use the WRKIMGCLGE CL command if you remember the image catalog name.

If you choose F11 you can also find out the maximum size and the percentage used of the volume.

Use the DUPTAP CL command to duplicate a catalog loaded in a virtual tape device.

8.6 Dependent catalog overview

A dependent image catalog is a copy of another (reference) image catalog. It contains the same virtual volumes as the reference image catalog, but it can be mounted in a different virtual device. This enables simultaneous reads from the same virtual volume.

See Figure 8-9 for an illustration of using the reference and dependent catalogs.

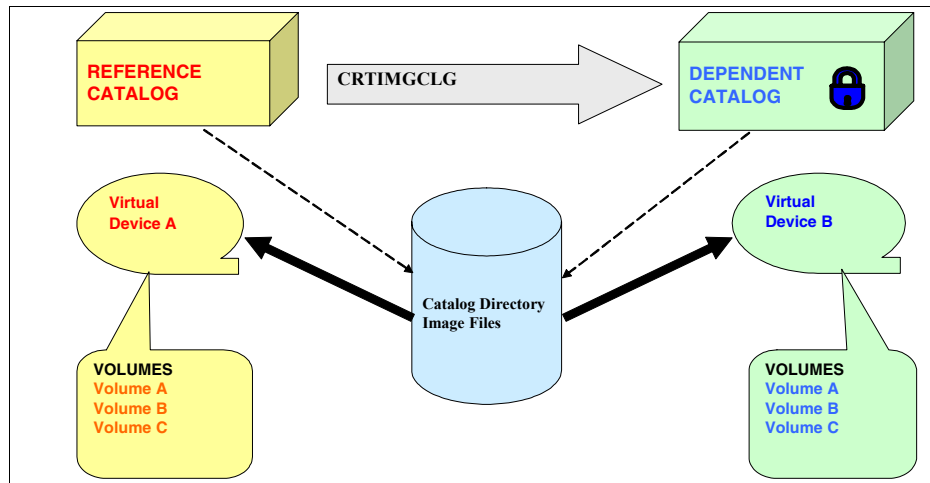


Figure 8-9 Dependent catalog overview

The following sections describe the benefits of using a dependent catalog and its reference catalog, the benefits of using them and how to set them up.

8.6.1 Using dependent catalogs

The following are benefits of using a dependent catalog:

- ▶ It allows more than one operation to be performed on a single volume. This includes displaying, restoring, installing, and duplicating.
- ▶ It helps to protect against accidental deletion of data, as all of the volumes in a dependent catalog are write protected.
- ▶ It helps to reduce the amount of system disk space need to perform some concurrent operations. Multiple true bit-for-bit copies are not needed in order to perform more than one operation on a single volume.
- ▶ It can be used to provide an unaltered view of the original image catalog. A dependent catalog is created as a copy of the reference catalog at a single point in time.

8.6.2 Reference catalog

An image catalog that shares its resources (catalog information, directory, and image files) with other image catalogs is called a reference catalog.

Reference catalog characteristics include:

- ▶ Object type *IMGCLG.
- ▶ A catalog becomes a reference catalog once a dependent catalog has been created.
- ▶ All dependent catalogs must be deleted before the reference image catalog files can be deleted.
- ▶ The reference catalog cannot be deleted unless all dependent catalogs have been deleted first.
- ▶ Optical type reference catalog image files can only be loaded in a virtual device as read only.

- ▶ Tape type reference catalog image files can be loaded in a virtual device as writeable or read-only.
- ▶ The reference catalog contains the actual image files.

A shared catalog is an image catalog that shares catalog information, the image catalog directory, and the associated images with another catalog. The catalog being shared with other catalogs is called a reference catalog, and an image catalog that uses a reference catalog's resources is called a dependent catalog.

A reference catalog is an image catalog that shares its resources (catalog information, directory, and image files) with other image catalogs. The shared catalog becomes the reference catalog once a dependent catalog has been created using the create image catalog command. The reference catalog cannot be deleted unless all dependent catalogs have been deleted first.

8.6.3 Dependent catalog

This is an image catalog that is dependent on the resources of a reference catalog. Dependent catalog characteristics are:

- ▶ Object type *IMGCLG.
- ▶ References another nondependent catalogs resources.
- ▶ Loaded as read only (write protect on).
- ▶ Does not set the allow save attribute for IFS files.
- ▶ Cannot delete the image catalog image files via the RMVIMGCLGE or WRKIMGCLGE commands.
- ▶ DLTIMGCLG will delete a dependent catalog as long as the *KEEP parameter is *YES.
- ▶ Snapshot of its reference.
- ▶ New reference catalog volumes are not seen by a dependent catalog.
- ▶ Optical catalog implement volume prefixing.
- ▶ LIC, OS, or PTF install not supported.
- ▶ LP install supported.

A *dependent catalog* is an image catalog that is dependent on resources of a reference catalog. A dependent catalog contains a snapshot of the reference catalog during the point in time when the dependent catalog was created. A dependent catalog is loaded as read only (write protect on) and can only be loaded as read-only media. You cannot delete virtual volumes in an image catalog using the *delete image catalog* command.

There can be up to 35 ready dependent catalogs pointing to one reference catalog because the virtual tape device can be 35. Each of the tape reference catalogs are read only and write accessible. However, optical reference and all dependent image catalogs are read only and can be used for restore operations.

In order to delete the reference catalog or to delete any one image file, all of the dependent image catalogs must be deleted first. The volume name for the optical dependent catalogs has a four-character prefix. The prefix is added once the optical dependent catalog is made ready. Use the *work with image catalog entries* (WRKIMGCLGE) to find the optical dependent volume name.

8.6.4 Creating a dependent catalog

To create a dependent catalog you can use the Create Image Catalog function in iSeries Navigator. In the directory path parameter, you should have to select a reference other than specifying the actual directory path. Then you can select the reference catalog from the list of image catalogs on your system in the reference catalog parameter.

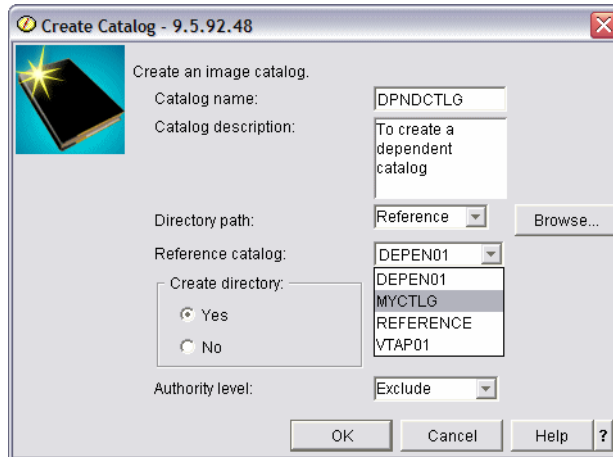


Figure 8-10 Creating a dependent catalog

You can also create a dependent catalog with the CRTIMGCLG CL command:

```
CRTIMGCLG IMGCLG(DEPENDENT1) DIR(*refimgclg) TYPE(*TAP) REFIMGCLG(MYCATALOG)
```


8.6.5 Shadow catalog views

There are Reference and Relationship columns, which show you the relationship of reference catalogs and dependent catalogs in the tape image catalog view using iSeries Navigator (see Figure 8-11).

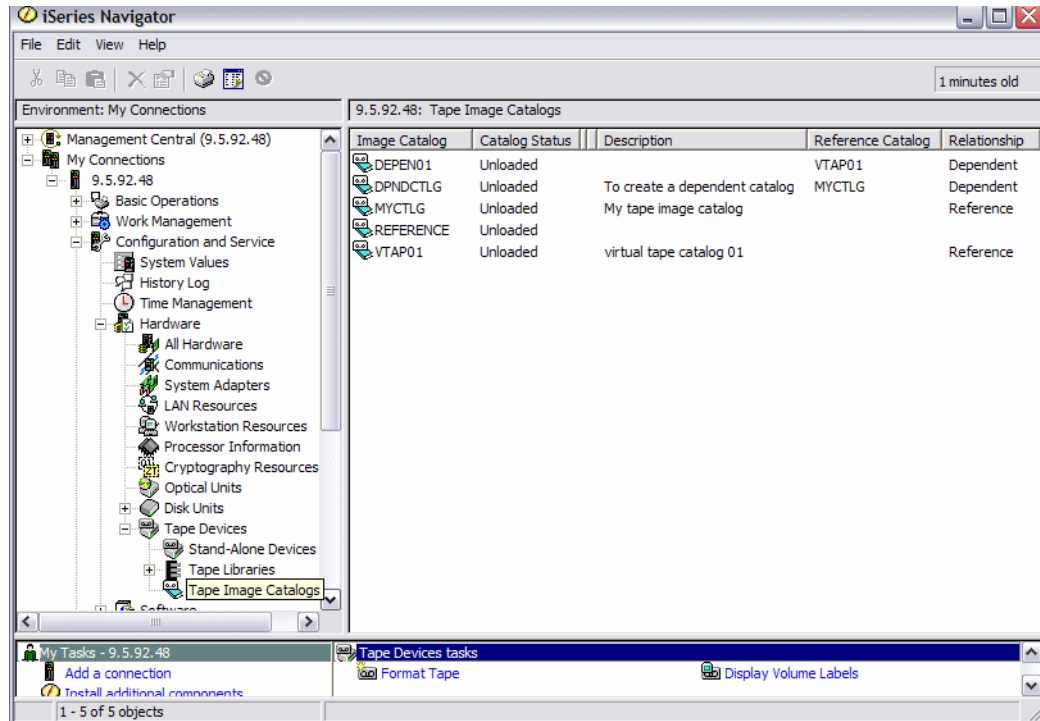


Figure 8-11 Shadow catalog view

In the green screen, at the Work with Image Catalogs display (WRKIMGCLG), you can also see the catalog relationship with F11(=View relationships).

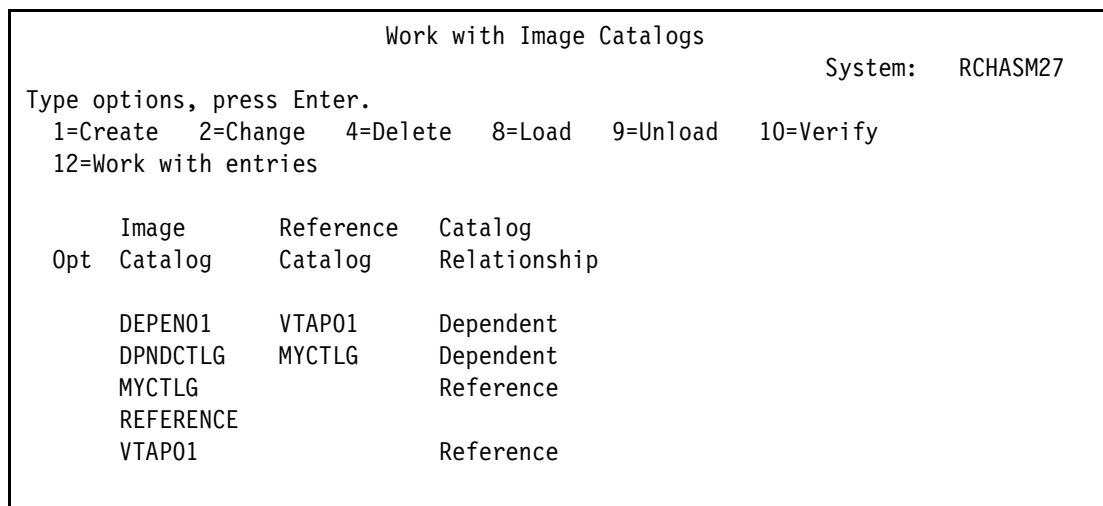


Figure 8-12 F11 view in WRKIMGCLG

8.6.6 Catalog security changes

Until V5R3, you must have security administrator (*SECADM) and all object (*ALLOBJ) special authorities to use image catalogs.

Customers have been reluctant to grant *ALLOBJ/*SECADM special authorities to image catalog users. With the addition of new save restore capabilities for optical and tape type catalogs, there is a need to reduce the security requirements.

In V5R4 it is no longer required to have *ALLOBJ and *SECADM authority to use image catalogs. An operator can use the image catalog function without special authorities.

Common catalog security requirements at V5R4:

- ▶ To use an existing catalog:
 - *USE authority to any desired image catalog commands.
 - Execute (*EXECUTE) authority to library QUSRSYS.
 - *USE authority to the image catalog.
 - Execute (*X) authority to each directory in the image catalog path name.
 - *USE authority to the virtual device description.
- ▶ To administer or create image catalogs:
 - *USE authority to any desired commands.
 - Execute (*EXECUTE) authority to library QUSRSYS.
 - *CHANGE authority to the image catalog.
 - Execute (*X) authority to each directory in the image catalog path name.
 - *USE authority to the virtual device description.
 - If adding an image file from physical media, *USE authority to the physical device is required.
 - If adding an image file from an existing image located in a different directory, the authorities required are the same as those required by the Copy Object (CPY) command.

8.6.7 New image catalog APIs

New image catalog APIs are:

- ▶ Retrieve image catalogs (QVOIRCLG)
Retrieves the list of image catalog names based on the image catalog type parameter.
- ▶ Retrieve image catalog details (QVOIRCLD)
Retrieves the contents of an image catalog. The QVOIRCLD API places the information in the receiver variable.

For more information search the iSeries Information Center with *Catalog API*.

8.7 V5R4 BRMS virtual tape support

We cover BRMS support in Chapter 7, “i5/OS-based backup and recovery” on page 147, and touch on virtual tape as a device and backup media in “Native virtual tape support” on page 162. Some highlights of BRMS support include:

- ▶ Virtual tape acts as both a device and a media for all save/restore, copy, tape utility functions, virtual I/O, and high-level language interfaces.
- ▶ Parallel save supported is supported.
- ▶ Once the backup is complete, can resume operations and do DUPMEDBRM in background.
- ▶ Can keep virtual tape for restoring files needed to recover during the week, send real tapes out for off-site storage.

Caution: Virtual tape can consume large amounts of disk space, as you may create and use many virtual tape catalog entries and could use dependent catalogs without realizing that you need to manage the associated disk space.

You should have a formal plan to manage this storage, such as removing old images no longer needed, to run your business.

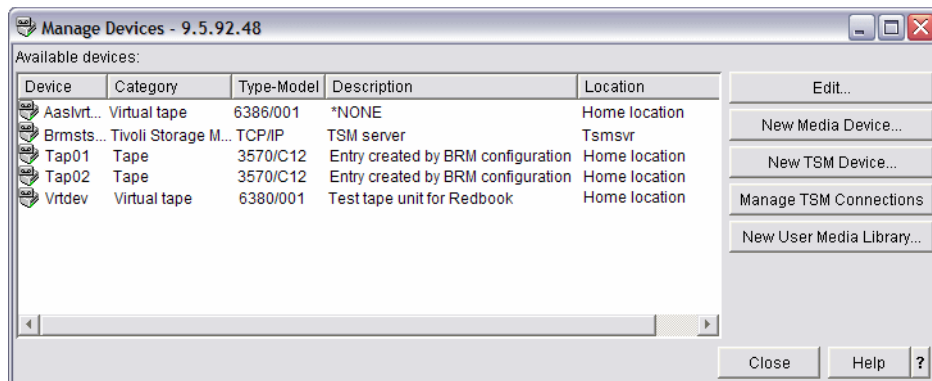


Figure 8-13 BRMS Manage Devices with virtual tape devices defined



Accessing 5/OS via the Web

This chapter discusses the enhancements in the iSeries Navigator Tasks on the Web and iSeries Access for Web in V5R4. We also show configuration and customization of some functions and features.

9.1 iSeries Navigator tasks on the Web

iSeries Navigator tasks on the Web, also referred to as iSeries Navigator on the Web, takes a major set of iSeries Navigator for Windows functions and provides browser-based access to them. This means that there is no application maintenance required on the client workstation, and you still have a lot of interfaces to i5/OS-based functions.

Originally introduced with V5R3, there is a major increase in iSeries Navigator for Windows functions ported to V5R4 iSeries Navigator tasks on the Web. This includes customizing the browser user's view and access to functions.

Basically, iSeries Navigator tasks on the Web or the iSeries Navigator URL Advisor, is a java servlet application that resides on an iSeries system and is accessed with a Web browser. A user can now access almost all iSeries Navigator data from any Web browser on any system or Internet-capable device instead of being limited to just the PC with a client application installed.

In general, V5R4 iSeries Navigator tasks on the Web functions are management (administration) functions, whereas iSeries Access for the Web, another Java servlet-based implementation, functions are oriented for the end users.

Additionally, you can access data from multiple iSeries systems in the network through a single iSeries Navigator Web server system. The browser interfaces with the Web server system to display data from multiple systems. Only the Web server system needs to be at V5R4. Other systems can be at different release levels (N-2, V5R3 or V5R2).

9.1.1 Product packaging and software requirements

This product is part of OPTION 3 of the operating system 5722-SS1. The setup and activation of this function in V5R4 requires that the following products are installed in the system:

- ▶ 5722-SS1 *BASE V5R4
- ▶ 5722-SS1 option 3
- ▶ 5722DG1: IBM HTTP Server for iSeries
- ▶ 5722SS1: (option 0012): OS/400 Host Servers
- ▶ 5722SS1: (option 0030): OS/400 Qshell
- ▶ 5722JV1: IBM Developer Kit for Java
- ▶ 5722JV1: (option 6): Java Developer Kit 1.4
- ▶ One of the following
 - 5733W60: WebSphere Application Server V6 (Base)
 - 5733W60: WebSphere Application Server V6 Express

The following products are recommended:

- ▶ 5722-SS1 option 34: Digital Certificate Manager (required for SSL)
- ▶ 5722-ACx: Cryptographic Access Provider (required for SSL)
- ▶ AFP™ Browser Plug-in (allows viewing spool files from a browser)

9.1.2 Activation and setup

The great thing about this product is that it is packaged as part of the operating system 5722-SS1 Option 3. It does, however, need to be set up via the HTTP Admin GUI. No installation is required.

To set up the iSeries Navigator tasks on the Web:

1. Ensure that IBM HTTP *admin* instance is running. Check that under the subsystem QHTTSPVR there are three ADMIN jobs running by using the command:

```
WRKACTJOB SBS(QHTTSPVR)
```

This command will give you the screen shown in Figure 9-1.

```
Work with Active Jobs                                RLTSYS
                                                    06/19/06 10:01:41
CPU %:      .0      Elapsed time:  00:00:00      Active jobs:  268

Type options, press Enter.
  2=Change  3=Hold  4=End  5=Work with  6=Release  7=Display message
  8=Work with spooled files  13=Disconnect ...

Current
Opt  Subsystem/Job  User      Type  CPU %  Function      Status
   QHTTSPVR      QSYS      SBS    .0     PGM-QZHBMAIN  DEQW
   ADMIN         QTMHHTTP  BCH    .0     PGM-QZSRLOG   SIGW
   ADMIN         QTMHHTTP  BCI    .0     PGM-QZSRLOG   SIGW
   ADMIN         QTMHHTTP  BCI    .0     PGM-QZSRHTTP  SIGW

                                                    Bottom

Parameters or command
====>
F3=Exit   F5=Refresh   F7=Find   F10=Restart statistics
F11=Display elapsed data  F12=Cancel  F23=More options  F24=More keys
```

Figure 9-1 ADMIN jobs in QHTTSPVR subsystem

2. Point your browser to `http://<web-server-system>:2001` and sign on as required, as shown in Figure 9-2.

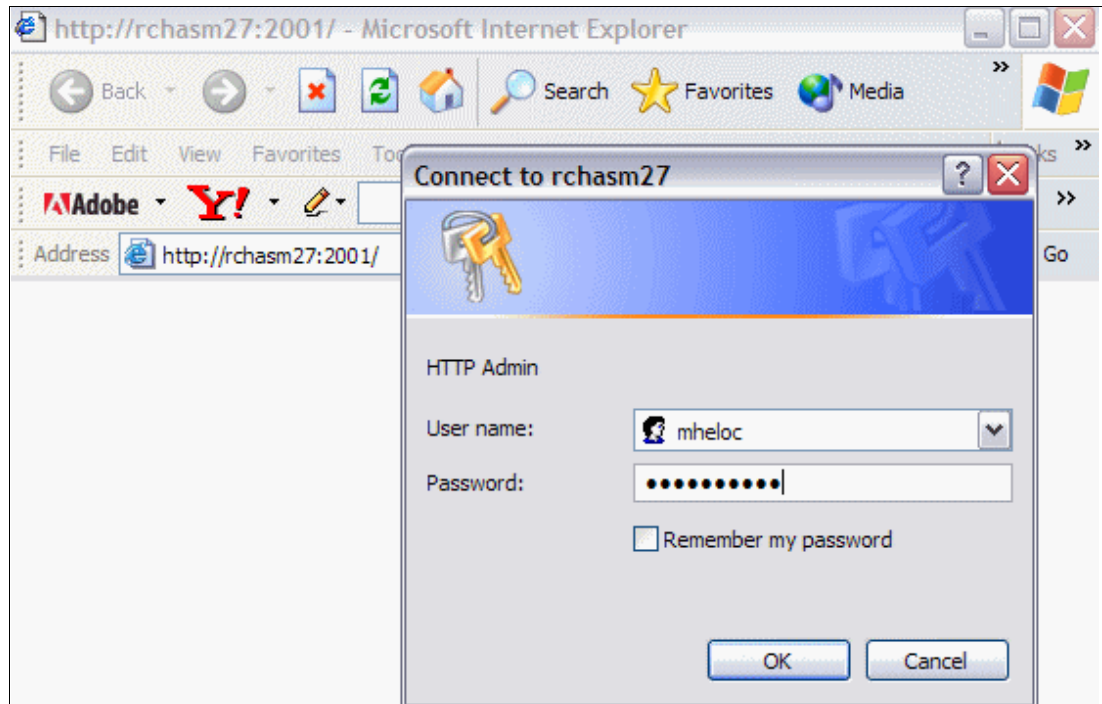


Figure 9-2 IBM HTTP Administration sign-on

3. Select the **IBM Web Administration for i5/OS** link in the i5/OS Tasks page (Figure 9-3).

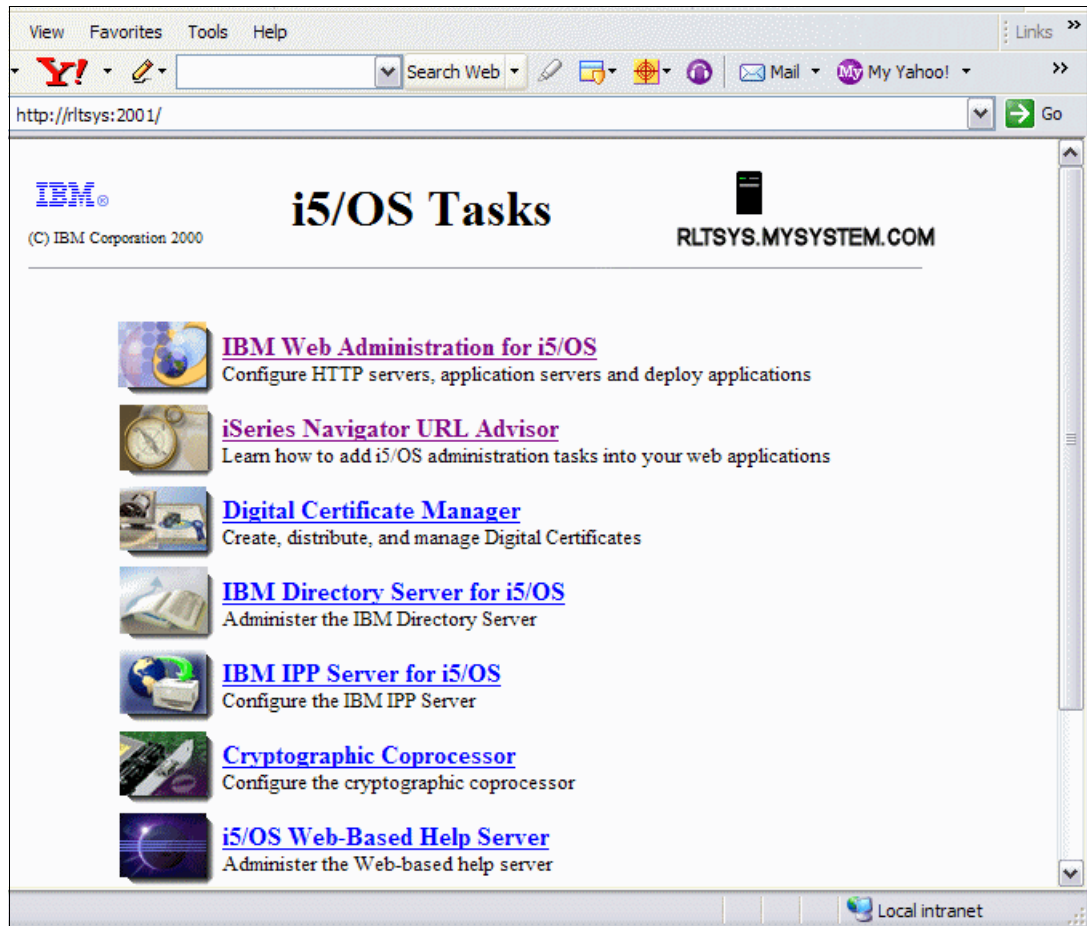


Figure 9-3 i5/OS tasks page

4. From the top set of tabs, select **Manage**, as shown in Figure 9-4.

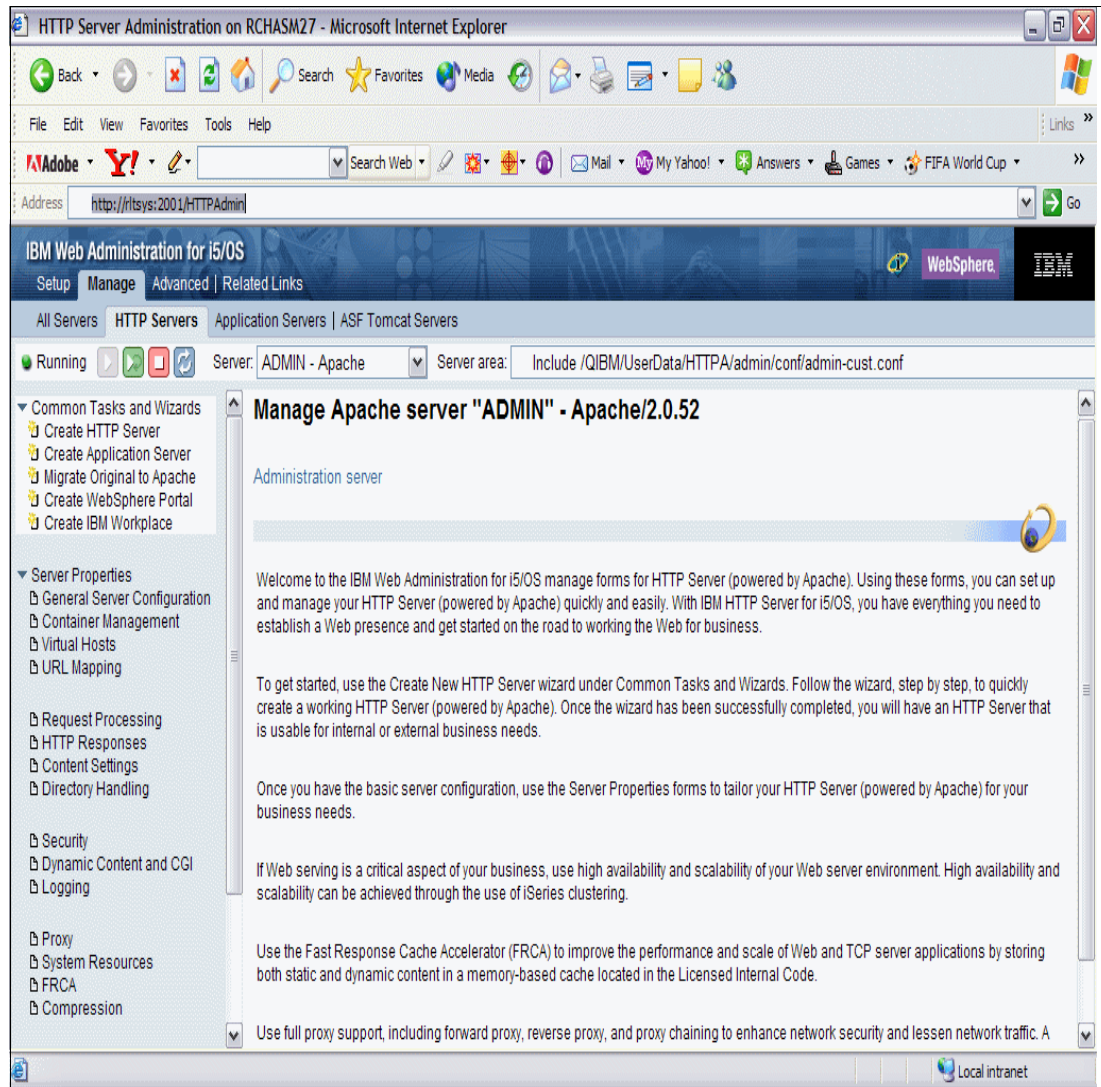


Figure 9-4 Managing the Apache server ADMIN

5. From the set of tabs that appears below those, select **HTTP Servers**.

6. From the Server drop down, select **ADMIN – Apache**.

- From the left-hand list of tasks select **General Server Configuration** under Server Properties, which gives you the screen shown in Figure 9-5.

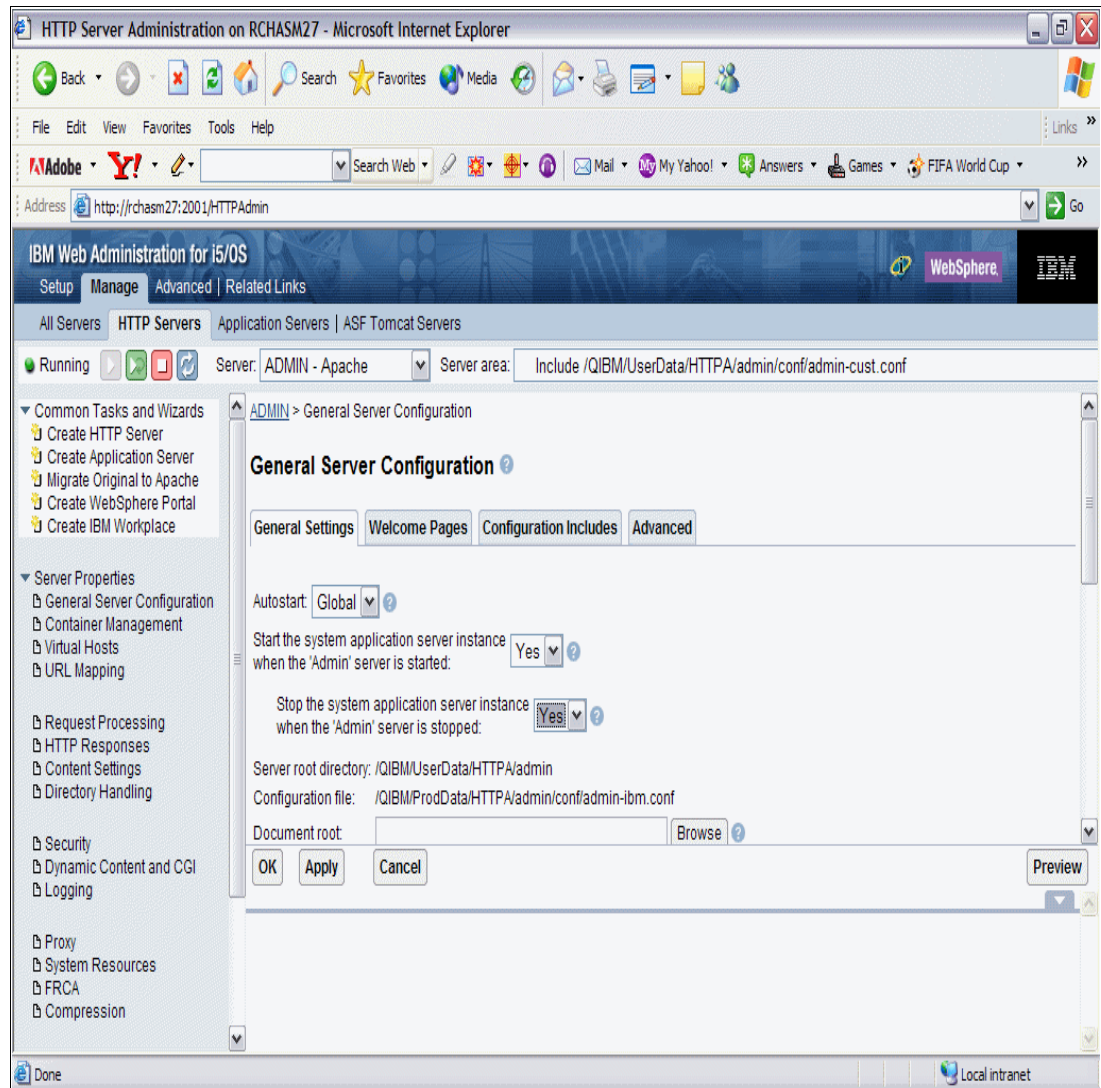


Figure 9-5 General server configuration of the APACHE admin

- On the General Settings tab, specify **Yes** for both “Start the system application server instance when the ‘Admin’ server is started” and “Stop the system application server instance when the ‘Admin’ server is stopped.”
- Click **OK**.

iiSeries Navigator Tasks for the Web runs in the WebSphere Application Server for i5/OS V6 provided server/profile named SYSINST/ADMIN. The steps we just completed ensure that this SYSINST/ADMIN is started whenever the HTTP admin server is started.

Now you must restart the HTTP admin server.

- ▶ To stop the server, use the command:
ENDTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)
- ▶ To start the server, use the command:
STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)

Note: During October 2006 IBM made available at no charge a set of Web enablement enhancements. These enhancements include the capability of iSeries Navigator (tasks) on the Web to operate under an integrated Web application server instead of SYSINST. For more information about this see Chapter 14, “i5/OS-based Web enablement enhancements” on page 337.

9.1.3 i5/OS V5R4 iSeries Navigator Tasks on the Web - basic functions

iSeries Navigator Tasks on the Web is more of an administration tool designed to manage your System i servers using a browser.

The very basic functions on the iSeries Navigator Tasks on the Web are discussed in this section.

For an in-depth discussion of these capabilities, refer to the iSeries Navigator Tasks on the Web topic in the iSeries Information Center V5R4 found at:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r4/topic/rzatg/rzatgoverview.htm>

URL wizard

The URL wizard guides users through creating a customized URL to an iSeries Navigator task. You can start the wizard by clicking the **Create a URL** button on the iSeries Navigator Tasks on the Web home page, as shown in Figure 9-6.

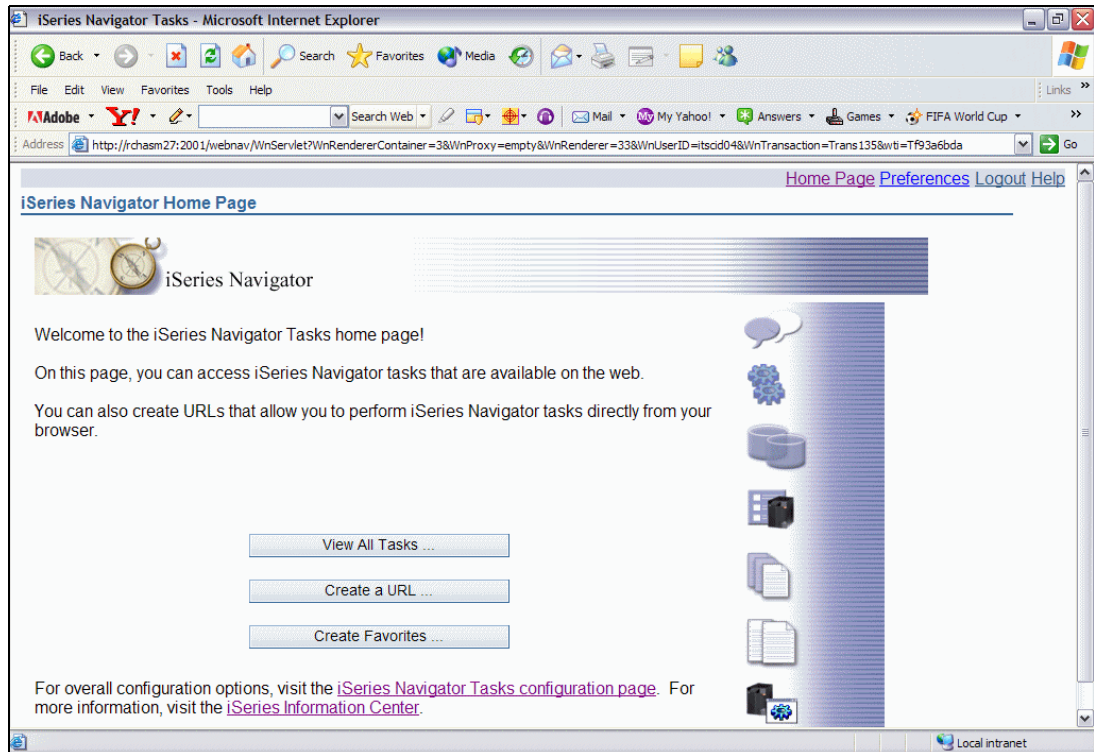


Figure 9-6 iSeries Navigator home page

This function will have you choose what type of task you want to run, as shown in Figure 9-7.

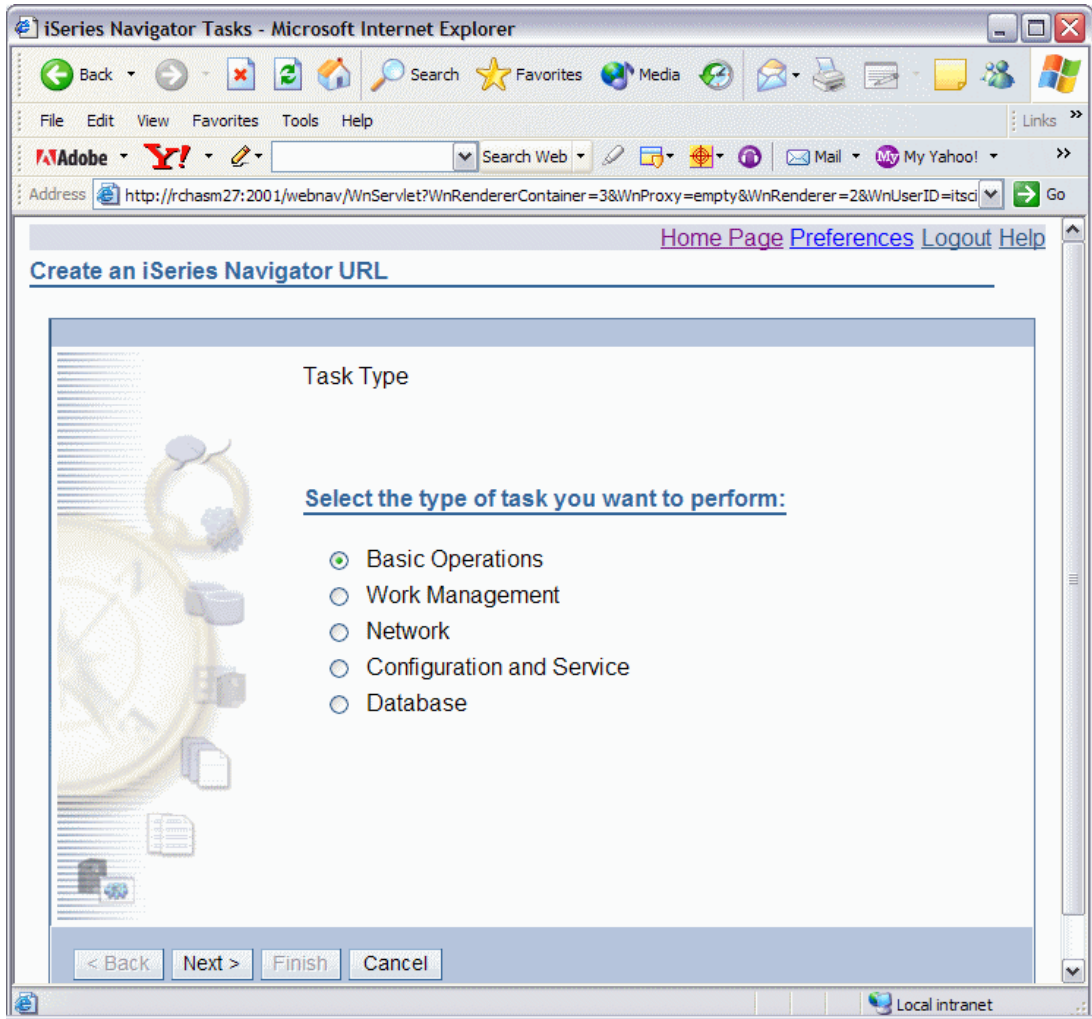


Figure 9-7 URL wizard: type of task selection

For our example we select the basic operations type of task (Figure 9-7 on page 197). We click **Next** to continue with the procedure. This will open another browser window where a list of tasks related to basic operations are listed, as shown in Figure 9-8.

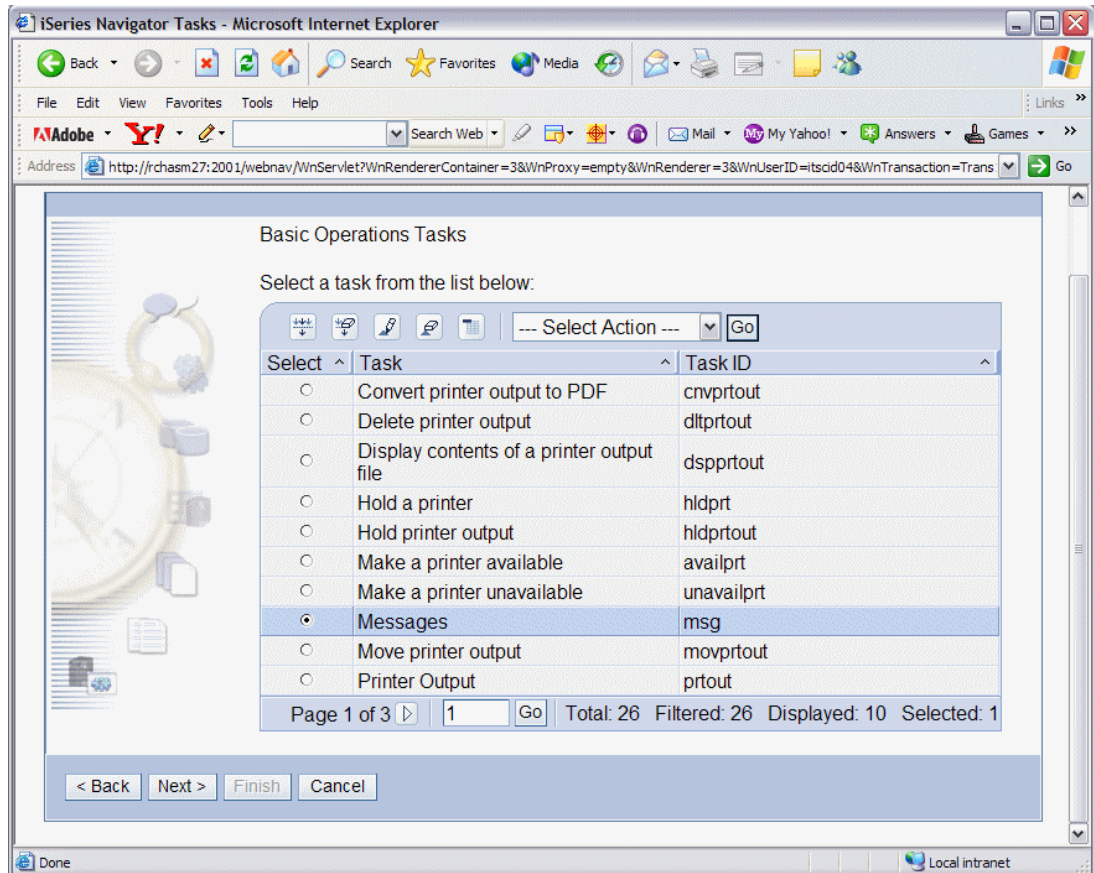


Figure 9-8 URL wizard: basic operations tasks list

After selecting your needed task, you are to choose on which machine or server you would like to run the task you wish to perform. You can also optionally choose a user ID (if a different user ID other than the current user ID will be used) to be used on the target system. A screen capture of this example is found in Figure 9-9.

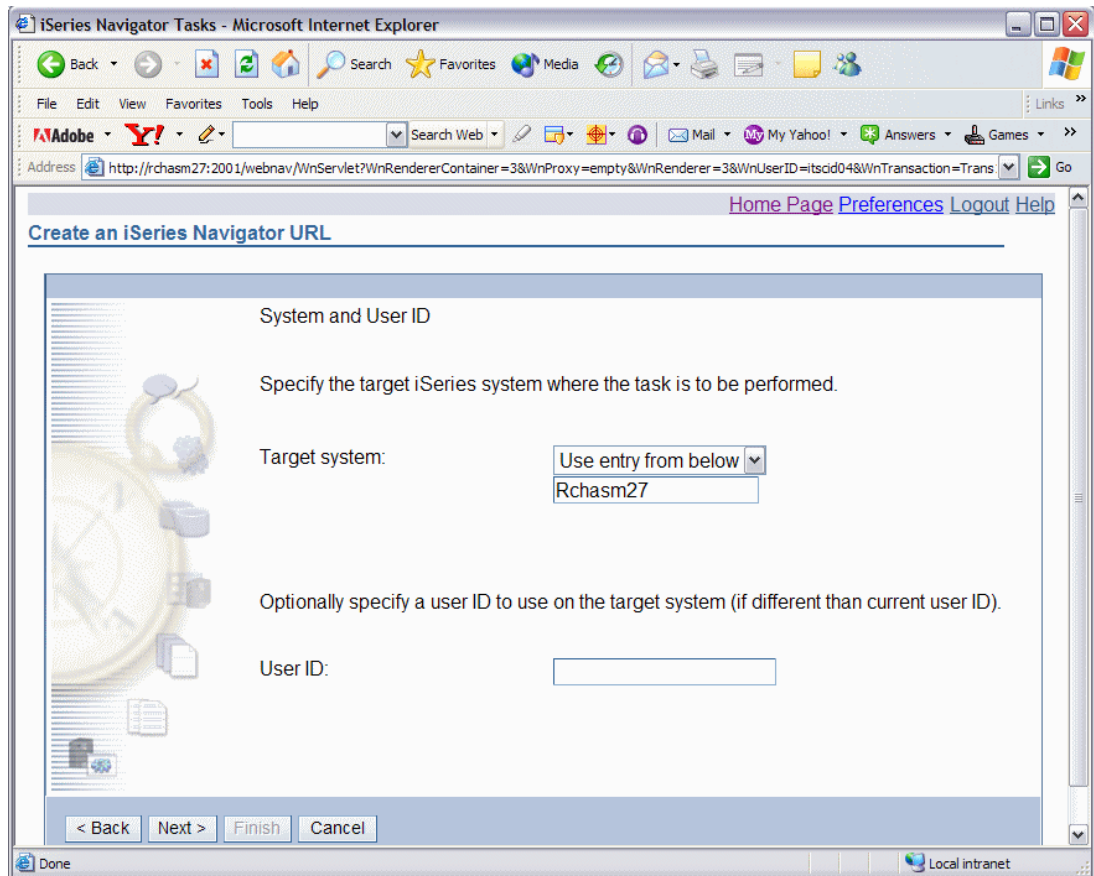


Figure 9-9 URL wizard: target system

After clicking **Next**, the browser will show the Task Customization window, as shown in Figure 9-10. This is where you can tailor how the final report of the task you have selected will look when the generated URL is opened.

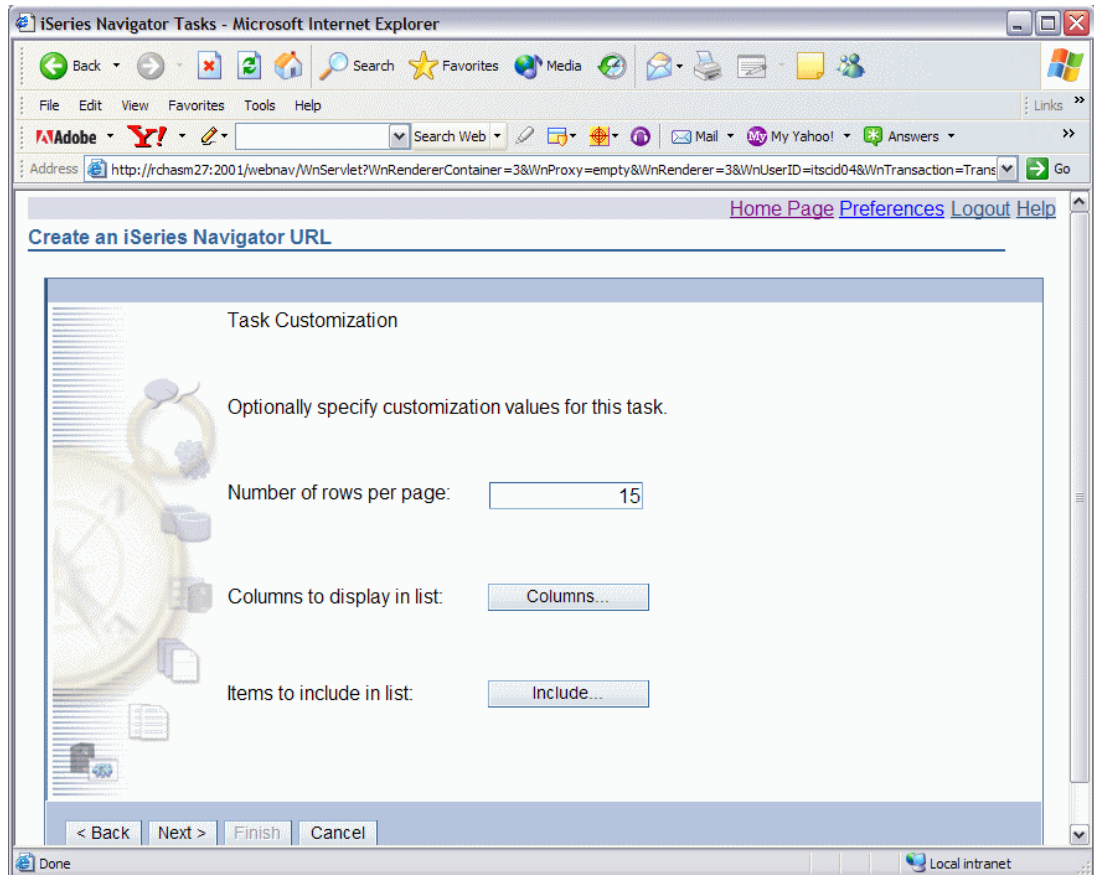


Figure 9-10 URL wizard: Task Customization window

In the URL wizard, the columns and include data specified is dynamic, meaning that it will be appended to the URL generated and will only be used for invocations of that URL. That Columns and Include data will also override any persistent Columns and Include data previously specified for the list. When Columns and Include data is specified on the URL, that data is used to determine the columns and list data shown.

This dynamic Columns and Include data allows the user to customize the list without making it their default for the list. Refer to Figure 9-11 and Figure 9-12 for a screen shot of how these two customization parameters are set in the wizard.

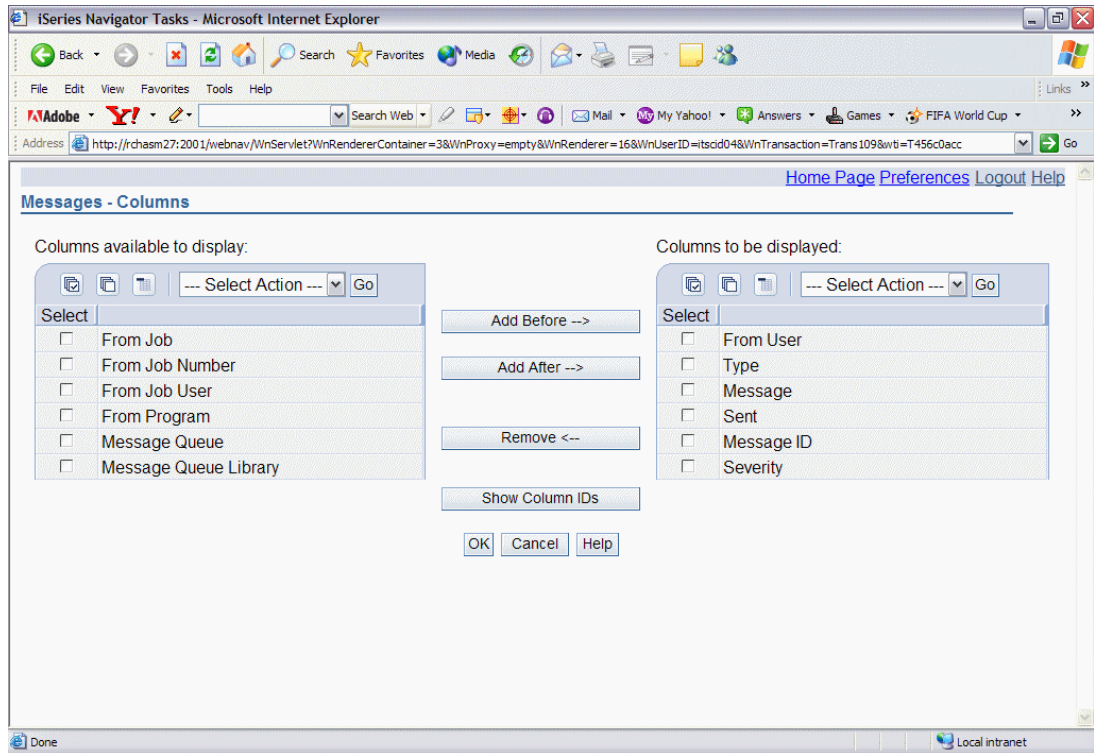


Figure 9-11 URL wizard: Columns

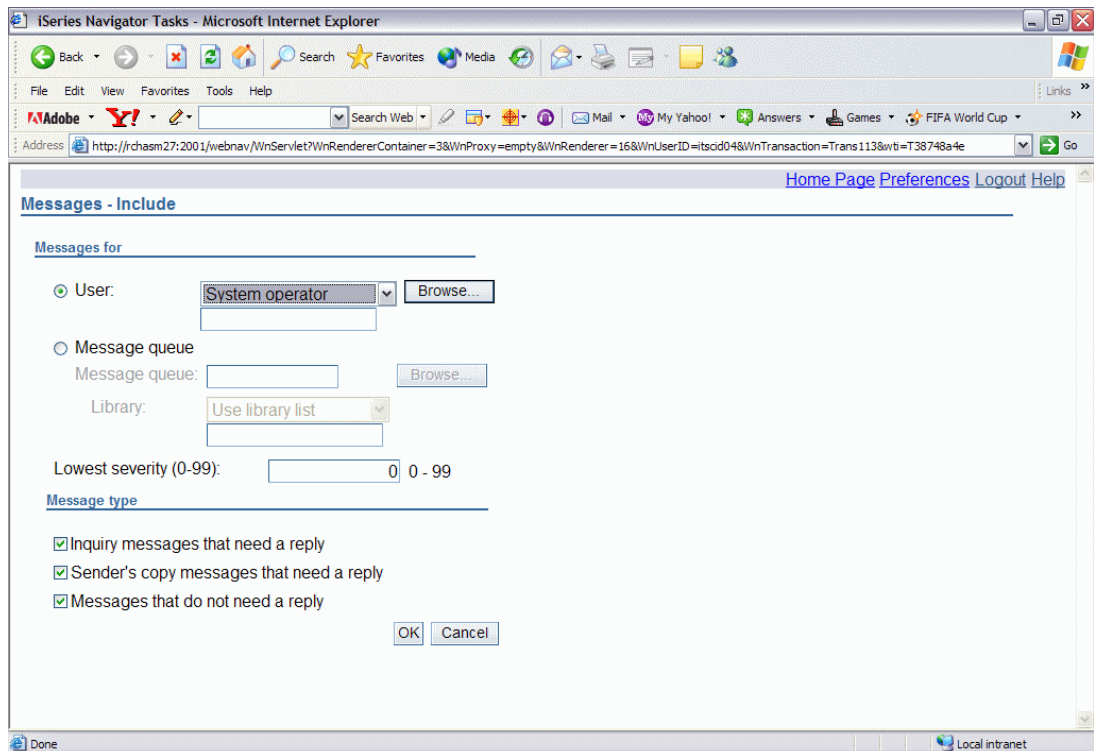


Figure 9-12 URL wizard - Include

After specifying the parameters in the Task Customization window and clicking the **Next** button, the browser opens a URL Complete window showing the URL of the task you just selected and customized, as shown in Figure 9-13.

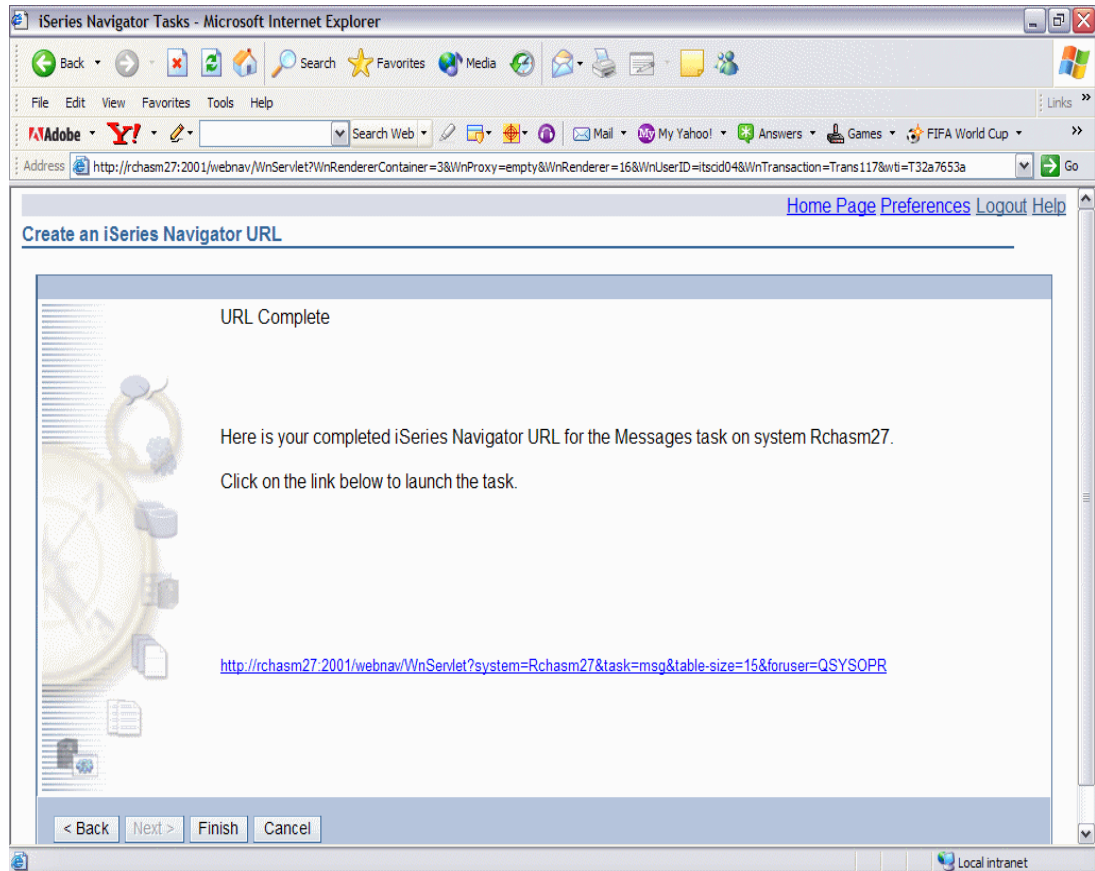


Figure 9-13 URL wizard - URL complete

The generated URL is displayed and can be invoked from this page. Once invoked, you can save it to your favorites.

The URL contains parameters with data from the wizard:

- ▶ Task to be invoked (task)
- ▶ System to run task for (system)
- ▶ User ID to use for the task (user)
- ▶ Any customizations (table size, columns, any include parameters)

Lists on the Web

In the iSeries Navigator Tasks on the Web home page, as shown in Figure 9-6 on page 196, you can click the **View All Tasks** button to see the type of tasks and the tasks under them. The tasks for each type are categorized into Lists and Actions, as shown in Figure 9-14.

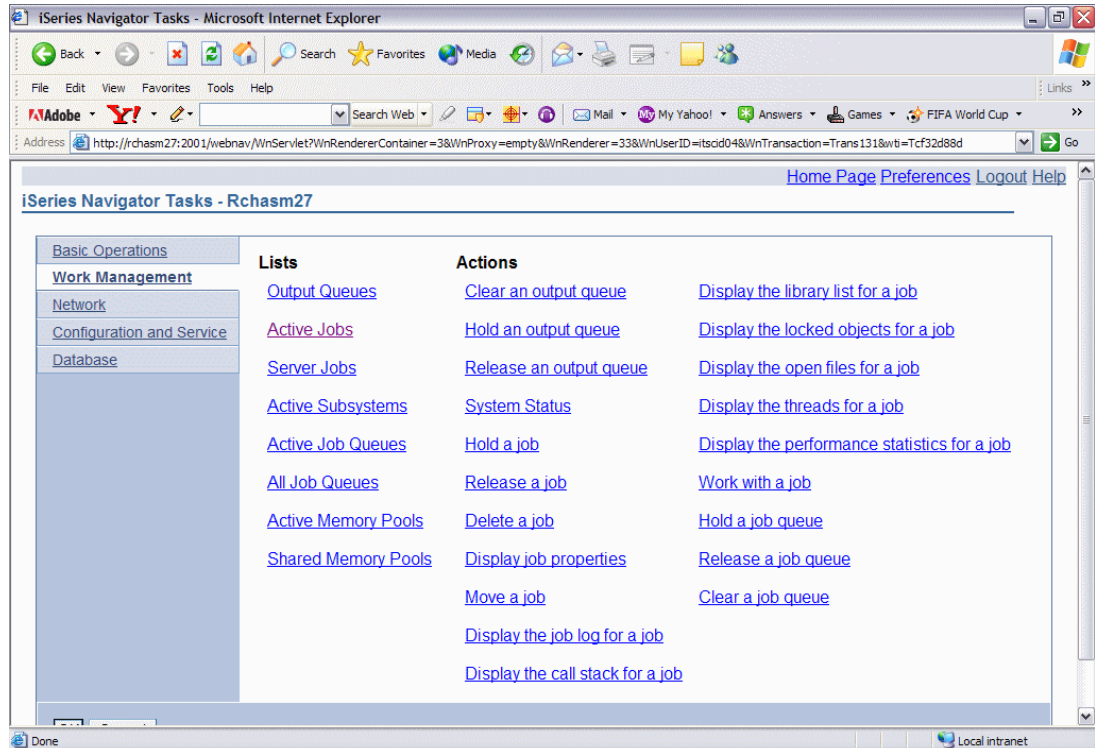


Figure 9-14 iSeries Navigator: view all tasks

Lists are rendered in a manner consistent with how a Web user would expect to see them. The functionality is the same. Only the look is different.

The drop-down menu provides actions that affect the entire list (that is, columns, include) as well as actions specific to the currently selected items (that is, printer output, job log), as shown in Figure 9-15.

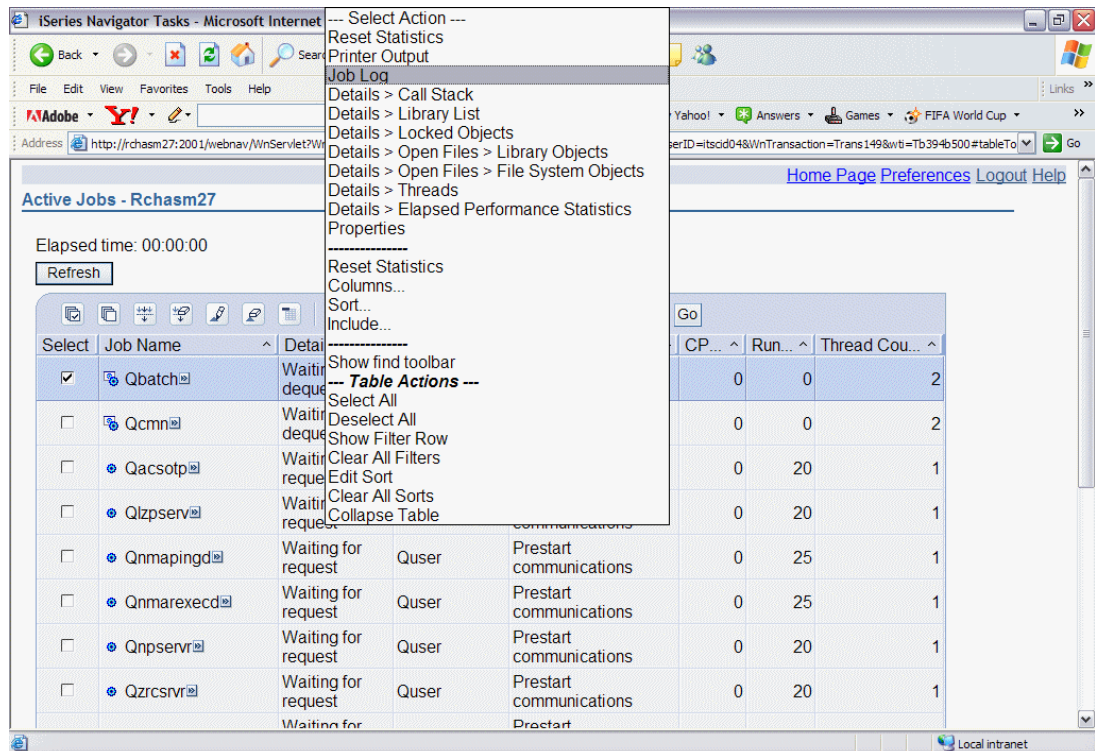


Figure 9-15 Lists on the Web

After selecting an action from the drop-down menu, the **Go** button must be clicked to invoke it. Clicking the arrow icon next to an individual item brings up the context menu for the item. This menu includes only those actions valid for that item.

Additional lists features

In Figure 9-15 you can also see that there are other features. These additional features make up the lists functionality of the iSeries Navigator Tasks on the Web as an administration tool.

The other features for lists include:

- ▶ Button to refresh list data
- ▶ Tool bar to manipulate data
 - Sort on multiple columns
 - Filter list data
 - Select/deselect all
- ▶ Navigation area
 - Page forward or backward
 - Go directly to a specific page

- ▶ List status information
 - Number of items in list
 - Number of items filtered
 - Number of items displayed
 - Number of items selected
- ▶ List *freshness indicator*
 - Indicates how recently the data was retrieved from the iSeries

iSeries Navigator preferences

Default preferences for iSeries Navigator Tasks on the Web can be set through the Preferences task link located at the top of any iSeries Navigator Web page.

- ▶ General preferences
 - Default number of list items shown per page
 - Indicate if new lists are shown in separate window
- ▶ Database preferences
 - Default naming option
- ▶ URL parameter defaults
 - Target system
 - User ID
 - Database
 - Schema

A sample screen of this function is shown in Figure 9-16.

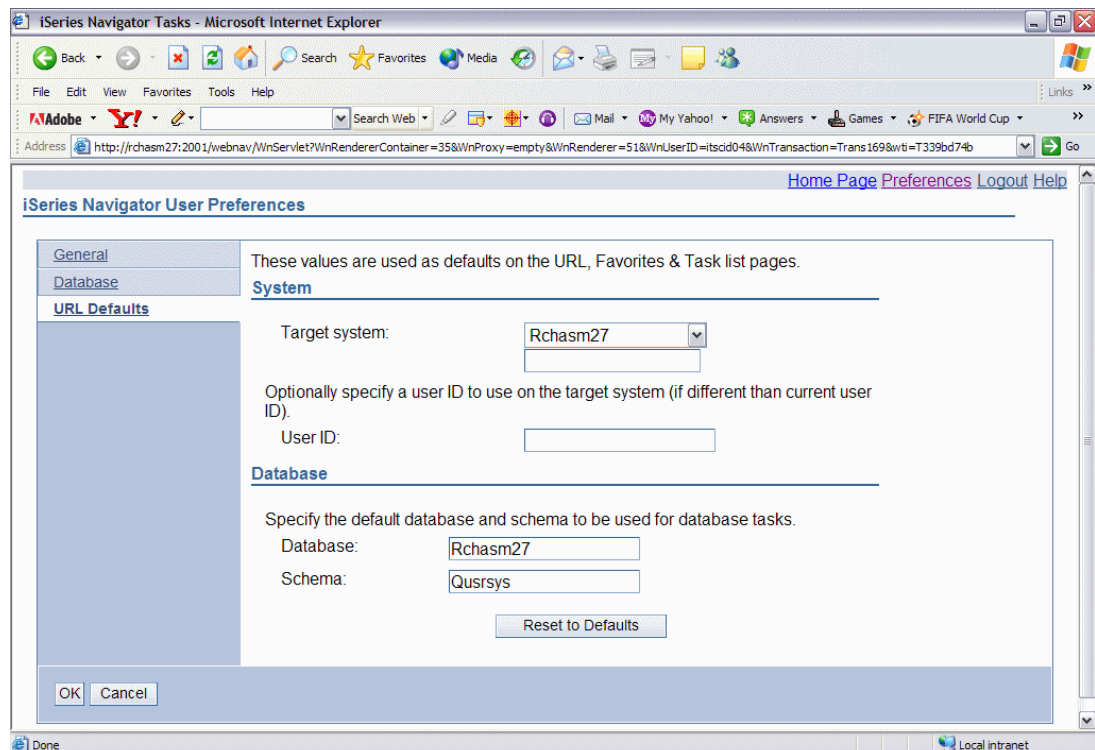


Figure 9-16 iSeries Navigator Tasks on the Web: preferences

These properties are stored in the Java registry and can only be changed through this interface. They are user specific, meaning changes made by a user only affect that user.

List of new functions for iSeries Navigator Task on the Web in V5R4

There are quite a number of functions that have been added to the iSeries Navigator Tasks on the Web in V5R4. See Figure 9-17 and Figure 9-18 for a list of these new functions.

V5R4 Tasks on the web

(Tasks with ** were added in V5R4)

- **Basic operations**
 - Messages (msg)
 - System Operator Messages (sysoprmsg)
 - **QSYSMSG messages (qsysmsg)
 - Send a Message (sndmsg)
 - Printer Output (prtout)
 - ** Display Printer Output (dspprtout)
 - ** Reply to a Message for Printer Output (rpyprtout)
 - ** Release Printer Output (rlsprtout)
 - ** Move Printer Output (movprtout)
 - ** Delete Printer Output (dltprtout)
 - ** Convert Printer Output to PDF (cnvprtout)
 - ** Printer Output Properties (prtoutprop)
 - User Jobs (usrjob)
 - ** Printers (prt)
 - ** Reply to a message for a Printer (rpyprt)
 - ** Hold a Printer (hldprt)
 - ** Release a Printer (rlsprt)
 - ** Start a Printer (startprt)
 - ** Stop a Printer (stopprt)
 - ** Restart a Printer (restartprt)
 - ** Printer Properties (prtprop)
 - ** Make a Printer Available (availprt)
 - ** Make a Printer Unavailable (unavailprt)
- **Work Management**
 - Active Jobs (actjob)
 - Server Jobs (svrjob)
 - Active Subsystems (sbs)
 - Output Queues (outq)
 - ** Clear and Output Queue (clroutq)
 - ** Hold an Output Queue (hldoutq)
 - ** Release and Output Queue (rlsoutq)
 - Delete a Job (dljob)
 - Hold a Job (hldjob)
 - Release a Job (rlsjob)
 - Move a Job (movjob)
 - Work with a Job (wrkjob)
 - Job Properties (jobprop)
 - Display Joblog for a Job (joblog)
 - Display Locked Objects for a Job (lockobj)
 - Display Call Stack for a Job (callstack)
 - Display Open Files for a Job (openfiles)
 - Display Library List for a Job (liblist)
 - Display Performance Statistics for a Job (perfstats)
 - Display Threads for a job (threads)
 - ** All Job Queues (alljobq)
 - ** Active Job Queue (actjobq)
 - ** Hold a Job Queue (hldjobq)
 - ** Release a Job Queue (rlsjobq)
 - ** Clear a Job Queue (clrjobq)
 - ** Active Memory Pools (actpool)
 - ** Shared Memory Pools (shrpool)
 - ** System Status (syssts)

Figure 9-17 V5R4 tasks on the Web functions

- **Database**
 - All Objects (db.allobj)
 - Aliases (db.alias)
 - Create an Alias (db.crtalias)
 - Constraints (db.cst)
 - Indexes (db.idx)
 - Create an Index (db.crtidx)
 - Sequences (db.seq)
 - Create a Sequence (db.crtseq)
 - Tables (db.tbl)
 - Create a Table (db.crttbl)
 - Distinct Types (db.typ)
 - Create a Distinct Type (db.crttyp)
 - Journals (db.jrn)
 - Journal Receivers (db.jrnrcv)
 - SQL Procedures (db.proc)
 - SQL Packages (db.pkg)
 - Create a Schema (db.crtschema)
 - Select Schema (db.selschema)
 - Triggers (db.trg)
 - Views (db.view)
 - Functions (db.func)
 - ** Table Partitions (db.tblpart)
 - ** Schemas (db.schema)
 - ** Databases (db.database)
 - ** Table Indexes (db.tblidx)
 - ** Performance Monitors (db.perfmon)
 - ** Create Performance Monitor (db.crtmon)
 - ** Import Data (db.import)
 - ** Export Data (db.export)
 - ** Index Advisor (db.idxadv)
- **Configuration and Service**
 - Work with System Values (sysval)
 - Time Management (timemgmt)
 - ** Disk Status (dsksts)
 - ** Change Password (chgpwd)
 - ** History Log (dsplog)
- **Network**
 - ** TCP/IP Servers (tcpsvr)
 - ** Universal Connection Wizard (ucw)
- **General iSeries Navigator Tasks**
 - iSeries Navigator Tasks Home Page (home)
 - ** View All Tasks (list)
 - ** Trace Levels (trace)
 - ** User Preferences (pref)
 - ** Configuration Options (config)
 - ** Command Prompter (runcmd)
 - ** Work with Jobs (wrkjobs)
 - ** Work with Messages (wrkmsgs)
 - ** Work with Printer Output (wrkprtout)

Figure 9-18 V5R4 tasks on the Web functions (continued)

iSeries Navigator Tasks on the Web - Security

iSeries Navigator Tasks on the Web can be secured by Secure Socket Layer (SSL).

There are two kinds of Secure Socket Layer connections that you need to consider configuring in order to run iSeries Navigator tasks on the Web securely:

- ▶ The first type of SSL connection is used in a connection between a Web browser and the iSeries system that is hosting iSeries Navigator tasks on the Web.
- ▶ The second type of SSL connection is used by the iSeries Navigator tasks on the Web interface to retrieve data from the local iSeries and any managed secondary systems.

Authentication filter and SSL warning

By default, customers receive a WARNING MESSAGE when first connecting to iSeries Navigator tasks on the Web. If the iSeries machine is safely behind a firewall, SSL can be switched off.

This can be controlled by the configuration parameters in the IFS file `webnav.properties` in `/QIBM/UserData/OS400/iSeriesNavigator/config/webnav.properties` or by using iSeries Navigator Tasks configuration. The link to this configuration page can be found in the iSeries Navigator Tasks on the Web home page (see Figure 9-6 on page 196).

Controlling access to iSeries Navigator Tasks for the Web

User access to iSeries Navigator Tasks for the Web is controlled by application administration in the PC Client version of the iSeries Navigator.

- ▶ Whether the user can configure iSeries Navigator Tasks for the Web on the server
- ▶ Whether the user can manage the server using iSeries Navigator Tasks for the Web
- ▶ Whether the user can use iSeries Navigator Tasks for the Web from this server

Initially this is set to all object access but can be set up to the desired access level for each desired user.

9.2 iSeries Access for Web enhancements

iSeries Access for Web, 5722-XH2, offers browser-based access to iSeries servers. iSeries Access for Web enables end users to leverage business information, applications, and resources across an enterprise by extending the iSeries resources to the client desktop through a Web browser. iSeries Access for Web is server based and implemented using Java Servlet technology using industry standard protocols (HTTP, HTTPS, and HTML). It runs batch commands, provides access to a database, an Integrated File System, printers, printer output, and messages, and a 5250 emulator.

iSeries Access for Web provides a lightweight user deployment, requiring only a browser on the client, as it installs on an iSeries server and eliminates the install, configuration, and management of any client code. iSeries Access for Web provides both a set of servlets (for use on Tomcat and WAS Web application servers) and a set of portlets (for use on WAS Portal products).

The iSeries Access for Web servlet code can be used with any of the following Web application servers:

- ▶ WebSphere Application Server V6.1 for i5/OS
- ▶ WebSphere Application Server Network Deployment V6.1 for i5/OS
- ▶ WebSphere Application Server - Express V6.1 for i5/OS

- ▶ WebSphere Application Server V6.0 for i5/OS
- ▶ WebSphere Application Server Network Deployment V6.0 for OS/400
- ▶ WebSphere Application Server V5.1 - Express for iSeries
- ▶ WebSphere Application Server V5.1 for iSeries
- ▶ WebSphere Application Server V5.0 - Express for iSeries
- ▶ WebSphere Application Server V5.0 for iSeries
- ▶ Apache Software Foundation (ASF) Tomcat Web serving environment
- ▶ The i5/OS integrated Web application server

This support became available November 2006 via PTFs. See the following Web site for information about getting this support:

<http://www.ibm.com/eserver/iseries/access/web/servicepacks.htm>

Additionally, iSeries Access for Web portlet code can be used with any of the following portal servers:

- ▶ IBM Workplace Services Express V2.5
- ▶ WebSphere Portal for Multiplatforms V5.1.0.1
- ▶ WebSphere Portal Express for Multiplatforms V5.0.2
- ▶ WebSphere Portal Express Plus for Multiplatforms V5.0
- ▶ IBM Workplace Services Express V2.6
- ▶ IBM WebSphere Portal V6.0

In V5R4 enhancements are made for certain functions on both the servlet and portlet versions of the iSeries Access for Web product.

Notes: Using iSeries Access for Web requires that you have a license for the iSeries Access Family (5722-XW1) and comply with the usage rules according to the 5722-XW1 license.

During October 2006, IBM made available a no charge a set of Web enablement enhancements. These enhancements include the capability of iSeries Access for Web to operate under an integrated Web application server instead of a WebSphere Application Server instance/profile. For more information about this see Chapter 14, "i5/OS-based Web enablement enhancements" on page 337.

Host Access Transformation Server Limited Edition (HATS LE) is no longer being bundled with the 5722-XH2 iSeries Access for Web product. For customers that wish to continue to use this type of functionality, the standalone HATS product will need to be purchased. Refer to the HATS Web site for more information:

<http://www-306.ibm.com/software/webservers/hats/>

9.2.1 Summary of V5R4 enhancements

This section lists a summarized version of the iSeries Access for Web V5R4 enhancements. Most of them have sections that follow, showing more details.

- ▶ *Single Sign-on:* Prior to V5R4, iSeries Access for Web used HTTP basic authentication to retrieve the iSeries user profile and password from the user. This profile information could not be shared with other iSeries Web-based applications. Sometimes this resulted in the user being prompted multiple times for the same information by different applications. In V5R4, iSeries Access for Web also supports the WebSphere security model to provide single sign-on support among WebSphere Application Server applications. This support is

referred to as application server authentication in the iSeries Access for Web documentation.

- ▶ *Style sheets*: The look of the iSeries Access for Web page content is now controlled by external style sheets. Users can provide their own style sheets if they desire a look different from the default iSeries Access for Web look.
- ▶ *5250 function*: There are several enhancements to 5250 sessions support, including Bypass sign-on, HTML keyword support, and hotspot support.
- ▶ *Customize function*: You can now transfer configuration data. This enables iSeries Access for Web administrators to transfer 5250 sessions and macros, saved commands, database requests, My Folder items, and policies from one user to another.
- ▶ *Database function*: This function has several enhancements:
 - *Open Office*: The Open Office spreadsheet format (.ods) is now supported in V5R4 iSeries Access for Web RUN SQL and Copy Data to Table functions, and also in Dynamic SQL, SQL Query, and the new SQL Requests functions.
 - *WebSphere Data Sources*: Customize now supports using WebSphere data source names to create database connections. Data source connections can be used with any of the iSeries Access for Web database functions.
 - *Integrated File System Destination*: Run SQL now has support for storing SQL results in the iSeries Integrated File System.
 - *Import Query*: Import Query support allows customers to import query for iSeries and DB2 UDB for iSeries Query Manager query file contents into iSeries Access for Web database requests.
 - *Find Record*: Find record provides a way to search for a record in a database table and display the contents of that record. Find record is available as an action of the tables functions.
- ▶ *Printer output*: Several preferences have been added or modified:
 - The Printer output list view and Printer output list filter preferences have been added.
 - The Restore original printer output status setting has been added to the Default PDF output settings preference.
 - The Bypass PDF output settings page setting has been added to the Default PDF output settings preference.
 - The Output queue list filter preference supports a new special value, *USRPRF. This limits the output shown to the signed on user to be only the output queue specified in that user's user profile.
- ▶ In the portlet environment the following enhancements or changes apply:
 - Database functions: Data source connections are not available in this environment. If you need data source connections, consider configuring a WebSphere Application Server environment where WebSphere data source names can be used to create database connections.
 - Web application server authentication: iSeries Access for Web will perform user authentication. If you require a single sign-on (SSO) environment, consider configuring a WebSphere Application Server environment where application server authentication options are available.

9.2.2 Servlet version enhancements

The following topics present more details on selected iSeries Access for Web enhancements under the servlets environment.

Style sheet support

The look of the iSeries Access for Web page contents is now controlled by externalized style sheets. Users, understanding style sheet support, can substitute their own style sheets if they want the iSeries Access for Web pages to appear differently from defaults.

Users that have customized templates being displayed by iSeries Access for Web will need to modify their templates to include the new style information.

Support for single sign-on

Previously, iSeries Access for Web used only HTTP basic authentication to retrieve the user ID and password from the user.

Since this profile information could not be shared with other iSeries Web-based applications, this often resulted in the user being prompted multiple times from different applications for the same user ID and password information.

iSeries Access for Web now supports the WAS security model so that single sign-on can be accomplished among the WAS applications.

This enhancement ties into the i5/OS single sign-on (SSO) support (that is, the use of identity tokens) that hold user credentials such as password certificates. EIM can be used to map a non-i5/OS user to an i5/OS user profile, with DAP as the user repository,

New configuration parameters are used to set up this type of environment and are described in the V5R4 Information Center.

Important: The bypass sign-on feature cannot be used if iSeries Access for Web is configured to use application server authentication.

5250 session enhancements

5250 session support has been available since iSeries Access for Web was initially made available. In V5R4, running an iSeries Access for Web 5250 session has more flexibility in:

- ▶ Bypassing 5250 session sign-on. This enables users to bypass the standard i5/OS sign-on screen.
- ▶ Target partition flexibility. With V5R4, in a network of at least two i5/OS partitions (or systems), the 5250 session target i5/OS partition does not need to have the iSeries Access for Web product installed. However, that target i5/OS partition must have Telnet server running and accessible.

This means at least one of the i5/OS partitions must have the iSeries Access for Web product installed. Then a browser user can run an iSeries Access for Web 5250 session to that i5/OS partition or any other partition with Telnet up and running.

Macro support is also included for 5250 functions in this environment. Macros can also be recorded and played in an active session. They can also be modified, copied, deleted, and renamed. These macros can be shared with other users through their shortcuts.

- ▶ Displaying HTML data in a field. This setting is used to display HTML data defined on application screens using the display file DDS HTML keyword. This setting cannot be changed once the 5250 session is active.
- ▶ Showing a hotspot link. This setting specifies whether an HTTP URL in a non-input field is displayed as a link. Choosing the link will open the HTTP URL in a different browser.

Figure 9-19 shows the new Bypass Sign-on and Display HTML options on Start Session.

The screenshot shows the 'Start Session' configuration page. The 'Bypass signon' and 'Display HTML data in fields' options are highlighted with a red rectangular box. The 'Start Session' button is located at the bottom center of the main content area.

Figure 9-19 Bypass sign-on support

Customize function

In addition to the preferences, policies, and settings customization support, the transfer configuration functionality has been added in V5R4.

This function is used to copy or move configuration data from one user (or group) profile to another. Configuration data includes 5250 session and macros, saved commands, database requests, folder items, and policy settings.

Important: The policies, settings, and transfer configuration functions can only be accessed by users with iSeries Access for Web administrator privileges or an i5/OS user profile with *SECADM special authority. A user profile can be given administrator privileges using the customize function and setting the grant administrator privileges policy to allow.

Database function enhancements

Some changes and additions to database functions have been made in iSeries Access for Web in V5R4. Each of these new enhancements are discussed in this section.

RUN SQL - new destination option

In this function a new destination option is added. In V5R4, you can now have the option to store your data in the Integrated File System (IFS), as shown in Figure 9-20.

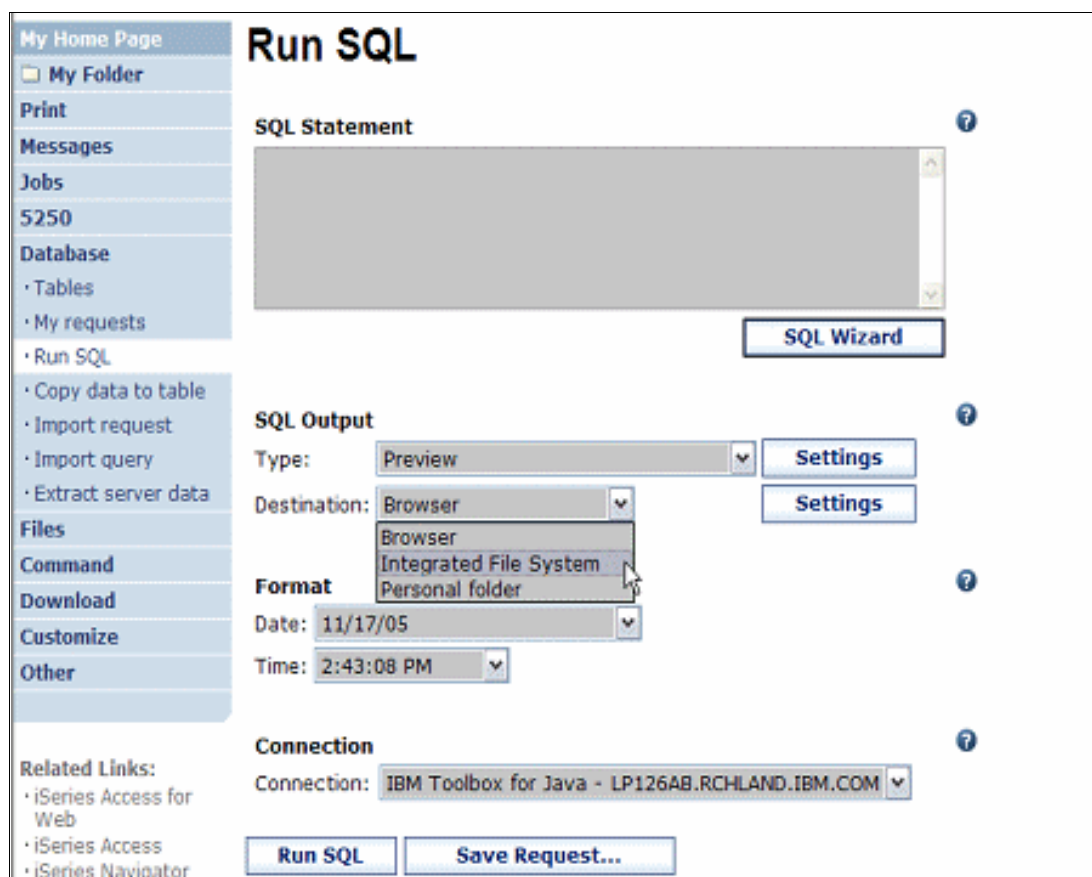


Figure 9-20 RUN SQL: new destination option

WebSphere data sources enhancement

You can now set up to use WebSphere data source names when creating database connections in the Customize tab, Database Connections link (see upper left window **1**) Figure 9-21 on page 213).

This enhancement enables many different applications running under WebSphere to use the same data source connections, which helps minimize Java heap size when many sessions are actively accessing DB2 for i5/OS. The windows in Figure 9-21 on page 213 show the progression through defining and saving a new database connection named IWAJCDB for user JCOOL.

Select **Add databases connection** in window **2**. In window **3** enter the connection name, driver class, and JDBC URL. We read the online help text for each field by clicking the question mark (?). You can select to use either a:

- ▶ Driver class and JDBC URL
- ▶ WebSphere Datasource

We also reviewed the iSeries Information Center documentation to get the JDBC URL correct. You see we specified to use libraries (schemas) PFREXP and QGPL.

Window 4 shows the newly defined connection joining the default one provided with i5/OS. In our simple example we chose the same JDBC servlet as the default but narrowed down a default library list for JCOOL that only has libraries PFREXP and QGPL.

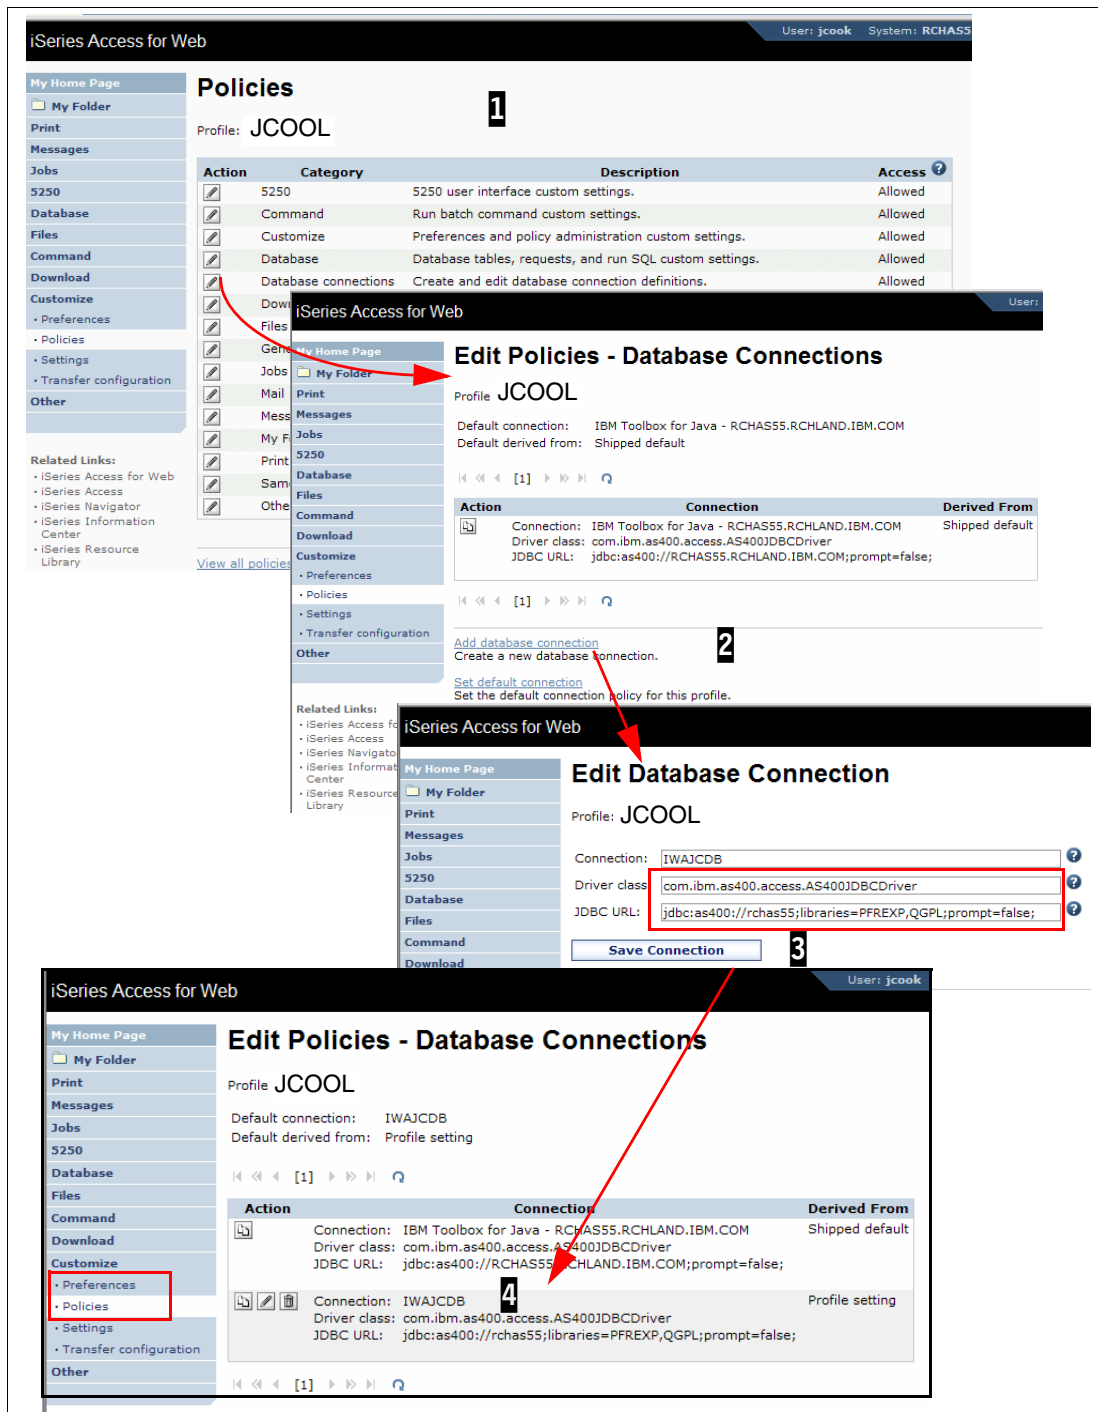


Figure 9-21 Edit policies: database connections

You need someone familiar with WebSphere to set up this specific database connection for maximum efficiency, and possible coordination with a data source that may be specified for the WebSphere Application Server profile in which iSeries Access for Web has been deployed.

The Import Query support tool allows you to import IBM Query for iSeries (5722-QU1) and DB2 Query Manager (5722-ST1) query file contents and converts them into iSeries Access for Web database requests, as shown in Figure 9-23. Both *QMORY and *QRYDFN query file types are supported.

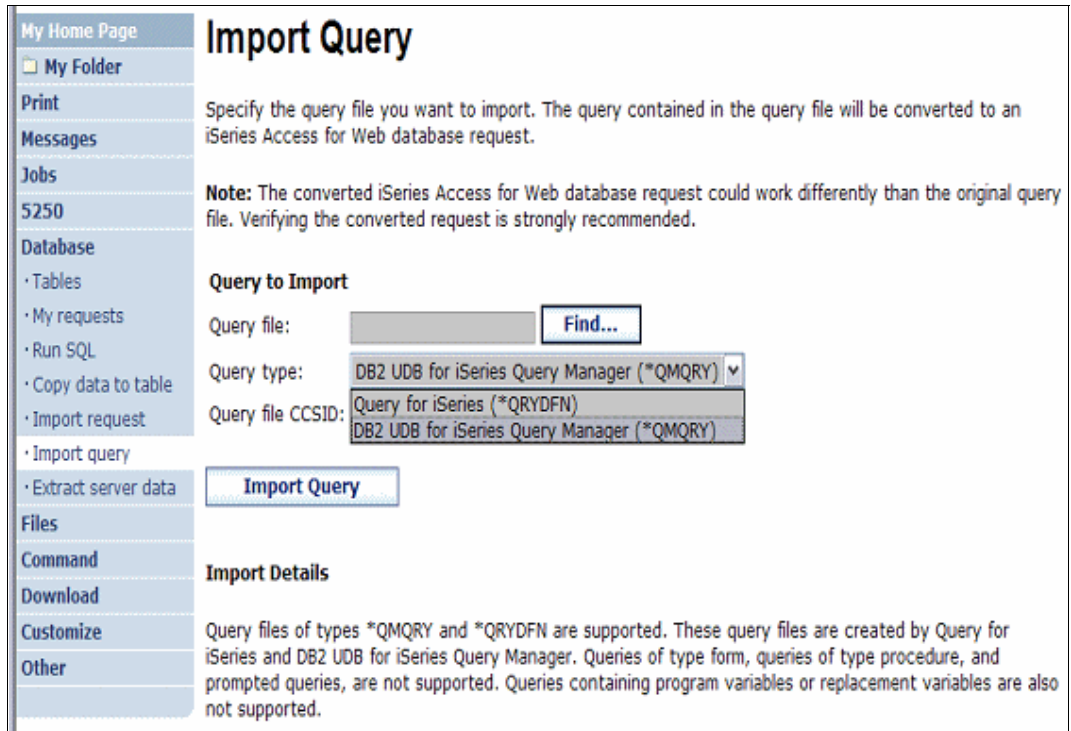


Figure 9-23 Import Query tool

Print function enhancements

With the availability of PTF SI23024, the iSeries Access for Web Print function includes the following enhancements:

- ▶ New printer output filtering preferences are available to control which spool files are displayed in the printer output list.
- ▶ The output queue filter preference supports a new special user profile value, *USRPRF. This value enables shows only the output queue that has been specified on the output queue (OUTQ) parameter of the authenticated user's i5/OS user profile.
- ▶ A new PDF output setting preference was added to control whether printer output status is changed to *saved* after viewing printer output as a PDF document.
- ▶ A new PDF output setting preference was added to control whether to bypass the intermediate PDF output settings page for selection. When this new setting is enabled, requests to view printer output as a PDF document automatically display the PDF file in the browser window.

9.2.3 Portlet version enhancements

Like the servlet version, iSeries Access for Web in V5R4 added new features to its functions. This section discusses the new features in each function.

Database function

In V5R4, an SQL requests portlet has been added so that create, save, manage, and run saved SQL requests are now possible. A viewer is also added so that results of running saved iSeries SQL requests can be viewed.

Figure 9-24 shows a screen capture of the portlet database window showing the groups of SQL functions now possible. No actual SQL statement or results are shown.

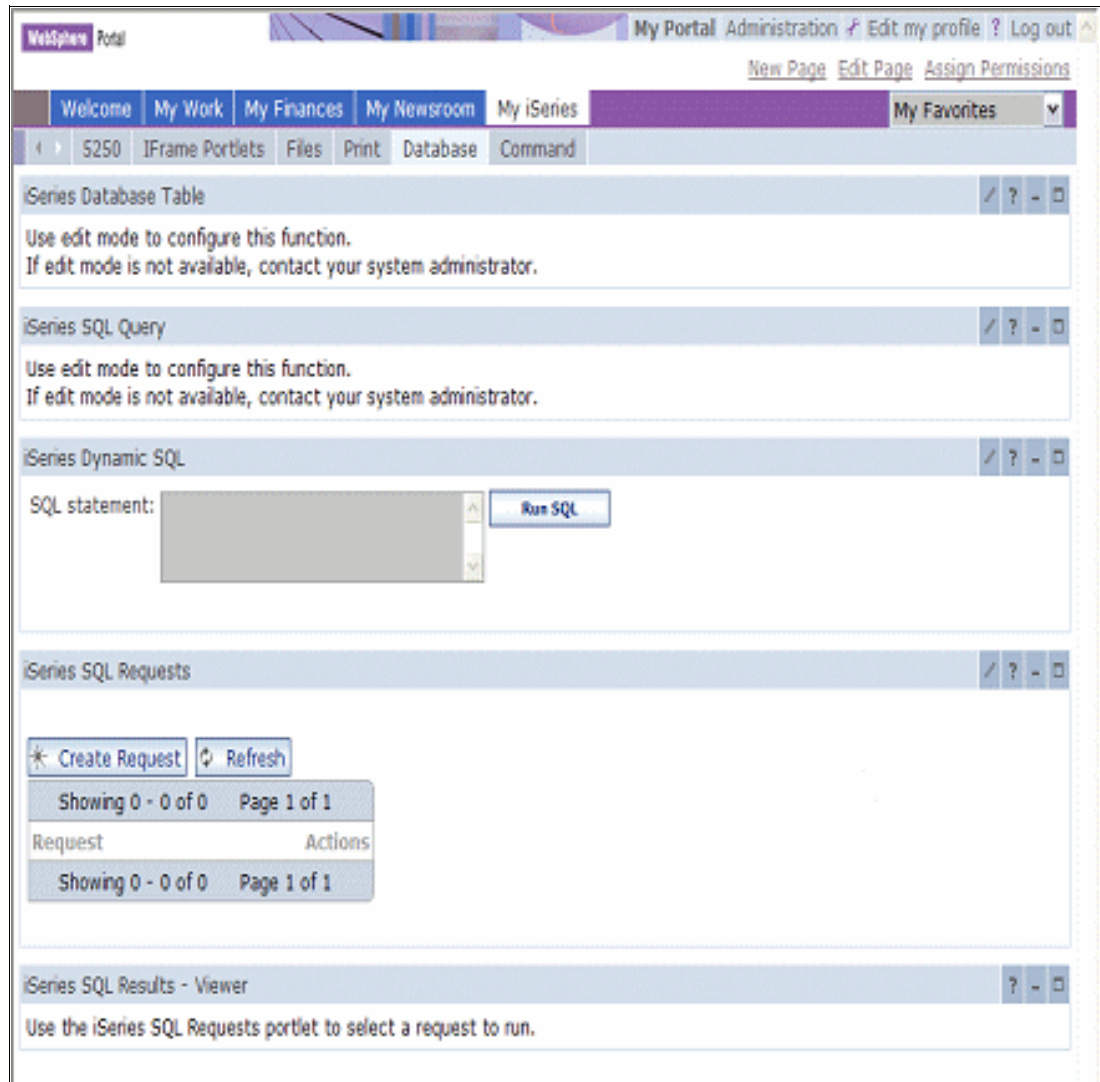


Figure 9-24 Database function enhancements: Portal version

File function

In the portlet version of iSeries Access for Web in V5R4, a zip function has been added so that you can zip files in the Integrated File System (IFS), as shown in Figure 9-25.

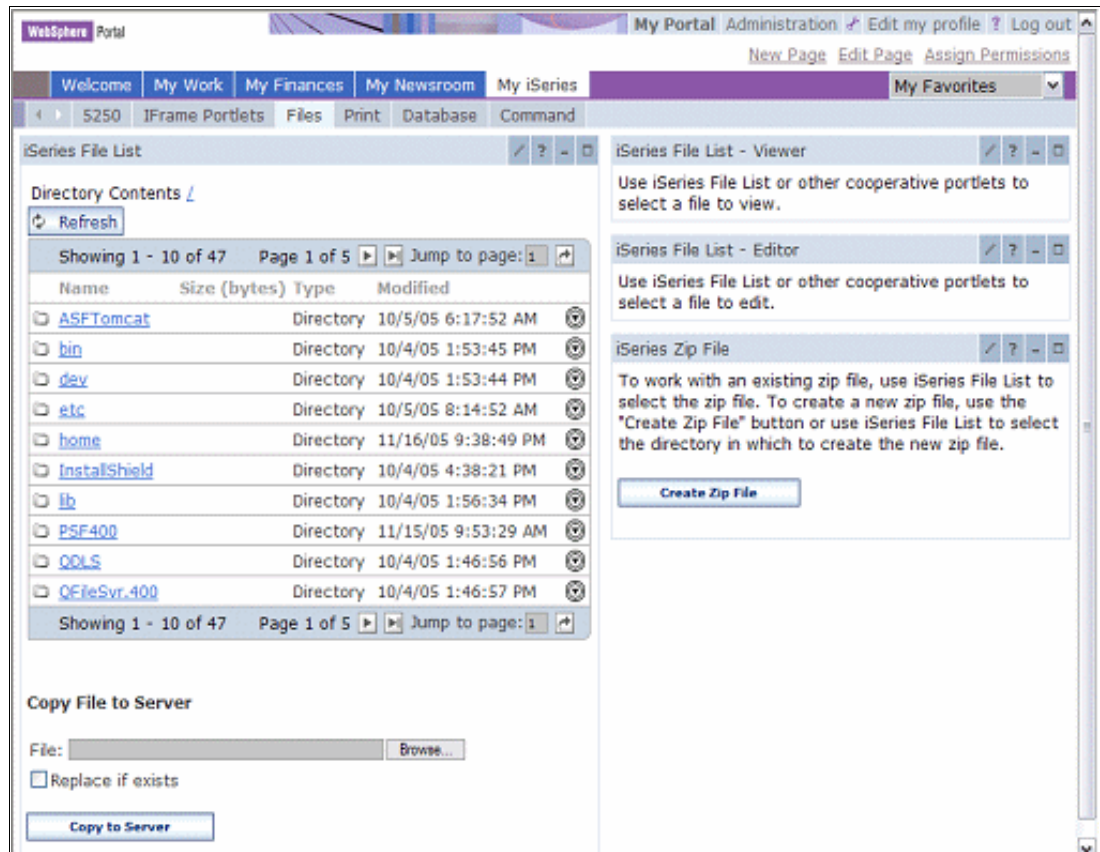


Figure 9-25 Zip of IFS files function

You can also add files to your zip file, extract them from an existing zip, and display the contents of a zip file in the IFS.

User profile credentials

Users can now change their i5/OS user profile password through the iSeries Access for Web portal version. Changing the i5/OS user profile password using this interface also provides a function to update currently defined vault credentials to match the new password value.

For more information about iSeries Access for Web and security under a portal server, refer to iSeries Information Center V5R4 at:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp>

Search with *portal +server +credentials*.

9.2.4 Extract server data

Though introduced in iSeries Access for Web V5R3, we mention this function here because many i5/OS users are unaware of the function. Depending on whether you already have a tool for managing object existence and size, this function can be very useful to you if, *ad hoc*, you want to look at objects and some of their attributes across the entire partition.

Use *extract server data* to select a specific object type and specify the library (schema) and database file (table) where those results can then be retrieved. You can access the data in this file in a variety of ways, including using iSeries Access for Web *table* functions (add/delete a row, quick view, find a record, and more).

Object-specific information can include:

- ▶ Library
- ▶ Name
- ▶ Type
- ▶ Status
- ▶ Size
- ▶ Object owner
- ▶ Description

Figure 9-26 shows two window examples:

- ▶ Selecting the object type (journal receivers) from a long list of possible object types, which include directory entries, messages, software fixes, software products, system pool, user profiles, and more.
- ▶ An excerpt of the Quick View of the results showing all journal receivers within our i5/OS partition. We found the library, size, and owner of many journal receivers in our i5/OS partition, including those we had no idea were taking up space and not being used.

The screenshot displays two windows from the iSeries Access for Web interface. The top window is titled "Extract Server Object Data" and shows a configuration screen for extracting data from a server. A red arrow points to the "Object Data to Extract" dropdown menu, which is currently set to "System object - journal receiver". Below this, there are sections for "Table to Extract", "Output", "Connect", and "Object Data". A warning message at the top states: "Error: Table PPREXP.EXTRACTINFO exists. Select the replace option to overwrite the table." Below the configuration options, there are instructions: "The table name must be specified using dotted SQL notation. For example, MYSCHEMA.MYTABLE is a valid table name." and "Select the **Replace table if table exists** option to overwrite an existing table. If the table exists, and the option is not selected, an error message will be displayed."

The bottom window is titled "Quick View of PPREXP.EXTRACTINFO" and shows a table of extracted journal receiver objects. The table has columns for LIBRARY, NAME, TYPE, STATUS, SIZE, OWNER, and DESCRIPTION. The list is limited to 1000 items. The table contains 17 rows of data, showing various journal receiver objects with their respective libraries, names, types, statuses, sizes, owners, and descriptions.

LIBRARY	NAME	TYPE	STATUS	SIZE	OWNER	DESCRIPTION
ITSOWSE6	QSQRN0008	JRNRCV		221184	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0009	JRNRCV		221184	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0010	JRNRCV		221184	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0011	JRNRCV		221184	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0012	JRNRCV		221184	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0013	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0014	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0015	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0016	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0017	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE6	QSQRN0018	JRNRCV		233472	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0001	JRNRCV		78327808	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0002	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0003	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0004	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0005	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0006	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0007	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0008	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0009	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0010	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0011	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0012	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0013	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL
ITSOWSE61	QSQRN0014	JRNRCV		200704	WSE6DBUSER	COLLECTION - created by SQL

Figure 9-26 Quick view example of extracted journal receiver objects



i5/OS-based communications and general connectivity

In this chapter we introduce the enhancements in the areas of communications and selected general networking connectivity in i5/OS V5R4, which includes the following major topics:

- ▶ Systems Network Architecture (SNA) Enterprise Extenders
- ▶ Internet Protocol (IP)
- ▶ HTTP support
- ▶ i5/OS support for Windows Network Neighborhood (iSeries NetServer™)
- ▶ iSeries Access family of products

10.1 SNA and the Enterprise Extenders

The i5/OS V5R4 supports SNA functions and protocols. As time goes on, direct support of SNA capabilities by new LAN adapters, new routers, and new switches is typically not available. As a result, IBM and specifically i5/OS has supported AnyNet for several years.

However, there are shortcomings with AnyNet, including more system overhead when exchanging data, and some SNA capabilities are not supported. High Performance Routing (HPR) is one of them.

V5R4 support of Enterprise Extenders (EE) provides expanded SNA capabilities and has improved performance characteristics over AnyNet.

Enterprise Extender

Enterprise Extender is a networking architecture that allows SNA applications to run over IP networks using HPR. This is the preferred way to run SNA applications over IP networks with communications input/output adapters (IOAs), such as Gigabit Ethernet, since these IOAs do not require an input/output processor (IOP). Gigabit Ethernet adapters do not automatically support SNA traffic. Enterprise Extender (or AnyNet) is required to allow SNA data to flow over a Gigabit adapter. IBM recommends that Enterprise Extender be used in place of AnyNet. For more information see the migration from AnyNet to Enterprise Extender in V5R4, iSeries Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzajt/rzajtanytoee.htm>

Enterprise Extender is an industry-standard solution defined by the APPN Implementer's Workshop (AIW) and the Internet Engineering Task Force (IETF) - RFC 2353. With Enterprise Extender, the Rapid Transport Protocol (RTP) endpoint views its interface with the UDP layer of the stack as just another data link control, and treats the connection across the IP network the same as if it would any SNA connection.

10.1.1 SNA function support under EE

i5/OS V5R4 offers full support for Advanced Program to Program Communication (APPC) and Advanced Peer to Peer Networking (APPN) over IP networks with EE (Figure 10-1). This also offers support for dependent logical unit (LU) communications to mainframes over Dependent LU Requester (DLUR). This enables businesses to run existing SNA applications over an IP network without the expense of a parallel or a separate network infrastructure.




Figure 10-1 SNA support with Enterprise Extenders

With this support functionality it is now possible for existing applications to pass through new generations of routers and network equipment that do not directly support SNA.

Aside from giving support for a broad range of SNA protocols and applications, it also has improved performance with High Performance Routing (HPR), as compared to AnyNet.

HPR's Rapid Transport Protocol component provides:

- ▶ Error detection with selective retransmission of lost packets.
- ▶ Nondisruptive reroute based on class of service requirements. HPR preserves the session without impact to the end user for planned and unplanned outages in the session path. In the event that no alternate path is available, HPR can even be configured to preserve the session while the failing component is recovered.
- ▶ Proactive congestion control. Enterprise Extender brings with it an enhanced version of HPR's Adaptive Rate-Based (ARB) congestion control algorithm. The new version, Responsive-Mode ARB, is more aggressive in using available bandwidth and more tolerant of variations in network latency. Responsive-Mode ARB was introduced with EE to better allow HPR traffic to coexist with native IP traffic in the backbone network.
- ▶ Prioritization. The SNA priority field is mapped to the IP Type of Service (TOS) byte, which is used by routing algorithms such as the Cisco Weight-Fair-Queuing algorithm. A set of standard UDP ports are also reserved based on priority with packets mapped to them according to the SNA priority field.

Figure 10-2 shows the summary of support functions for SNA under EE.

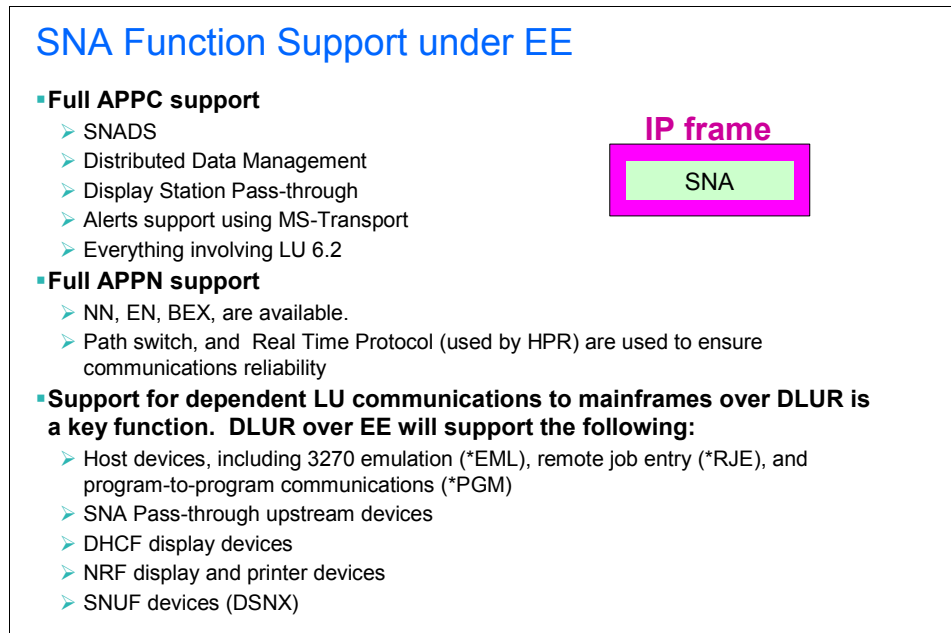


Figure 10-2 Summary of SNA support function under EE

i5/OS V5R4 does Enterprise Extender support does not provide:

- ▶ **IBM APPN Connection Network support:** An IBM APPN connection network is a logical construct used to provide direct connectivity between APPN end nodes (ENs) without the configuration overhead of defining direct connections between every pair of ENs. In general, the process of creating a connection network starts when a LOCATE request is received from an EN.

A network node (NN) is then used to locate the destination specified in the LOCATE request. If the NN sees that the two ENs (source and destination) are attached to the same transport medium (such as token ring), a virtual node (VN) is used to connect the two endpoints and form a connection network. The NN defines the session path as a direct connection from EN1 to VN to EN2, and then traffic is permitted to flow.

A connection network is sometimes also referred to as a version of a shared access transport facility (SATF). An example of a shared access transport facility is a local area network (LAN) where two nodes can communicate directly with each other without the need to use a router

- ▶ **SNA-based remote workstation support, such as 5294, 5394, or 5494 controller:** There are a number of OEM vendor solutions available that allow 5250 type workstations to attach to eServer i5 systems across a TCP/IP network. The 5250 workstations attach to OEM controllers, which encapsulate the workstation 5250 data stream into TCP/IP traffic. The Telnet server handles the communications on the i5/OS system (the 5250 workstations appear as Telnet clients).

There are multiple OEM vendors that provide this capability.

There are products from Perle and BosCom that allow customers to keep using twin-ax devices, but they connect into the system via TCP/IP. The following are Web sites that describe some of the details in this area:

<http://www.e-twinaxcontroller.com/?source=adwords?adcopy=CtwinaxGsnalreplacement20050713>

http://www.perle.com/products/prod_family/as_400/index.html

<http://www.affirmative.net/itwinax.html>

10.2 IPv6 support

Back in the mid-1990s, the Internet Engineering Task Force (IETF) determined that the continuing dramatic increase of the Internet and Internet-enabled devices would exceed the addressing and other capabilities in the then currently deployed IPv4 protocol in the not too distant future. IPv6, the next generation of the Internet protocol, was designed. Initial IPv6 implementations began to appear in the late 1990s, including the first UNIX IPv6 implementation in 1997 (IBM AIX).

IPv6 support on i5/OS is a multi-release project. This project began on V5R2 with a *developer's sandbox* release of IPv6. The V5R2 version included a majority of the socket IPv6 APIs and the base IPv6 protocol infrastructure. V5R2 support was sufficient to allow application developers and business partners the opportunity to start IPv6-enabling their network-based applications. A number of base network tools were also provided in this release including PING (Verify TCP/IP Connection), TRACEROUTE (Trace TCP/IP Route), and NETSTAT (Work with TCP/IP Network Status) i5/OS commands, in addition to the 5250 i5/OS command level. iSeries Navigator also provided interfaces to these tools.

Also in V5R2, the i5/OS communication trace (Start Communications Trace (STRCMNTRC) command) was updated to support capturing and formatting IPv6 network traffic that was sent or received. This developer's sandbox release came with a number of hardware restrictions. For example, IPv6 could only be enabled on one physical Ethernet adapter, which had to be feature code 2838 or 2849, and that adapter had to be dedicated to running only IPv6 traffic.

i5/OS V5R4 removes the hardware restrictions that had previously existed with the earlier V5R2 and V5R3 releases. All models of Ethernet adapters (10 MB, 100 MB, 1 GB, and 10 GB) are now supported, including virtual Ethernet between different partitions on the System i platform. The Ethernet adapter can now be shared between IPv4, IPv6, and PPPoE traffic. The restriction of being able to use only a single Ethernet adapter running IPv6 has also been removed. IPv6 can be enabled on all Ethernet adapters that you have installed on your system.

i5/OS V5R4 also provides additional IPv6 programming APIs and infrastructure enhancements. These enhancements include multicast and fragmentation support for applications using User Datagram Protocol (UDP), complete support of RFC3493: Basic Socket Extensions for IPv6, and other IPv6 protocol improvements.

The IPv6 support on i5/OS V5R4 is *IPv6 Ready*, certified by the IPv6 Ready organization (<http://www.ipv6ready.org>). In order to become certified, i5/OS V5R4 had to pass a series of IPv6 protocol conformance tests as well as interoperability testing with other vendors' implementations of IPv6. Future releases of i5/OS will also be certified as they are released.

Prior to i5/OS V5R4, IPv6 support had to be manually configured and enabled. Starting with V5R4, the IPv6 protocol and the IPv6 loopback interface are automatically started whenever

TCP/IP is started (unless you specify STRIP6(*NO) on the STRTCP command). Manual configuration using iSeries Navigator is still required in order to enable IPv6 support on each of your physical Ethernet adapters and over virtual Ethernet.

As was stated earlier, the IPv6 support on i5/OS is a multi-release project. V5R4 is the latest installment, but it is not the final release. The IPv6 protocol itself is still being refined by the IETF. New protocol specifications and RFCs are still being created. These new enhancements will be provided in future releases of i5/OS. In addition, there are a number of protocol and infrastructure items that were unable to be included in the V5R4 release. These items will also be shipped in future releases of i5/OS.

V5R4 IPv6 support does not include IPv6 enablement of many of the system network applications, such as FTP. Every single application on the system that sends data over the network needs to be examined, modified to use IPv6, tested, and documented. A summary of what is and is not supported through V5R4 is included in the following paragraphs.

The standard IPv4, 32-bit, dotted decimal IP address is used by applications in many different ways: as parameters or options at startup, stored away in databases or configuration files, written out to log files, displayed in error and diagnostic message, and so on. In each of these cases, the application must be updated to allow the longer IPv6 addresses to be entered, stored, and displayed correctly.

There are nearly 100 network-enabled applications that ship with the base i5/OS system. In the V5R4 release, only a small subset of these applications have been enabled. The remaining applications are being examined and prioritized to determine which should be updated to support IPv6.

However, while IBM is updating these provided applications, you should consider updating your own applications to handle IPv6. Since this environment will likely be new to your business, you must start learning about IPv6 and start modifying your own applications as soon as possible.

Refer to Appendix G, "i5/OS V5R4 IPv6 restrictions" on page 543, for a summary list of i5/OS IPv6 support through i5/OS V5R4.

The following section provides additional background information about IPv6.

10.2.1 More IPv6 summary information

The addressing capacities of IPv4 and IPv6 are:

- ▶ IPv4 has 32-bit addresses, the majority of which have already been assigned.
 - Exhaustion of the IPv4 unallocated address pool is predicted to occur in 2011. See the following text for some debate on this date.
 - Additional addresses cannot be retrofitted into IPv4.
- ▶ IPv6 has 128-bit address, which enables an almost limitless number of addresses: 340,282,366,920,938,463,463,374,607,431,768,211,456.

Some issues with current IPv4 protocol are urgent, whereas others are simply logistical. Thanks to technologies such as Classless Inter-Domain Routing (CIDR) and Network Address Translation (NAT), the unallocated IPv4 addresses pool has continued to last longer than any of the IETF designers back in the mid-1990s had ever anticipated when they were working on IPv6. Unfortunately, these IPv4 *add ons*, while allowing increased usability, have also resulted in performance degradation in end-to-end network traffic. The IETF has taken the knowledge gained from the last thirty years of networking and designed the IPv6 protocol

to be more efficient and safe. In other words, whereas IPv4 was designed to communicate in an experimental academic setting, IPv6 was designed with the ever-evolving world business model in mind.

In addition to *patched up* deficiencies of IPv4, the simple size of an IPv4 address is severely limiting. Its 32-bit address gives the world about four billion addresses, most of which are already taken. The date the remaining IPv4 unallocated addresses will be depleted is under constant debate. Various organizations use many calculations and weight certain factors differently. Some projections state that the unallocated IPv4 address pool will run out as early as 2008. Others are more optimistic and think that the unallocated IPv4 address pool may last until 2014 or even later. The Web site <http://bgp.potaroo.net/ipv4/>, an often quoted Web site that tracks the actual IPv4 address allocations on a day-to-day basis, currently projects that the eventual depletion of the IPv4 unallocated address pool will occur in June of 2011.

Many countries in Asia and Pacific countries cannot wait for 2011. These countries were unable to reserve much of the original allocations of the IPv4 address space and have undertaken major efforts to develop, test, and implement IPv6 networks. They have little choice in the matter, as there are nowhere near enough IPv4 addresses for such a large population of new and expanding businesses.

The United States Government Accounting Office published a report in May 2005 titled "Internet Protocol Version 6: Federal Agencies Need to Plan for Transition and Manage Security Risks" (<http://www.gao.gov/new.items/d05471.pdf>). As a result of this report, the U.S. Office of Management and Budget issued a mandate that all U.S. Federal Agencies must be running IPv6 by June 2008. More information about this can be found at:

<http://www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf>

The June 2008 date has become the new *Y2K deadline* that is impacting the entire computer and network industry. This deadline is forcing many U.S. companies to accelerate their internal plans for IPv6, especially those companies that sell products to the U.S. Government.

Although existing IPv4 networks and the IPv4 protocol will remain around for many years to come, IPv6 is here and it is time to move forward. Now is the time to start planning your transition to IPv6.

Although IPv4 will remain in the network scene for a long time, IPv6 should eventually become more common throughout the Internet as a majority of impacted products and functions are modified to use this protocol.

10.3 TCP/IP configuration enhancements

A new parameter was added to TCP/IP interfaces in V5R4. The parameter alias name provides administrators with an option to define a name for an IP address. i5/OS commands and interfaces also support the alias name. For example, interfaces can be started or stopped via a name rather than an IP address. Similarly, they can be displayed and changed in the Work with TCP/IP Interfaces screen, as shown in Figure 10-3. This is especially useful when dealing with IPv6 interface addresses.

Work with TCP/IP Interfaces				System: RLTSYS
Type options, press Enter.				
1=Add 2=Change 4=Remove 5=Display 9=Start 10=End				
Opt	Internet Address	Subnet Mask	Interface Status	Alias Name
	172.5.92.48	255.255.255.128	Active	PROD
	10.1.1.1	255.255.255.0	Active	LAPASTORA1
	10.1.1.2	255.255.255.0	Active	LABNET2
	10.1.1.3	255.255.255.0	Active	*NONE
	10.1.1.4	255.255.255.0	Active	*NONE
	10.1.1.50	255.255.255.255	Inactive	SHEILA
	127.0.0.1	255.0.0.0	Active	LOCALHOST

Bottom

F3=Exit F5=Refresh F6=Print list F11=Display line information
F12=Cancel F17=Top F18=Bottom

Figure 10-3 Work with TCP/IP interfaces showing ALIAS NAME

Control Language (CL) commands (for example, Start TCP/IP Interface STRTCPIFC) has a new parameter (ALIASNAME) that can be used instead of the IP address. With this enhancement you can start TCP interfaces in two ways:

```
STRTCPIFC INTNETADR(10.1.1.1)
```

And:

```
STRTCPIFC ALIASNAME(LAPASTORA1)
```

Running these commands on the command line accomplishes the same when done on the Work with TCP/IP Interfaces screen shown in Figure 10-3.

Similarly, you can use the iSeries Navigator to view or change this property, as shown in Figure 10-4.

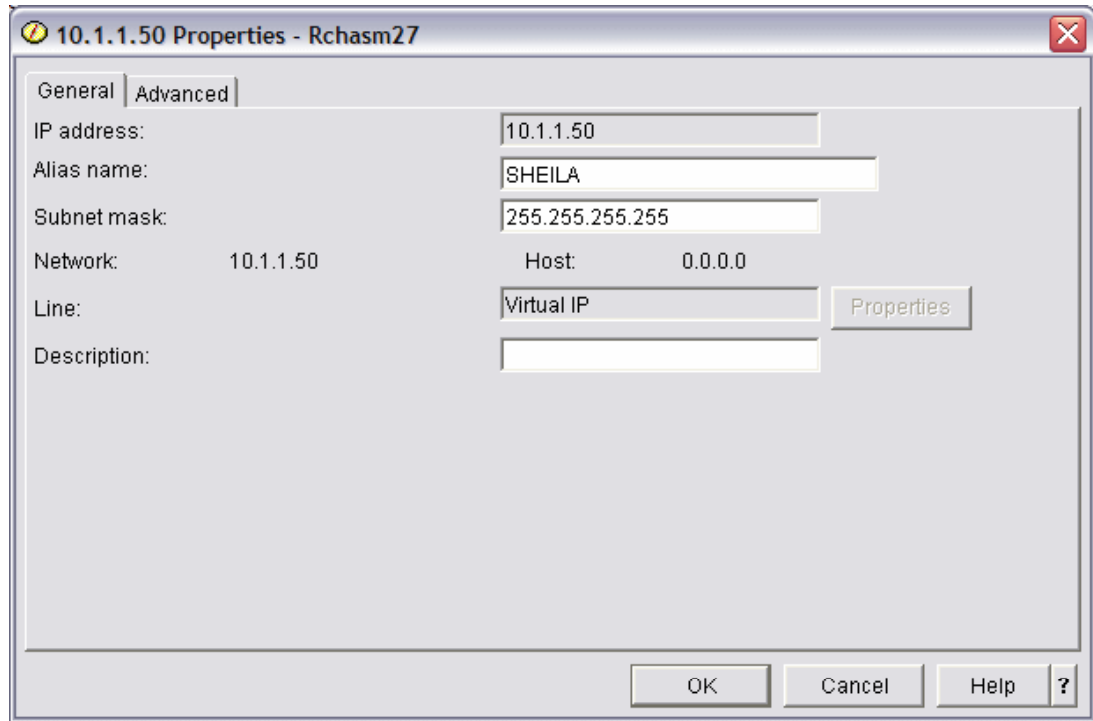


Figure 10-4 Interface properties as viewed in iSeries Navigator

With the IPv6 unique address format, this feature will come in handy for both system and network administrators. For details on i5/OS V5R4 IPv6 address format refer to:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzai2/rzai2ip6addrformat.htm>

10.4 Virtual IP and Virtual Ethernet

In i5/OS V5R4 you can now create a preferred interface list to select which adapters and IP addresses will be the preferred interface for the virtual IP address proxy Address Resolution Protocol (ARP) agent selection. This is available for both virtual IP addresses and virtual Ethernet.

10.4.1 Virtual IP enhancements

Starting with i5/OS V5R4, you can manually select which adapters and IP addresses are to be the preferred interface for virtual IP address (VIPA) proxy Address Resolution Protocol (ARP) agent selection. You can select which interface to use by creating a preferred interface list if an adapter failure occurs. A preferred interface list is an ordered list of the interface addresses that will take over for the failed adapters. You can use either iSeries Navigator (via the interface properties) or the Change TCP/IP IPv4 Interface (QTOCC4IF) application programming interface (API) to configure a preferred interface list. The preferred interface list is also configurable for both virtual Ethernet and virtual IP address interfaces.

A manual selection mechanism was added to give you control over which adapters and IP addresses will be the preferred interface for virtual Ethernet and VIPA proxy ARP agent

selection. You can now create a preferred interface list to select the adapters and IP addresses for the ARP proxy agent. The preferred interface list is configurable for both virtual Ethernet and virtual IP address interfaces.

Figure 10-5 shows the iSeries Navigator's interface properties window and how it relates to VIPA proxy ARP agent configuration.

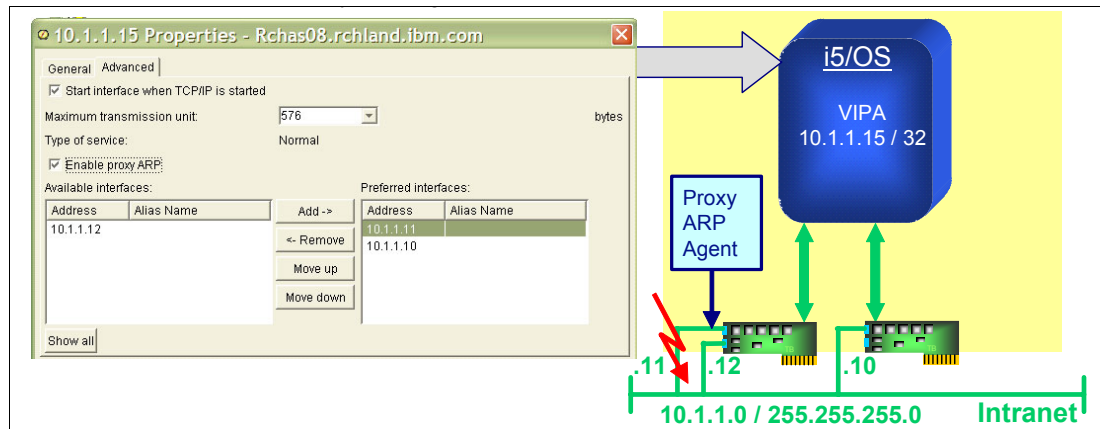


Figure 10-5 Virtual IP and proxy ARP

If the first IP interface fails, the proxy ARP agent will start listening on the next interface in the list. As shown in Figure 10-5, a VIPA (10.1.1.15) is configured in the i5/OS partition. Three physical interfaces (10.1.1.10, 10.1.1.11, 10.1.1.12) provide access to the system. The proxy ARP agent starts responding on interface 10.1.1.11, because this is the first interface in the preferred interface list. When this interface fails, the agent will start listening on the second interface 10.1.1.10. If, for example, the second interface would fail, too, the system will no longer answer any ARP requests for the VIPA. Even though there is another interface (10.1.1.12) in the same subnet, the system will not use this interface, because it was not defined in the preferred interface list. Once a preferred interface list is used, the system will only use the interfaces for proxy ARP selection that are defined in the list.

Failover using a preferred interface list

You can create a preferred interface list to control the order in which the local interfaces are used when an adapter failure occurs.

To create a preferred interface list, follow these steps:

1. In iSeries Navigator, expand **Network** → **TCP/IP Configuration** → **IPv4**.
2. Click **Interfaces**.
3. From the lists of interfaces that are displayed, select an interface for the virtual IP address or virtual Ethernet for which you want to create the preferred interface list.
Using Figure 10-5 as an example, you will select the virtual IP address 10.1.1.15.
4. Right-click the interface, and then select **Properties**.
5. Click the **Advanced** tab.
6. In the panel, select the interface addresses from the available interface list, and click **Add**.

Using Figure 10-5 as an example, you will select interfaces 10.1.1.11 and 10.1.1.10 and add them to the preferred interface list one by one.

You can also remove an interface from the preferred interface list in the right pane by using the **Remove** button, or move an interface up and down the list to change the order by using the **Move up** and **Move down** buttons.

7. Select the **Enable proxy ARP** check box above the Available interfaces list to enable the list.
8. Click **OK** to save the preferred interface list that you have just created.

Attention: You can only include 10 interfaces in the preferred interface list. If you configure more than 10, the list is truncated to the first 10.

The interface for which you want to create the preferred interface list must be inactive for the list to be configured. Interfaces listed in the preferred interface list do not need to be inactive at the time the list is configured.

10.4.2 Virtual Ethernet and the Multiple Proxy ARP support

Similar to the scenario with a virtual IP address (VIPA), a preferred interface list can also be used with IP interfaces that are defined for a virtual Ethernet line. Prior to V5R4, when a system was partitioned and you wanted to provide access to the various partitions through a single i5/OS partition, transparent subnetting could be used to provide the connection from the external LAN to the internal partitions via virtual Ethernet.

The proxy ARP or transparent subnetting method requires that you have a firm understanding of subnetting and TCP/IP. You need to obtain a contiguous block of IP addresses that are routable by your network. You subnet this block of IP addresses.

As an example (refer to Figure 10-6) a contiguous block of eight IP addresses (10.1.1.32 through 10.1.1.39) is used. Since it is a block of eight IP addresses, the subnet mask for these addresses is 255.255.255.248. You assign one to each of the virtual TCP/IP interfaces on your partitions. You cannot use the address 10.1.1.32 (is the network itself) and 10.1.1.39 (broadcast address for network 10.1.1.32).

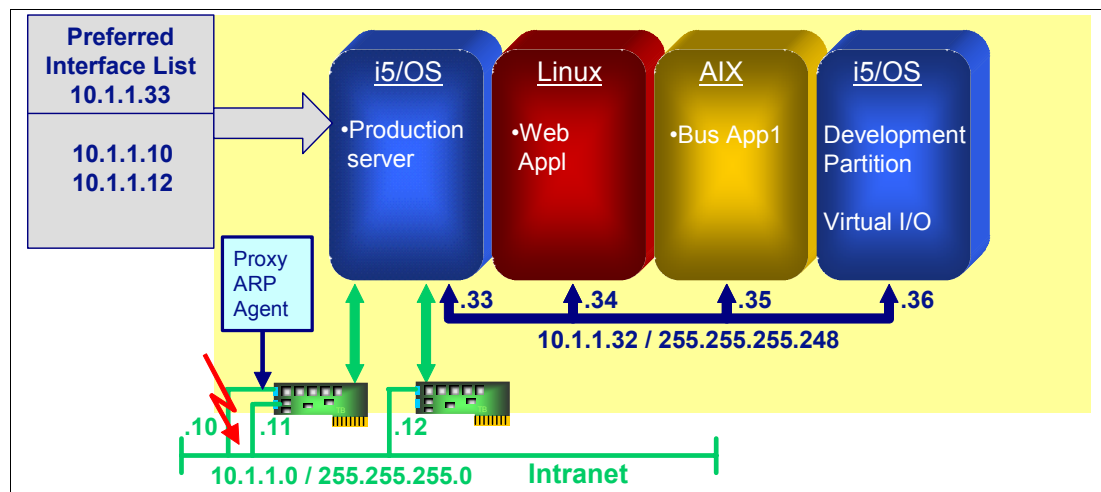


Figure 10-6 Virtual Ethernet with transparent subnetting

Prior to V5R4, transparent subnetting (proxy ARP) was accomplished by associating the external IP interface with the IP interface of the virtual Ethernet connection. If, for example, the external interface went down, the system would not answer any ARP requests for the IP interfaces on the virtual Ethernet anymore, even though other interfaces to the external

network were available. With V5R4, instead of defining the associated local interface on the virtual Ethernet interface, you can use the preferred interface list. When using the list, the proxy ARP agent will answer ARP requests for the virtual Ethernet IP interfaces in the defined order.

10.4.3 Proxy ARP support for a Virtual IP address on a different subnet

In OS/400 V5R2 and i5/OS V5R3, proxy ARP is supported for virtual IP addresses that are in the same subnet as the IP interface addresses on the physical LAN adapters. For instance (see Figure 10-7 on page 233), when configured for proxy ARP, an Ethernet IP interface 10.10.10.10/24 would answer ARP requests for a VIPA 10.10.10.20 /24, but it would not answer ARP requests for 10.20.20.5 /24.

With i5/OS V5R4, support was added to provide proxy ARP support for VIPAs that are not in the same subnet as local interfaces. This allows, for example, an Ethernet IP interface of 10.10.10.10/24 to answer ARP requests for a VIPA 10.20.20.5 /32.

This enhancement can be leveraged by a technology that Cisco introduced for mobility in local area networks. The feature that is only supported by Cisco routers is called Local Area Mobility (LAM) and is part of Cisco's router operating system IOS.

Consider the diagram shown in Figure 10-7 on page 233. Assume that the iSeries system with its VIPA 10.10.10.75 /32 used to be installed in network A. When a client in network A (10.10.10.0 /24) wanted to communicate with the iSeries system, which was on the same subnet, the client would issue an address resolution protocol (ARP) request to obtain the MAC address for the iSeries interface. At a certain point in time it was necessary to move the iSeries system to a different building. This building had an IP subnet of 10.20.20.0 /24 (network B). However, many applications and clients had hard-coded the iSeries IP address 10.10.10.75, so a quick change of the iSeries IP address was not possible. The answer to this problem is LAM. LAM allows you to implement a migration scenario like the one described above. When LAM is used, the router (LAM router) to which the iSeries system is now connected will be enabled for proxy ARP and LAM (mobility) on the LAN interface. This allows the router to answer ARP requests for addresses in the 10.10.10.0/24 network. In addition, route redistribution has to be enabled on this router using one of the interior gateway protocols (IGPs), such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP). The host route for 10.10.10.75 is then distributed to other routers in the network. On the original router in network A, proxy ARP has also to be enabled. The original router's routing table now contains an entry for network 10.10.10.0 /24 for its LAN interface and a more specific host route to 10.10.10.75 on the external interface. When a client now wants to communicate with 10.10.10.75 (VIPA), the original router answers the ARP request, thus pretending to be the interface the client wants to talk to and routes the packets to the LAM router, which in turn sends the packet to the network B. In order for the LAM router to send a packet to the iSeries, the router first sends an ARP request for 10.10.10.75 to the local network. This is where the V5R4 enhancement comes into the picture. The LAN interface on the iSeries will now answer the ARP request for the VIPA even though the VIPA is in a different subnet than the LAN IP interface (10.20.20.30).

Figure 10-7 shows the diagram for the example in the previous paragraph.

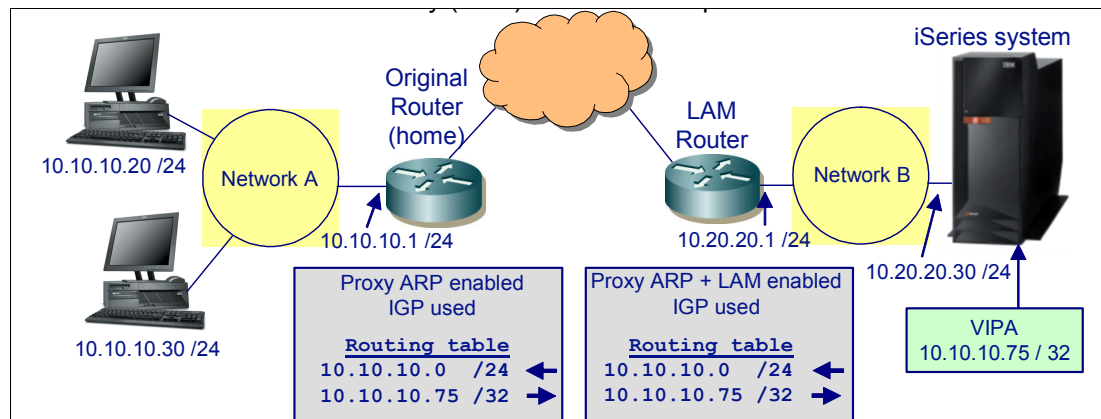


Figure 10-7 Proxy ARP support for Virtual IP address on different subnet

For more information about Cisco's LAM support, search for Local Area Mobility at the Cisco Web site at:

<http://www.cisco.com>

10.5 Virtual Private Networking and the i5/OS V5R4

In V5R2 support was added to the iSeries to act as a client for User Datagram Protocol (UDP) encapsulation. This means that an iSeries system can act as an initiator of a VPN connection that used UDP encapsulation, but it cannot process incoming VPN connection requests with UDP encapsulation.

Also in V5R2, support to VPN was added to allow a system to initiate VPN connections that traverse network address translation firewalls. Network Address Translation (NAT) allows you to hide your unregistered private IP addresses behind a set of registered IP addresses. This helps to protect your internal network from outside networks. NAT also helps to alleviate the IP address depletion problem, since many private addresses can be represented by a small set of registered addresses. Unfortunately, conventional NAT does not work on IPsec packets because when the packet goes through a NAT device, the source address in the packet changes, thereby invalidating the packet. When this happens, the receiving end of the VPN connection discards the packet and the VPN connection negotiations fail.

In V5R4, server-side NAT traversal support has been added, and with this enhancement, i5/OS has a complete solution for establishing VPNs through NAT firewalls. Here UDP encapsulation support is added on the responder side of a VPN connection, thus providing full UDP encapsulation support for i5/OS VPN. A graphical view of this end-to-end connection is shown in Figure 10-8.

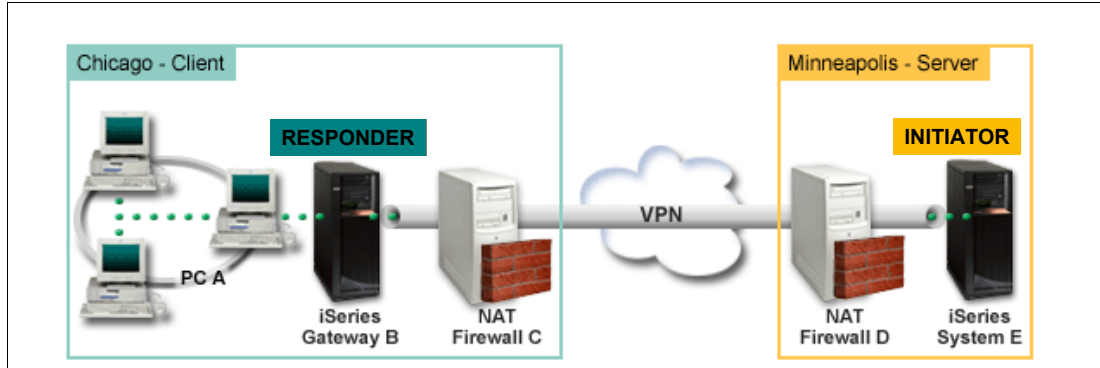


Figure 10-8 Server side NAT traversal support

This UDP encapsulation (also known as NAT friendly IPsec or NAT traversal support) allows IPsec traffic to pass through a conventional NAT device.

UDP encapsulation wraps an IPsec packet inside a new, but duplicate, IP/UDP header. The address in the new IP header gets translated when it goes through the NAT device. Then when the packet reaches its destination, the receiving end strips off the additional header, leaving the original IPsec packet, which will now pass all other validations.

You can only apply UDP encapsulation to VPNs that will use IPsec ESP in either tunnel mode or transport mode.

Once the packet is encapsulated, the iSeries sends the packet to its VPN partner over UDP port 4500. Typically, VPN partners perform IKE negotiations over UDP port 500. However, when IKE detects NAT during key negotiation, subsequent IKE packets are sent over source port 4500, destination port 4500. This also means that port 4500 must be unrestricted in any applicable filter rules. The receiving end of the connection can determine whether the packet is an IKE packet or a UDP encapsulated packet because the first 4 bytes of the UDP payload are set to zero on an IKE packet.

Another addition to VPN in i5/OS V5R4 is the Traffic Flow Confidentiality (TFC) and the Extended Sequence Number (ESN) support.

10.5.1 Traffic Flow Confidentiality

Traffic Flow Confidentiality (TFC) allows you to conceal the actual length of the data packets transferred over the VPN connection. Use TFC for extra security against attackers who might guess the type of data being sent over the VPN connection from the length of the data packet. You can only use TFC if your data policy is set up for Tunnel Mode.

If your data policy is configured for tunnel mode you can use Traffic Flow Confidentiality (TFC) to conceal the actual length of the data packets transferred over a VPN connection.

TFC adds extra padding to the packets being sent and sends dummy packets with different lengths at random intervals to conceal the actual length of the packets. When you enable TFC you gain more security, but at the cost of system performance. Therefore, you should test

your system's performance before and after you enable TFC on a VPN connection. TFC is not negotiated by IKE, and the user should only enable TFC when both systems support it.

TFC is configured per VPN connection in the VPN connection properties on the General tab (see Figure 10-9).

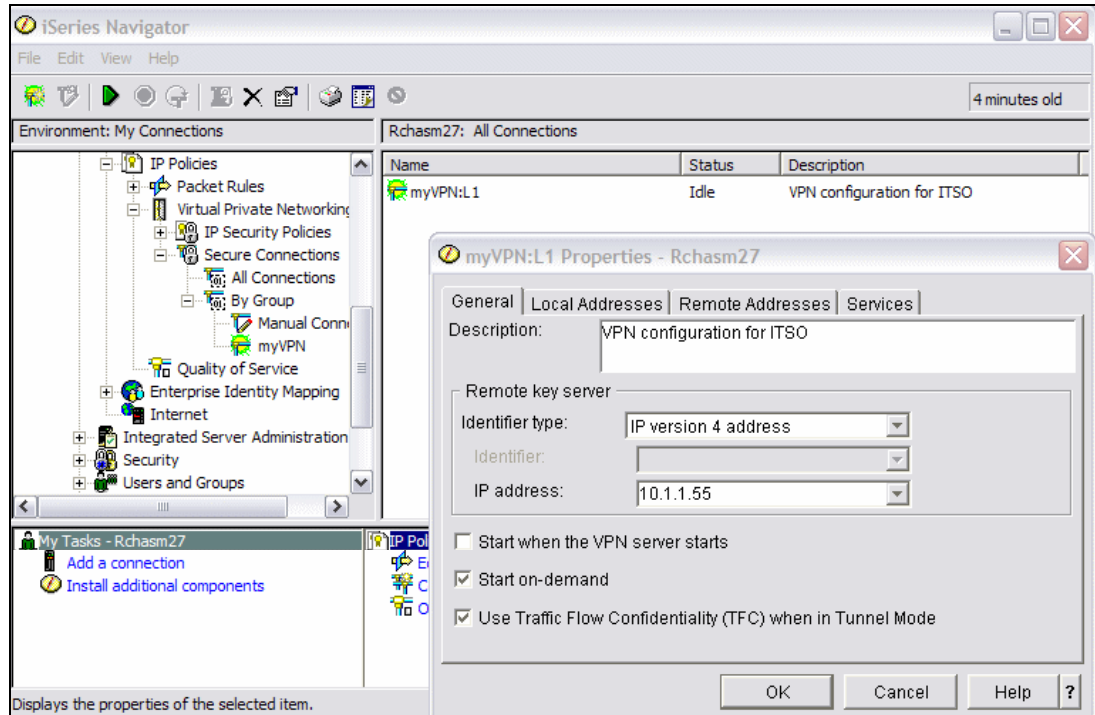


Figure 10-9 TFC configuration for VPN

10.5.2 Extended Sequence Number (ESN)

The IP Security Authentication Header (AH) and Encapsulating Security Payload (ESP) protocols use a sequence number to detect replay of packets.

ESN allows your VPN connection to transmit large volumes of data at a high speed without having to frequently re-key. You can only enable ESN when your data policy uses the Authentication Header (AH) protocol or the Encapsulation Security Payload (ESP) protocol and AES as the encryption algorithm.

The VPN connection uses a 64-bit sequence number instead of 32-bit numbers over IPsec. Using 64-bit sequence numbers allows more time before re-keying, which prevents sequence number exhaustion and minimizes the use of system resources.

ESN is a standard that is described in RFC4304.

ESN is turned on as a global value in the Virtual Private Networking properties on the General tab (refer to Figure 10-10).

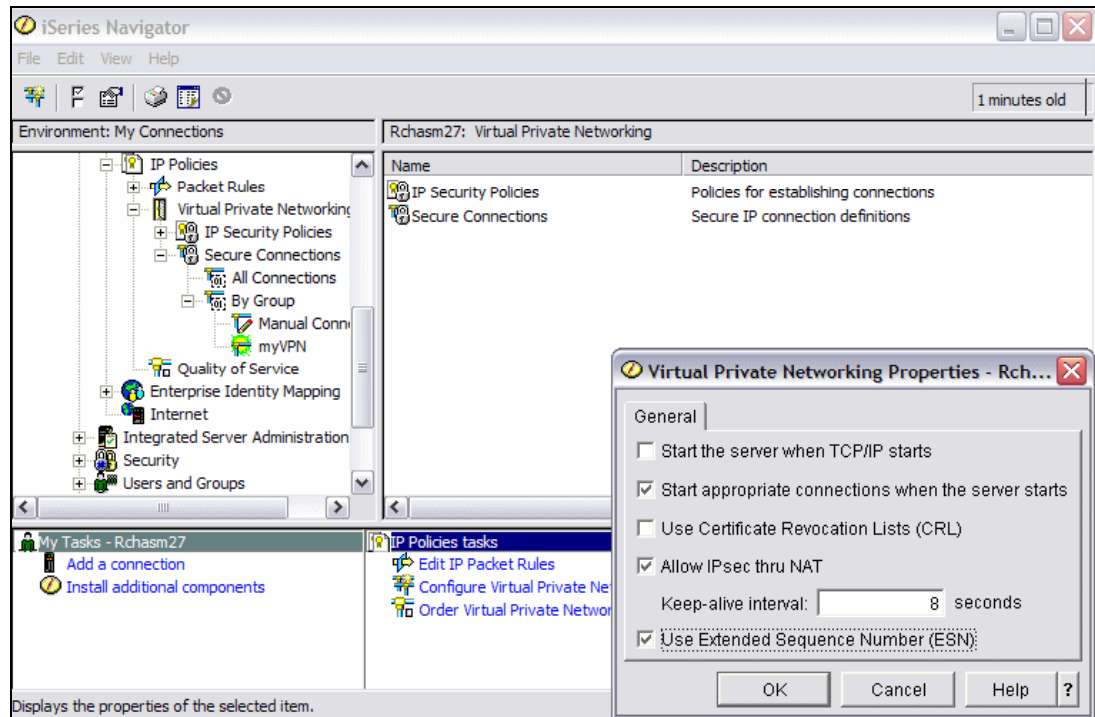


Figure 10-10 ESN configuration in VPN

10.6 Point-to-Point Protocol and Layer 2 Tunnel Protocol

i5/OS V5R4 also offers changes to the Point-to-Point Protocol (PPP) and the Layer 2 Tunnel Protocol (L2TP). One such change is that the QTPPPSSN and the QTPPPL2SSN (L2TP) jobs are no longer used in V5R4. Prior to V5R4, these jobs were started and ended with the STRTCPPTP and ENDTCPPTP commands or by QTPPPL2TP when a tunnel was established or ended. They can also be started or ended automatically as links were started or ended by the multilink protocol.

PPP no longer uses these QTPPPSSN and QTPPPL2SSn jobs in V5R4. Sessions run as threads in QTPPPCTL.

In releases prior to i5/OS V5R4, Work with Point-to-Point TCP/IP Profiles (WRKTCPPTP) command option 14 (Work with job) brought up the active session job. Occasionally it will start QTPPPL2TP if there was no active PPP session for the L2TP profile.

In V5R4, WRKTCPPTP option 14 brings up QTPPPCTL if a session thread is active in that job.

In addition to these changes, there is a new message log and call log spooled file for session messages. It collects messages from the session thread, messages from the initial thread that are the result of work on behalf of the session, and messages from spawned processes into one spooled file.

Since a PPP session is run in a thread, the message log and call log spooled file for this session are named MLmmmmmmmm and CLmmmmmmmm, respectively, where

mmmmmmmm is the thread ID. This allows the easy matching of call logs, message logs, and session messages in the QTPPPCTL job log.

Another notable change is that the list of Ethernet adapters that support Point-to-Point Protocol over Ethernet (PPPoE) is expanded to include type 2743, 2760, 2838, 2849, 287F, 5700, 5701, 5706, 5707, and 573A Ethernet adapters. And not only this, but PPPoE can now share the same adapter as IPv4 and IPv6 traffic.

10.7 HTTP server enhancements

In i5/OS V5R4, porting open source modules to i5/OS is now easier than before. However, the already existing modules must be recompiled. To set up third-party modules for HTTP Server (powered by Apache), see the iSeries Information Center V5R4 for details at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzaie/rzaiethrdprtymod.htm>

Also in V5R4, logs are now tagged with CCSID UTF8 (1208). All messages are encoded in UTF8 regardless of where they are contained (custom, error, FRCA, and script logs). Saving logs that are currently active is now possible, causing no error messages to be sent to the job log.

In addition, data is maintained in its original encoding until the CGI job is ready to be launched. This is a new enhancement in the CGI job handling function by avoiding unnecessary conversions until the CGI CCSIDs and conversion mode are determined specific to the configuration.

More on CGI enhancements, new and existing directives now allow CGI jobs to run in more than one language from a single Apache server. The CGIJobCCSID directive sets the CGI job CCSID, which can differ from the server job CCSID. The CGIJobLocale directive sets the locale environment variable for the CGI job to retrieve. The DefaultNetCCSID directive is now allowed within a container. The combined use of these directives in a directory or virtual host container allows one CGI program to support language A, while in another container a CGI program can support language B.

In V5R4 support to activate the Enterprise Workload Management (eWLM) instrumentation module for HTTP Server (powered by Apache) has been added and can be configured using the IBM Web Administration for i5/OS interface. Application Response Measurement (ARM) 4.0 APIs are used to classify requests and record the time spent for each one.

For more details on HTTP Server enhancements, see the iSeries Information Center V5R4 at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzaie/rzaiemain.htm>

10.8 iSeries NetServer changes

i5/OS support for Windows Network Neighborhood (iSeries NetServer) is an IBM i5/OS function that enables Windows 2000, Windows XP, and Windows Server 2003 clients to access i5/OS shared directory paths and shared output queues. Windows clients on a network utilize the file and print sharing functions that are included in their operating systems. You do not need to install any additional software on your PC to use and benefit from iSeries NetServer.

NetServer also supports Linux access to the Integrated File System:

- ▶ Linux PC can mount NetServer Share.
- ▶ Supports Linux kernel Version 2.4.4+.
- ▶ Supports Samba Versions 2.0.7+ or 2.2+.

Linux workstation command access to IFS via:

- ▶ `smbmount` – mounts a share to a local directory
- ▶ `smbclient` – like an FTP client

Linux workstation Graphical User Interface (GUI) access to IFS via:

- ▶ Konqueror – file browser

In V5R4 changes have been made to the iSeries NetServer code to support the following:

- ▶ Multi-threaded support

iSeries NetServer now processes client requests in a single, multi-threaded job (QZLSFILET) in the QSERVER subsystem. The new iSeries NetServer configuration now allows for a QZLSFILET job in other subsystems. QZLSFILET is the new multithreading job in V5R4.

- Scalability and performance changes.
- Reduce contention on key table/lock.
- Threading support enables use of techniques to remove two task switches for requests.
- Opportunistic lock support added, which allows clients to cache data.

- ▶ Message authentication

Message authentication allows the Common Internet File System (CIFS) protocol to be protected against various attacks: man in the middle, hijacking, replay

- ▶ Windows-style messages

iSeries NetServer has supported Windows-style pop-up messages, but the configuration for this support is now in the APIs and Navigator

- ▶ Instrumented i5/OS support for Start Trace (STRTRC) command

Greatly improved tracing of the i5/OS support using the STRTRC command

- `TRCTYPE(*FILESVR)` for the file serving component
- `TRCTYPE(*NETSVR)` for the NetServer specific component

More on multi-threaded support

Threads cannot access non-thread safe file systems, which are QDLS, QNetWare, and QSYS.LIB usage of save files and device descriptions. For file systems that are not thread safe, NetServer will launch the session into a QZLSFILE job if it is the first drive mapped by the session.

Access-denied errors occur when using non-thread safe file systems or functions from a threaded job. Examples of this include:

- ▶ Map a drive to the root file system (thread safe).
- ▶ Drill into QDLS (not-thread safe) results in access denied error.
- ▶ Map a drive to share for path `/dir1/dir2`.
- ▶ In `/dir1/dir2` is `qdlslnk`, which is a symbolic link to `/qdl`.

Drilling into `qdlslnk` will get an access-denied error.

Subsystem configuration is recommended to alternatively segregate thread safe/non-thread safe clients.

Scalability and performance enhancements

In i5/OS V5R4, internal changes to the iSeries Navigator have been made for scalability and performance.

Comparison for response time and throughput have been made between the previous version and the i5/OS V5R4 iSeries Navigator and the V5R4 version has proved to be better in both criteria. Figure 10-11 shows the performance charts comparing the V5R4 iSeries NetServer with the previous releases.

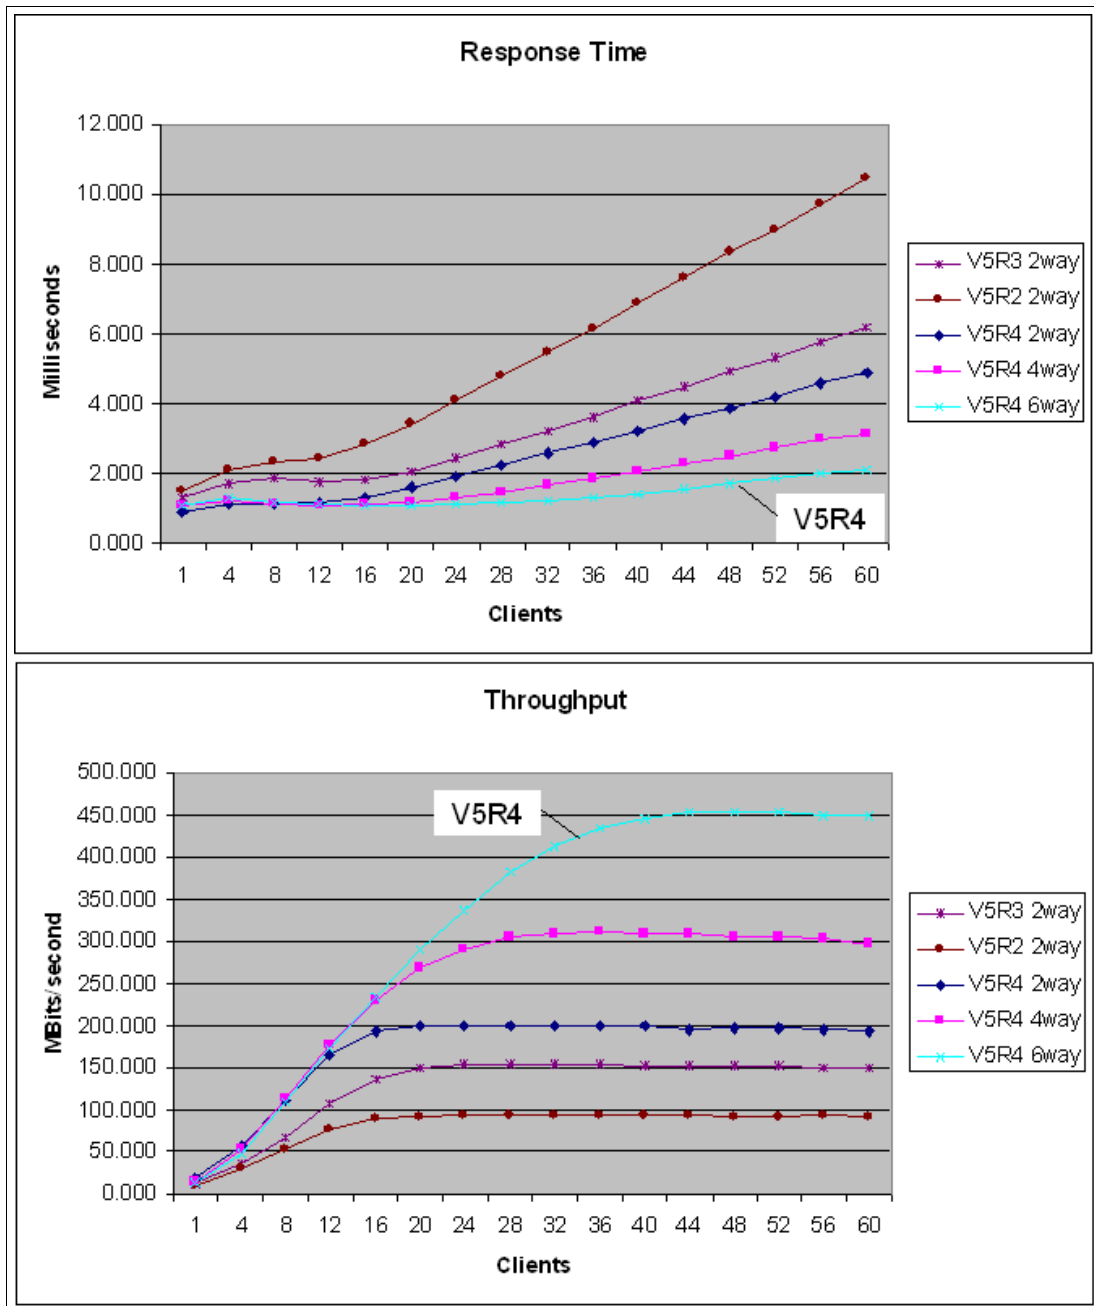


Figure 10-11 iSeries NetServer performance comparison

10.9 iSeries Access family in V5R4

The iSeries Access family of products is the solution for PC, Web browser, or wireless connectivity to your server. With the iSeries Access, you can access and administer your servers from the various iSeries Access interfaces.

Packaging

iSeries Access is packaged with the licensed program 5722-XW1 together with the licensed programs 5722-XE1 (for Windows), 5722-XH2 (for Web), 5722-XL1 (Linux), and 5722-XP1 (Wireless). The chart in Figure 10-12 shows the iSeries Access packaging since V5R2.

V5R4 5722-XW1 iSeries Access Family	V5R3 5722-XW1 iSeries Access Family	V5R2 5722-XW1 iSeries Access Family
Series Access for Windows, 5722-XE1, V5R4	Series Access for Windows, 5722-XE1, V5R3	Series Access for Windows, 5722-XE1, V5R2
Series Access for Web, 5722-XH2, V5R4	Series Access for Web, 5722-XH2, V5R3	Series Access for Web, 5722-XH2, V5R2
		WebSphere Host Publisher, 5724-B81, V4.0, 5724-B81, V4.01
	HATS Limited Edition V5.0, 5724-F97-01	HATS Limited Edition V4.0 5724-D34-01
Series Access for Linux, 5722-XL1	Series Access for Linux, 5722-XL1, V1.10	Series Access for Linux, 5722-XL1, V1.0
Series Access for Wireless, 5722-XP1, V5R4	Series Access for Wireless, 5722-XP1, V5R3	Series Access for Wireless, 5722-XP1, V5R2
	V5R3 customers not wanting to upgrade to i5/OS V5R4 but want the new V5R4 iSeries Access Family clients can order no-charge Feature No. 2648 of Product No. 5722-XW1	V5R2 customers not wanting to upgrade to i5/OS V5R3 but want the new V5R3 iSeries Access Family clients can order no-charge Feature No. 2647 of Product No. 5722-XW1

Figure 10-12 iSeries Access family packaging

Note that HATS Limited Edition cannot be ordered new as of May 2006.

10.9.1 iSeries Access for Windows

In i5/OS V5R4 iSeries Access for Windows runs on Windows 2000, Windows XP Professional, Windows XP Tablet PCs, and Windows 2003 Server and supports 64-bit Processors like Advanced Micro Devices (AMD64), Intel with EM64T0, and the Intel with Itanium®.

For detailed information about what functions support 64-bit, refer to:

<http://www.ibm.com/eserver/iseries/access/supportedos.htm>

iSeries Access for Windows has many components, each of which may have its own enhancements. The following sections are organized according to the various components or sub components. Note that there is no specific section on iSeries Navigator under iSeries Access for Windows. iSeries Navigator is a large *superset component* under iSeries Access for Windows.

Enhancements to the iSeries Navigator functions are covered in other chapters in this book under the respective support area, such as database, communications, security, work management, Backup Recovery and Media Services (BRMS), and so forth.

iSeries ODBC driver

The V5R4 ODBC driver of iSeries Access is compliant with Microsoft ODBC Version 3 specification. Workstation applications can now take advantage of new functions included in the 3.5 specification.

The chart in Figure 10-13 shows the V5R4 ODBC enhancements as well as the enhancements in V5R3 and V5R2.

ODBC Enhancements		
V5R4 Enhancements	V5R3 Enhancements	V5R2 Enhancements
<p>Requires i5/OS V5R4</p> <ul style="list-style-type: none"> ▪ 128-byte column names ▪ Maximum SQL statement lengths of 2,097,152 bytes or 1,048,576 characters ▪ Support for IBM Enterprise Workload Manager (eWLM) correlator ▪ Support for lock sharing between loosely coupled transactions <p>OS/400 V5R2 or later</p> <ul style="list-style-type: none"> ▪ Improved support for delimited names 	<ul style="list-style-type: none"> ▪ Support for BINARY / VARBINARY data types ▪ Support for UTF-8 / UTF-16 data ▪ Support for increased precision of decimal numbers ▪ Enhanced MTS support ▪ ANSI / ISO (American National Standards Institute / International Standards Organization) Core Level SQL standard of 1999 	<ul style="list-style-type: none"> ▪ 64K SQL Statements ▪ Additional descriptor information ▪ MTS Support (actually shipped in V5R1)

Figure 10-13 ODBC enhancements in V5R4

iSeries OLE DB provider

The iSeries OLE DB Provider in V5R4 supports OLE DB 2.5. Being the Microsoft-recommended access for 32-bit Windows applications, it can give you record-level access and access to data queues and stored procedures, and can be used to run SQL, remote commands, and ODBC.

Figure 10-14 shows the chart of the enhancements in the iSeries OLE DB Provider in V5R4.

OLE DB Enhancements		
V5R4 Enhancements	V5R3 Enhancements	V5R2 Enhancements
<p>Requires i5/OS V5R4</p> <ul style="list-style-type: none"> 128 byte column names Maximum SQL statement lengths of 2,097,152 bytes or 1,048,576 characters Support for IBM Enterprise Workload Manager (eWLM) correlator <p>OS/400 V5R2 or later</p> <ul style="list-style-type: none"> System naming and library list Improved support for delimited names 	<ul style="list-style-type: none"> New SQL-only provider (IBMDASQL) <ul style="list-style-type: none"> SQL commitment control using IBMDASQL MTS support using IBMDASQL SQL <ul style="list-style-type: none"> Custom blocking in SQL SQL data compression SQL package support New Record-Level Access-only provider (IBMDARLA) <ul style="list-style-type: none"> Record-level access support for forward-only cursors and blocked reads using IBMDARLA Supports updatable cursors for the SQL dialect Database BINARY and VARBINARY data types Database larger decimal precision support Unicode support <ul style="list-style-type: none"> UTF-8 & UTF-16 support 	<ul style="list-style-type: none"> Custom properties added <ul style="list-style-type: none"> Force Translate (translate CCSID 65535 data) Default Collection Catalog Library List Convert Date Time To Char Supports updatable cursors for the SQL dialect Supports ROW-ID 64K SQL Statements Additional descriptor information Is thread safe OLE DB provider work with Windows products, such as: <ul style="list-style-type: none"> Visual Basic 6.0 OLE DB controls and wizards ADO 2.5 – ships with Microsoft Windows 2000 ADO 2.1 - ships with Microsoft IE 5.0 and Office/2000 products An OLE DB interface is provided to support ADO recordset Seek (ADO 2.x)

Figure 10-14 OLE DB provider enhancements

Visual Basic programming

iSeries Access in V5R4 provides plug-ins for Visual Basic. Additionally, it provides wizards to aid in developing Visual Basic applications, sample programs for IE, PowerBuilder, Delphi, Visual C++®, and Lotus 1-2-3®. Figure 10-15 shows Microsoft Visual Basic with Express Toolkit.

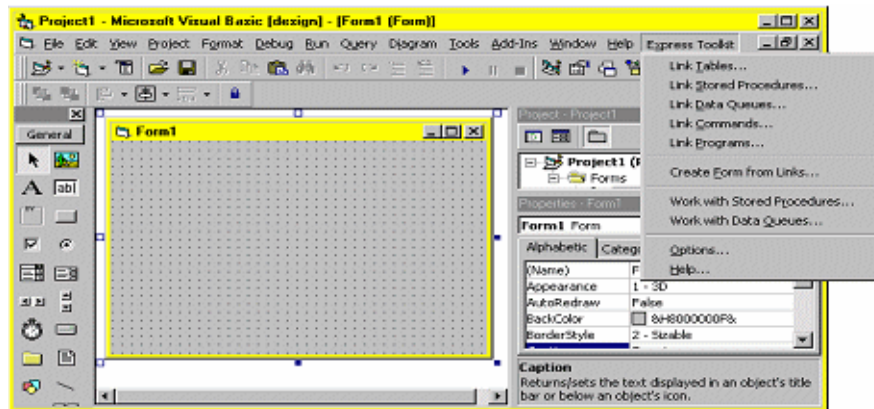


Figure 10-15 Visual Basic support

.NET support

In i5/OS V5R4 it is now possible in .NET to have 128-byte column names and SQL statements with maximum lengths of 2,097,152 bytes or 1,048,576 characters. It now has support for IBM Enterprise Workload Manager (eWLM) correlator.

.NET Enhancements	
V5R4 Enhancements	V5R3 Enhancements
<ul style="list-style-type: none">▪ System naming and library list support¹▪ LOB data type columns support¹▪ MS firewall 2.0 compatibility¹▪ Multiple active result sets per connection▪ Customizable String processing for Char for Bit Data, Date, Time, Timestamp, Decimal, and Numeric data types▪ Support for IntelliSense™▪ Additional sample programs▪ Improved support for delimited names <p>Requires i5/OS V5R4</p> <ul style="list-style-type: none">▪ 128 byte column names▪ Maximum SQL statement lengths of 2,097,152 bytes or 1,048,576 characters▪ Support for IBM Enterprise Workload Manager (eWLM) correlator <p>¹ Available with latest V5R3 service pack</p>	<ul style="list-style-type: none">▪ SQL (INSERT,UPDATE, DELETE)▪ Commitment Control▪ Connection Pooling▪ SQL naming▪ Unicode▪ Tracing▪ Threads▪ IASPs (multiple databases)▪ Stored Procedure Support▪ iSeries-specific Properties▪ User-Defined Types <p>Supported later via SP SI15176</p> <ul style="list-style-type: none">▪ System Naming (/)▪ Library List▪ Large Objects (LOBs)

The .NET provider is named IBM.Data.DB2.iSeries. It allows applications using Microsoft's .NET framework to access DB2 UDB for iSeries databases. For complete documentation of the .NET Data Provider, see IBM DB2 UDB for iSeries .NET Provider Technical Reference.

Figure 10-16 .NET enhancements

For more information about .Net support in this book refer to 4.3.1, “iSeries .NET provider” on page 86, in the Database chapter.

Data transfer

Data Transfer in V5R4 iSeries access now supports MS Excel® XML Spreadsheet format and has improved support for delimited names. It now supports 128-byte column names as well as PC selection of an independent auxiliary storage pool (IASP). It now supports creating and overwriting of empty query result sets.

10.9.2 iSeries Access for Linux

iSeries Access for Linux includes a full function 5250 emulator, ODBC driver, and iSeries Access APIs (in support of various functions like RmtCmd, NLS, System Object, and so on). This is product can be used on Linux operating systems with Intel processors and on Power PCs or iSeries server logical partitions (LPARs).

iSeries Access for Linux comes in two versions:

- ▶ 32-bit Linux operating systems
 - Available August 2, 2005: Version 1.10.
 - The 32-bit version provides an ODBC driver for accessing the DB2 Universal Database™ (UDB) for iSeries and a 5250 emulator.
 - Version 1.10 includes new support for Kerberos and single sign-on (SSO), bypass sign-on, and an example SSL configuration.

- ▶ 64-bit Linux operating systems
 - Available February 21, 2005: Version 1.2.
 - The 64-bit version provides an Extended Dynamic Remote SQL (EDRS) driver for Power PC®.
 - It is supported only on SuSE SLES 9 and requires the iSeries Access for Linux 32-bit product to be installed first.
 - For more information about ERDS, refer to the XDA Web Site and iSeries Infocenter, Extended Dynamic Remote SQL (EDRS) APIs in the **Programming** → **APIs** → **Database and File** → **Database** section.

The iSeries Access for Linux 64-bit package has been enhanced to include a 64-bit ODBC driver. This driver is functionally equivalent to the 64-bit ODBC driver currently included in iSeries Access for Windows. The 64-bit ODBC driver will be included in the Web download in 1H06.



i5/OS-based consoles and logical partitioning

This chapter contains the V5R4 updates for supported consoles and logical partitioning.

In this chapter we discuss the new features added for the supported consoles (that is, an HMC with 5250 emulation, ASMI, Operations Console with 5250 emulation, 5250 twinax), as well as the new with V5R4 Thin Console for System i5 models.

We also:

- ▶ Summarize non-HMC-managed logical partition support available prior to i5/OS V5R4.
- ▶ Cover the latest tool updates to simplify logical partitioning ordering and configuration.
- ▶ Review information about how Reserve Capacity on Demand for processors works.

11.1 V5R4 console support

With V5R4 you have the first opportunity to use the two port Ethernet LAN adapter features 5706 and 5707 for Operations Console (LAN) and the corresponding embedded two-port Ethernet LAN adapter in the IBM System i5 520, 550, and 570 models.

When a system is managed by a Hardware Management Console (HMC), you specify which 5250 console type you want to use as the main console to an i5/OS partition through the HMC interface.

When your system is not managed by the HMC, you have ordered the system to be managed by either a twinaxial attached 5250 workstation or Operations Console (Direct attached), or Operations Console (LAN attached). When shipped new, there are default algorithms used by the system to find the card or port to be used as the console. These rules should be understood.

Existing console options also remain available:

- ▶ Twinaxial console (IOP-based)
- ▶ Operations Console (direct) attached to an asynchronous port of WAN adapter (IOP-based)
- ▶ Operations Console (LAN) to LAN adapter (IOP-based)
- ▶ 5250 emulation interface via HMC

The Thin Console, made by Neoware, Inc. for IBM, joined the available 5250 consoles customers can choose from. It is intended for usage with small and intermediately configured System i5 520 and 550 models.

11.1.1 Operations Console support

Operations Console support, requiring a PC workstation running a Windows operating system, has been available for several years, originally supported over Direct attachment, IBM Token Ring, or TCP/IP-based Ethernet LAN connections.

With V5R4 you have the first opportunity to use the two-port Ethernet LAN adapter features 5706 and 5707 for Operations Console (LAN) and the corresponding embedded two-port Ethernet LAN adapter within an IBM System i5 520, 550, and 570 model processor enclosure, sometimes also referred to a system unit.

When a system is managed by a Hardware Management Console (HMC), you specify which 5250 console type you want to use as the main 5250 console to an i5/OS partition through the HMC interface.

When your system is not managed by the HMC, you can order (or change) the system to be managed by Operations Console (Direct attached), or Operations Console (LAN attached). When shipped new, there are default algorithms used by the system to find the card or port to be used as the console. These rules should be understood.

Embedded LAN ports support

V5R4 and V5R3 with microcode level V5R3M5 now support using the first of the two embedded LAN ports (not HMC ports) in 520, 550, and 570 processor enclosures for Operations Console (LAN).

This same Operations Console over the LAN support is also available for the first port on the individually ordered 5706/5707 2 line LAN adapters, which can be plugged in various processor enclosures and I/O towers and drawers. For information about 5706/5707 support see “Two-port Ethernet LAN support” on page 247.

This option can save on cost of a separate LAN IOA (and possibly an IOP controlling the IOA) and the associated PCI slots being occupied.

The embedded ports support for Operations Console on the LAN is for V5R3 i5/OS and V5R3M5 LIC and V5R4 i5/OS with V5R4M0 LIC as follows: 5553 feature code for embedded Ethernet for 520, 550, and 570.

There are two Ethernet LAN ports of this type available within the processor enclosure. However, if the first of the two ports is used for the LAN Operations Console, that port must be dedicated for that use only and nothing else. The second port can be used for normal LAN connectivity to i5/OS and running applications. Since these two ports are managed under a single adapter, both ports have to be owned by a single partition.

Using any one of these ports for a LAN Operations Console will require an i5/OS V5R4 if your system is a 1.5/1.65 GHz model 520 or any 550, 570, or 595. On the other hand, you can also use i5/OS V5R3 and LIC level V5R3M5 if you are using a 1.9 GHz model 520.

Two-port Ethernet LAN support

Another new V5R4 i5/OS feature is the use of the first port of features #5706/5707 LAN adapter for Operations Console (LAN). This #5606/#5705 LAN adapter has been available for several years but could not be used for Operation Console support. The #5607/#5707 has never required a controlling IOP. This already available feature is a IOP less IOA Ethernet 2-port LAN adapter, which in V5R4 and V5R3 with V5R3M5 LIC can now be used for Operations Console (LAN). Like the processor enclosure embedded ports, if the first one is used for the console, it must be dedicated to the console function only while the second port can be used for normal utilization under one partition only.

This is supported as 5557 feature code for 5706/5707 Ethernet IOA for 595 or any I/O tower/drawer.

There are specific LAN adapter (IOA) considerations for Operations Console (LAN) when an HMC is not used to specify which adapter to select. The next paragraphs describe the considerations.

Special considerations

The following are special considerations:

- ▶ 520, 550, 570 models: If present (all models are shipped with this), *only* the first port of the embedded 2-port Ethernet LAN adapter is used as the default for the console. You must turn off this support if you plan to use a different supported LAN adapter in a card slot either in the 520, 550, or 570 processor enclosure or a separate I/O tower or drawer.
- ▶ 595 models: This model requires an HMC, and therefore tagging the desired IOA directly through the HMC interface is used to specify the console instead of by card location.

Operations Console adapter placement rules

Each System i5 model has special rules in terms of placement of adapters for Operations Console use. In Figure 11-1 it shows these adapter placement rules for i520, i550, and i570.

<p>Model i520</p> <p>Operations Console (LAN): first embedded Ethernet port (T5) unless support has been turned off. Card slots: C2 or C5 (no IOP) Card slots: C2 driven by IOP in C1 Card slots: C5 driven by IOP in C3 or C6</p> <p>Operations Console (Direct): Asynchronous communication adapter Card slots: C3 (no IOP) Card slots: C3 driven by IOP in C6 If an IXS is present it will occupy C5/C6 LAN or Asynchronous communication adapter goes into C2 and driven by IOP in C1, if needed.</p> <p>Model i550</p> <p>Operations Console (LAN): First embedded Ethernet port (T9) unless support has been turned off. Card slots: C4 (no IOP) Card slots: C4 driven by IOP in C3</p> <p>Operations Console (Direct): Asynchronous communication adapter Card slots: C2 (no IOP) Card slots: C2 driven by IOP in C1 If an IXS is present it will occupy C2/C3 Asynchronous communication adapter goes into C5 and driven by IOP in C4, if needed. Note: Console support using a LAN connection is provided by the embedded port only.</p> <p>Model i570</p> <p>If multiple Processor Enclosure Units, the one with the load source DASD is used for console support.</p> <p>Operations Console (LAN): First embedded Ethernet port (T6) unless support has been turned off. Card slots: C4 or C6 (no IOP) Card slots: C4 and C6 can be driven by an IOP in C3</p> <p>Operations Console (Direct): Asynchronous communication adapter Card slots: C2 (no IOP) Card slots: C2 driven by IOP in C1</p>
--

Figure 11-1 System i5 Operations Console adapter placement rules

If you have selected to use Operations Console (LAN) we recommend that you install the server using the embedded adapter port and then migrate the console to use another location if you wish. If you turn off the embedded port support for console, an adapter card located in a specific location, by model, can be used for console use. The only IOP-less LAN adapter is the 5706/5707 10/100/1000 Mbps Ethernet adapter. All other LAN adapters will require an IOP. Like all other console supported adapters, even though the 5706/5707 runs without an IOP it still must be located in a location specific to a LAN console. Refer to Figure 11-1 for the proper adapter locations.

When upgrading to a new System i5 520, 550, 570, or 595, your upgrade instructions should contain documentation on how to continue using your current console hardware or change to one of the POWER5-based options discussed here.

For a comprehensive description of IBM System i5 and eServer i5 PCI card slot plugging rules, refer to *PCI and PCI-X Placement Rules for IBM System i5, eServer i5, and iSeries servers with i5/OS V5R4 and V5R3*, REDP-4011. This can be found at the IBM Redbooks Web site:

<http://www.ibm.com/redbooks>

For detailed cable connection information refer to the eServer IBM Systems Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>

For full details on Operations Console support on IBM System i5 and eServer i5 refer to the IBM System i5 and eServer i5 Managing Operations Console PDF found at the iSeries V5R4 Information Center. This PDF includes information about how to turn off the POWER5-based default algorithm for selecting an Operations Console (LAN) adapter port. This document can be found in the iSeries Information Center, V5R4:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/index.jsp>

11.1.2 Hardware Management Console

The Hardware Management Console (HMC) is a pre-installed Linux-based workstation connected to the System i5 server via two HMC dedicated ports. This console is required on all POWER5 servers to create and change logical partitions (LPAR) or to use Capacity on Demand (COD), but is never required to run the partition.

The HMC has full function enablement of capacity on demand processor and memory movement capabilities. However, permanent processor activation can also be accomplished through the system's Advanced System Management Interface (ASMI), which is provided through the system's service processor.

Each HMC must have access to a graphics display, a keyboard, and a mouse. In addition to the displays offered as features of this HMC, the IBM 7316-TF2 Flat Panel Console Mounting kit and 7316-TF3 Flat Panel rack-mounted display, keyboard, and mouse and associated VGA switch are supported as well as existing IBM T541H, P76, P77, P260, and P275 graphics displays. The 7316 configuration must be specifically ordered and is recommended for users who are in a space-constrained environment.

The desktop model HMCs come with one Ethernet LAN adapter port for external communications. If you want additional connections, such as using the WebSM tool from another workstation, you need to order a second physical Ethernet card (comes with two LAN ports).

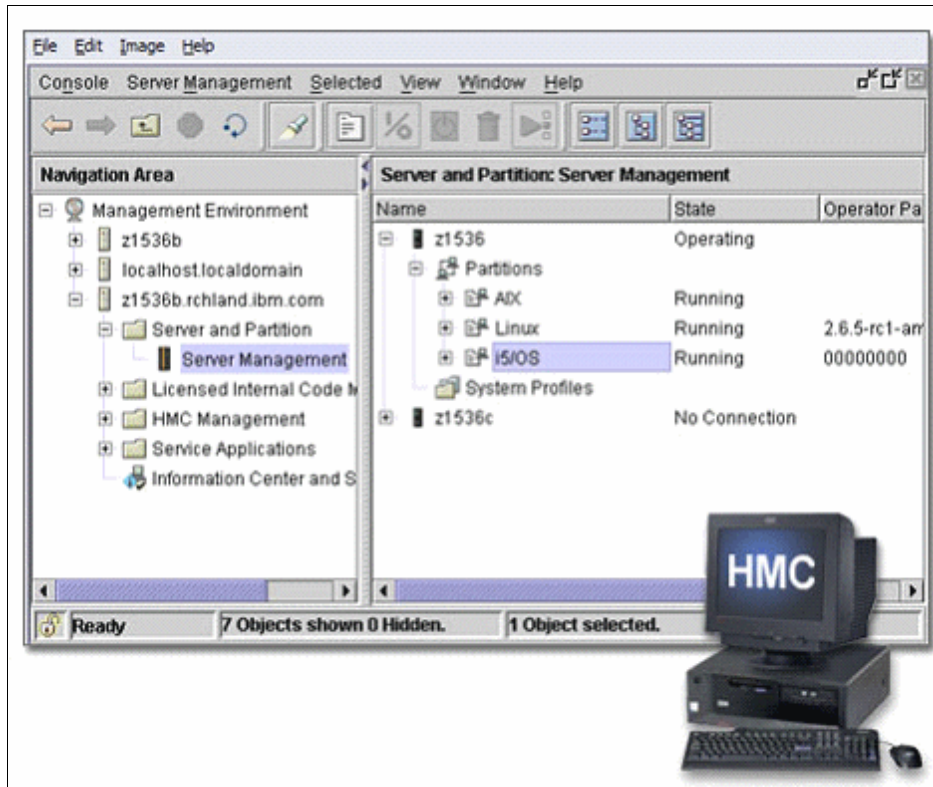


Figure 11-2 Hardware management console

The 7310-C05 Deskside HMC is a new model available first quarter 2006. Similar to the 7310 C04, the C05 comes with one physical Ethernet port. The IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter (#5706) should be ordered to provide a second physical Ethernet connection. Up to two feature number 5706s can be ordered for a maximum of five Ethernet ports. Each server partition will have a physical connection to the network, and the HMC will be logically connected to each partition via its connection to the network.

Over time, new HMC models will be announced and older models will be withdrawn from marketing. Watch for 2007 hardware announcements in this area, including 907-026 and 107-093.

One Ethernet connection is required for each POWER5 or POWER5+ server controlled by the console. Either a serial port for modem attachment or an Ethernet connection is required if the Service Agent call-home function is implemented. Higher availability configurations may require the use of an Ethernet hub. An Ethernet hub will provide a physical connection from the HMC to each server rather than a logical Ethernet connection over a LAN.

To back up HMC data, a DVD-RAM media is required. Other media are not usable.

A System i5 that is managed by HMC should have at least a server firmware at the SF235_160 level to be able to run V5R4 on logical partitions that are in full-production mode.

Advanced system management interface

The advanced system management interface (ASMI) is an interface to the service processor that lets you access detailed service functions like auto power restart and view information about the system such as error logs and vital product data.

ASMI is required to do some functions without HMC, such as activating capacity on demand or activating the new i5 520 accelerator function on specific models.

Web System Manager

The Web System Manager (WebSM) has been enhanced to support an expanded set of languages. WebSM can be downloaded from your HMC and installed on any Windows or Linux PC with a Java enabled browser. Figure 11-3 shows the set of supported languages.

▪ Austrian German	▪ Dominican Spanish	▪ Hungarian
▪ Luxembourg German	▪ Ecuadorean Spanish	▪ Japanese
▪ German	▪ El Salvadorian Spanish	▪ Swiss Italian
▪ Swiss German	▪ Guatemalan Spanish	▪ Italian
▪ Australian English	▪ Honduran Spanish	▪ Russian
▪ Belgian English	▪ Mexican Spanish	▪ Slovak
▪ Belgian English	▪ Nicaraguan Spanish	▪ Belgian Dutch
▪ Canadian English	▪ Panamanian Spanish	▪ Dutch
▪ Hong Kong English	▪ Paraguayan Spanish	▪ Brazilian Portuguese
▪ Indian English	▪ Peruvian Spanish	▪ Czech
▪ Irish English	▪ Puerto Rican Spanish	▪ Korean
▪ New Zealand English	▪ Spanish	▪ Simplified Chinese
▪ Philippine English	▪ American Spanish	▪ Hong Kong Traditional Chinese
▪ Pakistani English	▪ Uruguayan Spanish	▪ Singaporean Simplified Chinese
▪ Singaporean English	▪ Venezuelan Spanish	▪ Traditional Chinese
▪ South African English	▪ French	
▪ American English	▪ Belgian French	
▪ British English	▪ Canadian French	
▪ Argentine Spanish	▪ Swiss French	
▪ Bolivian Spanish	▪ Luxembourg French	
▪ Chilean Spanish		
▪ Colombian Spanish		
▪ Costa Rican Spanish		

Figure 11-3 Supported languages by WebSM

11.1.3 Thin Console

This new console type (feature #9944) comes as an option for i520 and i550 models. An alternative to twinax-attached 5250, HMC, and Operations Console (LAN or Direct attach), it is designed to connect to partition one and provide a 5250 interface to a single i5/OS partition.

IBM does not support environments with both an HMC and Thin Console concurrently active on the same machine. If you desire a different type of console interface at a later time, use normal console type switching procedures.

Figure 11-4 depicts an example connection of the Thin Console attached to one of the HMC ports on the IBM System i5 model 520 or 550. The HMC ports support communications with the system's service processor.

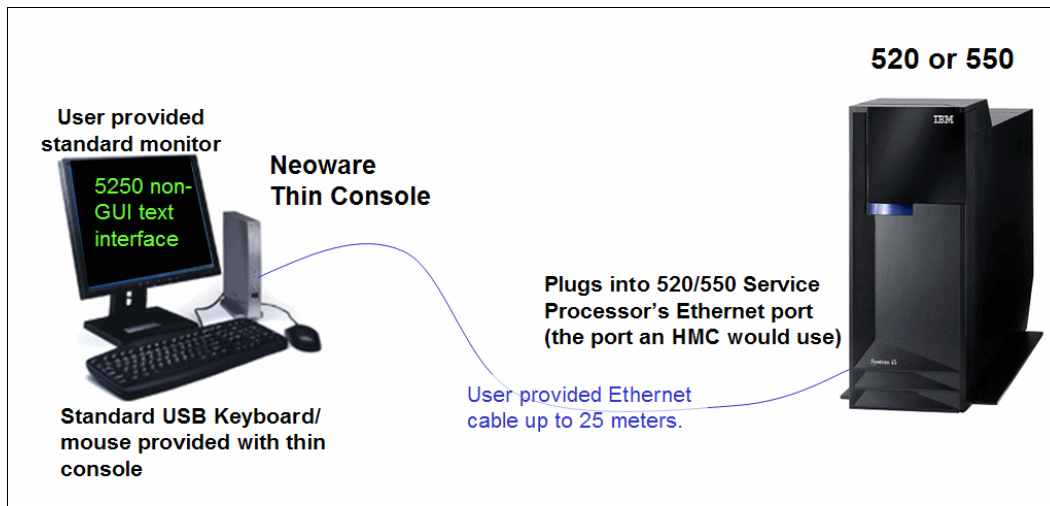


Figure 11-4 Thin Console: Network configuration example

Figure 11-5 (model 520) and Figure 11-6 on page 253 model (550) show you the back of each model and highlight where the two HMC port connections are in the processor enclosure. Either HMC port can be used as the install processes detects the connection. As stated earlier, you cannot operate with both an HMC and a Thin Console attached to the system at the same time.

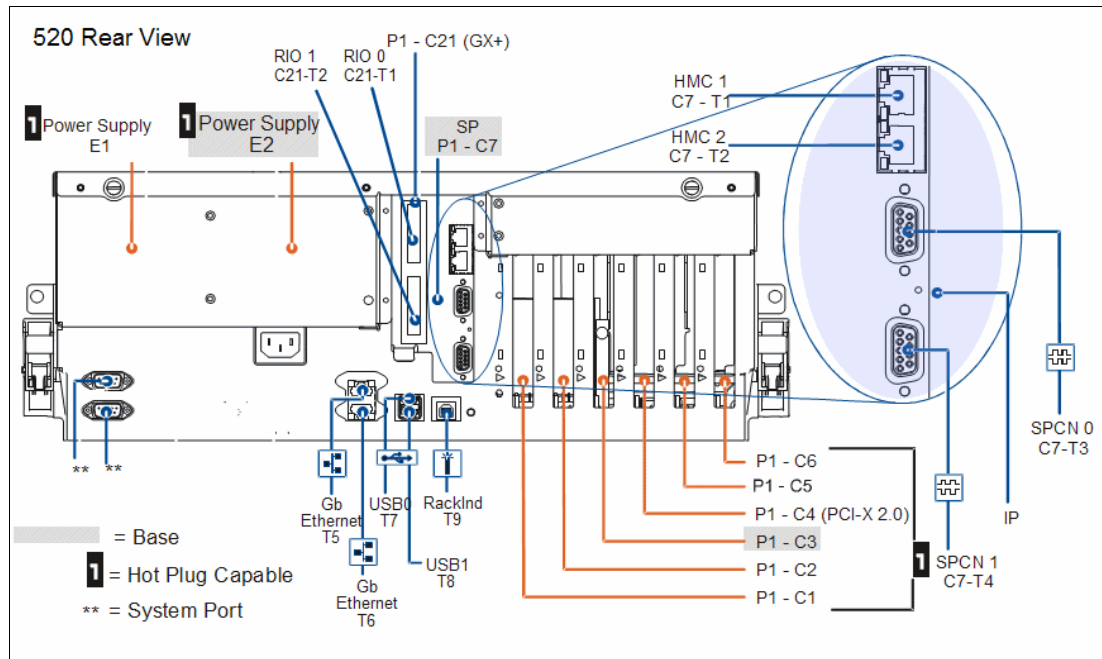


Figure 11-5 i520 processor enclosure: rear view

Figure 11-6 shows the rear view of the i550 processor enclosure.

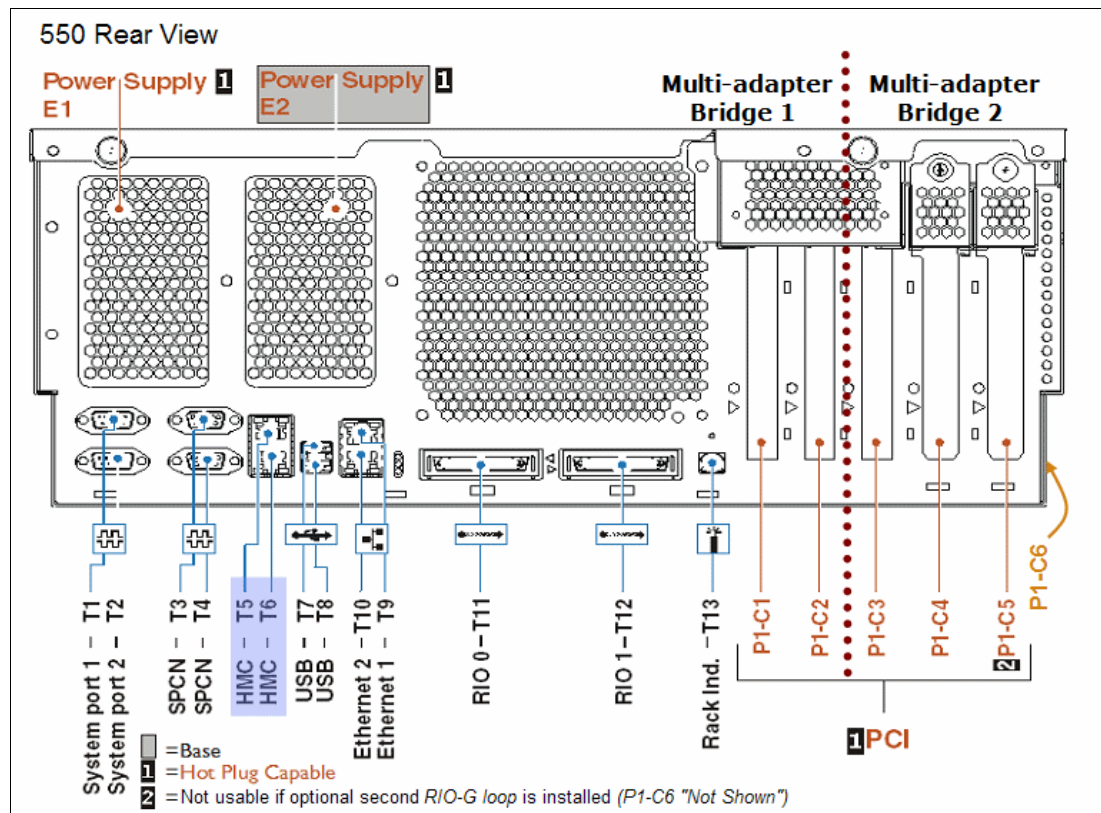


Figure 11-6 i550 processor enclosure: rear view

A Thin Console is recommended only for i5/OS partition use. It can also be used with 5250 applications, as it supports i5/OS commands.

This customer setup product includes a logic unit, power cord, keyboard, and mouse (can be used for some 5250 emulation functions). It does not, however, include a standard monitor that supports 1024 x 768 resolution at 75 Hz to plug into the logic unit of the Thin Console and a standard Ethernet cable up to 7.62 m (25 ft.) that attaches the logic unit to the 520 or 550 processor enclosure.

Neoware, Inc., builds the Thin Console for IBM and places its software on the logic unit for IBM. This product includes a standard two-year hardware warranty from Neoware, which can be upgraded to a three-year hardware warranty by registering at the Neoware Web site. Neoware supports software preloaded on the Thin Console.

If you later need software updates, you can download the updates to any other workstation from the Neoware Web site:

<http://www.neoware.com/support/thinclient.html>

Place the downloaded software onto a USB-attached memory stick and then place the memory stick in the Thin Console logic unit. Follow the update installation instructions.

Thin Console setup information

The two-step setup information is stored in the Thin Console. After the initial setup, the 5250 operator console screen is all that is seen when the Thin Console is turned on. If you move the Thin Console to a different system unit, or if language changes are required, you can

reset the stored setup information back to the initial factory settings by pressing Ctl+Alt+End. You can then make changes by repeating the setup process.

The Thin Console is a single-byte character set (SBCS) device. It does not support double-byte character set (DBCS) or bidirectional languages. The keyboard languages shown in Figure 11-7 are supported.

Thin console supported languages

Country	Keyboard
United States	US English
Italy	Italian, Swiss French/German
France	French, Swiss French/German
Germany	German/Austrian, Swiss French/German
Spain	Spanish
United Kingdom	UK English
Canada	US English, Canadian French
Netherlands	Dutch
Belgium	Belgium/French, Dutch
Switzerland	Swiss French/German, Italian
Sweden	Swedish/Finnish
Austria	German/Austrian
Denmark	Danish
Portugal	Portuguese
Norway	Norwegian

Figure 11-7 Thin console supported keyboard languages

11.1.4 Service support level console function changes with V5R4

Here are some service level enhancements to console functions made available with i5/OS V5R4:

- ▶ The system no longer requires a password when creating a System i5 Service Tools device ID. For those familiar with the System i family of systems, the Service Tools are accessed either via the Dedicated Service Tool (DST) or System Service Tools (SST) interfaces.
- ▶ The server can accept console service functions in D-mode even in the event that there is an uninitialized hard disk drive present.
- ▶ The server can force an exit at C6004508 (console not found) in D-mode through function 21.
- ▶ The DST sign-on window associated with the takeover and recovery function (Note: This device can become the console.) no longer has any PF keys because the only function allowed is sign-on.
- ▶ The takeover and recovery option takes effect immediately.
- ▶ The takeover option is now supported in D-mode. The Console Information Status window that displays during console takeover or recovery changed to make it easier to see if the *Take over the console* option is set to YES or NO.

11.2 Enhancements to logical partitioning

LPAR capabilities associated with IBM System p5 and eServer p5 systems can now be ordered for partitions running on an IBM System i5 and eServer i5 system via the Advanced POWER Virtualization feature.

11.2.1 Advanced Power Virtualization

The Advanced POWER Virtualization feature (based on the IBM System p5 Advanced POWER Virtualization feature) provides two capabilities to IBM System i5 and eServer i5 systems (POWER5-based systems):

- ▶ A Virtual I/O server for AIX and Linux partitions

In October 2005 IBM announced a new release of the Virtual I/O server for IBM System p5 and eServer p5 systems that support LPAR without requiring and HMC — the Integrated Virtualization Manager (IVM). This function does not work on IBM System i5 and eServer i5 systems, but the base virtual I/O server does.

- ▶ Partition Load Manager (PLM) for AIX 5L partitions

Note: IBM System p5 Advanced POWER Virtualization feature is required for all of the following capabilities:

- ▶ Virtual Ethernet
- ▶ Shared Ethernet Adapter
- ▶ Virtual SCSI Server
- ▶ Micro-Partitioning™ technology
- ▶ Partition Load Manager

Virtual I/O server

The virtual I/O server makes it possible to share hardware that is not supported by the i5/OS. This technology is the answer should you choose not to have the i5/OS hardware provide the virtual I/O to an AIX 5L or Linux partition, as shown in Figure 11-8.

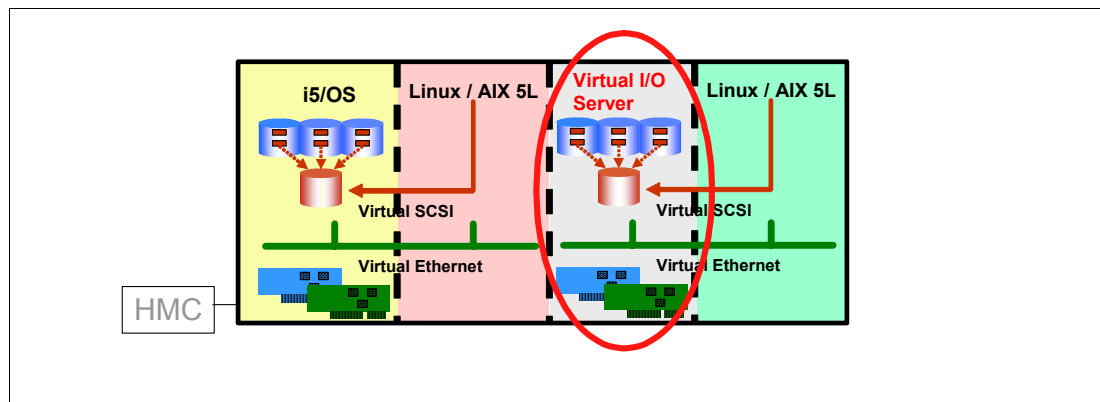


Figure 11-8 Virtual I/O server

Partition Load Manager

Partition Load Manager (PLM) provides automated processor and memory distribution between dynamic AIX 5L 5.3 partitions. The PLM application is based on a client/server model for the sharing of system information, such as processor or memory events, across multiple partitions on the same server. PLM allows AIX 5L resources within one server to be

automatically adjusted between the LPARs based on schedules, policies, priorities, and utilization. However, this feature is only available for AIX 5L partitions. It provides a higher level of automation management across partitions in a single POWER5 server.

PLM differs most from Enterprise Workload Manager (eWLM) in that eWLM focuses on multi-server environments.

i5/OS does not use Partition Load Manager at this time.

Advanced POWER Virtualization - ordering

The Advance POWER virtualization feature can be ordered as follows:

- ▶ Optional hardware features
 - \$nnn per processor #7940 (520), #7941 (550)
 - \$yyy per processor #7942 (570), #7992 (595)
- ▶ Software features (PID and SWMA)
 - No charge: Virtual I/O server with one-year SWMA and PLM with one-year SWMA
 - Chargeable: three-year SWMA

Check the sales manual pages for the latest prices.

- ▶ Quantity features required

One Advanced POWER Virtualization feature is required for every processor's worth of workload by partitions that are using the Virtual I/O Server or PLM functions during peak utilization.

For example, four AIX 5L/Linux partitions all use one or more of these functions. These four partitions have a maximum total of two processors worth of work during the maximum peak time. In addition, the Virtual I/O Server partition uses a maximum of ½ processor during the same time. The customer should order three Advanced POWER Virtualization features ($2 + 0.5 = 2.5$, rounded up to 3).

11.2.2 Virtual Partition Manager

Virtual partition manager is an i5/OS V5R3 based-tool to create simple Linux partitions on an IBM System i5 or eServer i5 (POWER5-based) system without having to have HMC. You can have a maximum of one i5/OS partition and up to four Linux partitions and four virtual Ethernets. However, these four Linux partitions must use virtual I/O.

This makes LPAR configuration straightforward simply by using a 5250 workstation interface to i5/OS DST-like interfaces. Although Virtual Partition Manager does not support dynamic LPAR, it does on the other hand support uncapped partitions.

For more information see redpaper *Virtual Partition Manager A Guide to Planning and Implementation*, REDP-4013. You can access this redpaper from:

<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4013.html?Open>

Note: When you have partitioning, generally you do not want IBM Manufacturing to complete the i5/OS implementation (unless you are using custom manufacturing options #8453, 0454 and 0455). When IBM loads the server, it assumes that all of the disk drives are owned by one i5/OS. But if partitioning into multiple i5/OS partitions, where some of these disk drives will be used by the second partition, it is different. In this case, now the you have to undo things that IBM Manufacturing did. You have to remove the disk drives from one i5/OS partition and move them to another i5/OS partition. The same is true for disk drives that will be owned/managed by Linux.

But if all of the drives will be used by i5/OS (for example, any Linux or AIX 5L partition will leverage i5/OS virtual I/O or will be supported by I/O ordered separately), then you do want IBM Manufacturing to finish loading i5/OS on the server. That is why the new #0496 Force i5/OS Preload feature has been announced.

11.2.3 LPAR simplification update

For several years the LPAR Validation Tool (LVT) has been successfully used to help you plan a logical partition's hardware configuration with output that helps you order the hardware configuration for one or more partitions on the same system. Reports (system plans) generated by the LVT can be input to the hardware order process. Additionally, with LVT output submitted with the order, you can specifically order one of several LPAR configuration and software pre-load features to ensure that your delivered hardware configuration matches your LVT report output.

These features simplify the ordering and installation process, which not only increases customer satisfaction but also reduces IBM service costs to resolve partition hardware configuration problems.

You can order your system with the following feature codes should you have decided to do logical partitioning and provide output from the Logical Partition Validation Tool (LVT):

- ▶ #8453 customer-specified placement
 - Based on LVT output, place cards/disk drives in requested physical location
 - No charge from IBM Manufacturing (chargeable if done on site)
- ▶ #0454 LPAR partitioning initialization
 - Configures partition and assigns correct resources based on LVT output
 - Chargeable — one feature per partition initialized
- ▶ #0455 LPAR OS preload
 - Load i5/OS or AIX 5L into partition created by #0454
 - Chargeable — one feature per partition loaded

For more complete information about these partition configuration options in this book, refer to 2.19, “Customer Specified Placement and LPAR preload” on page 55.

More LPAR simplification tools - 2006

The LPAR Validation Tool (LVT) available October 2005 was another incremental step in getting its output as input to the IBM econfig order process. Output from the Performance Management for System i service offering could be used as input to the IBM Workload Estimator (WLE) sizing tool. Using WLE helps you configure a sufficient amount of processor capacity, main memory, and disk arms to yield satisfactory performance in a system with one or more logical partitions.

Figure 11-9 shows an overall process of integrating support for multiple IBM system platforms into the size, order, and configuring logical partitions process. This minimizes reentering information when using the different tools during the size, order, and configuration process.

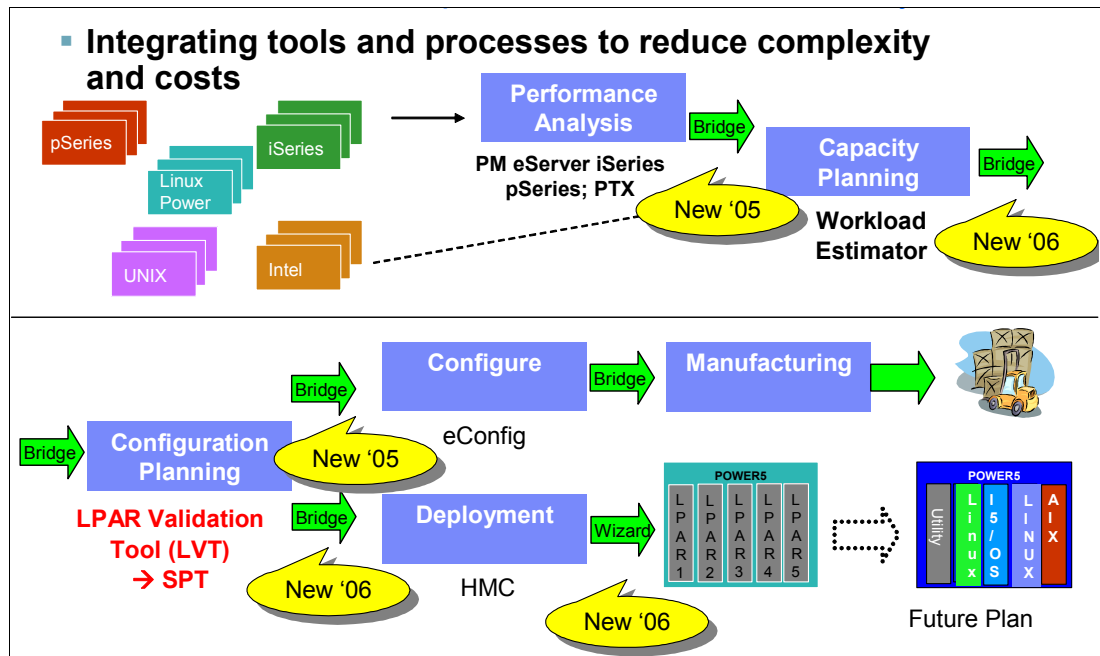


Figure 11-9 LPAR simplification process

In 2006 the LVT was replaced with the Systems Planning Tool (SPT).

11.2.4 IBM System Planning Tool

The IBM System Planning Tool (SPT) is the next generation of the IBM LPAR Validation Tool. It contains all of the functions from the LVT and has more integration with output from the IBM Performance Management (PM) for the IBM System i and System p systems offerings. The SPT can also take output from the capacity sizing tool — the IBM Systems Workload Estimator (WLE).

The output of the LVT is a file of type .lvt. The output of the SPT is a file of type .sysplan. LVT is no longer available for downloading. SPT converts .lvt files to .sysplan files.

System plans generated by the SPT can be deployed on the system by the Hardware Management Console (HMC). This means that the HMC can use the SPT system plan as a base for your partition configuration. Your HMC must be at release level V5R2 (available February 2006) to support this input.

The SPT is the next step in assisting you in system planning, design, and validation, and to provide a system validation report that reflects the user's system requirements, while not exceeding system recommendations. The SPT is a PC-based browser application designed to be run in a standalone environment.

For more information than we present in this appendix, refer to:

- ▶ The SPT Web site at:
<http://www-03.ibm.com/servers/eserver/support/tools/systemplanningtool/>
- ▶ IBM Redbook *LPAR Simplification Tools Handbook*, SG24-7231-00, at:
<http://www.ibm.com/redbooks>

Support is available via e-mail at rchspt@us.ibm.com.

The latest version of SPT, 1.00.020 includes the following enhancements over earlier versions:

- ▶ Input from the Performance Manager for the i and p systems and OpenPower™ systems, and the Workload Estimator (WLE) integration to aid in sizing workloads for your system
- ▶ System plans enabling deployment using the HMC
- ▶ Web Browser Interface
- ▶ Interactive Report
- ▶ Support for files created via the LVT
- ▶ Expanded online help

You can download the SPT to your PC workstation (from the SPT Web site listed above) running either a Microsoft Internet Explorer® or a Mozilla Firefox browser. SPT requires JavaScript™ enabled under your browser. During the SPT install wizard process you identify a TCP/IP port on your PC workstation that can be used to display Web pages on your workstation.

Note: To run SPT:

- ▶ Both pop-up windows and JavaScript must be enabled for your browser.
- ▶ Make sure that your firewall is configured to allow access to the SPT.exe program and is not blocking the port number specified in your SPT settings.
- ▶ You are not required to be connected to the Internet. If you want to access the Internet version of the Work Load Estimator you must be connected to the Internet.
- ▶ As of October 2006 there is no direct support of a .txt file version of the system plan. However, the SPT .sysplan file does contain a .txt file.

Until that .txt format option becomes available, you can extract the .txt file from the .sysplan file:

- Copy the .sysplan file and rename the file extension from .sysplan to .zip.
- Using an unzip utility, extract the entry Report.txt.
- Open Report.txt with a text editor to view its contents.

The October 2006 version has hardware navigation enhancements and:

- ▶ Interval planning

The estimator supports the capability to consider multiple time intervals in a single sizing. Often, customer workloads vary by hour of the day, by shift, by day of the week, by season of the year, or by whatever interval you require. You can vary the existence and the throughput level of any workload across the defined intervals. The Estimator will recommend a solution that meets the requirements for all intervals.

- ▶ Solution Visualizer

The estimator provides a visual option to help illustrate the immediate and growth solutions in terms of LPAR. This feature will assist the user to better understand the size of the LPARs, LPAR utilizations, and utilization variation among sizing intervals.

- ▶ Support for October 2006 Capacity Backup, Telephony, and Solution editions

The SPT can be used to plan partition solutions for the following hardware systems:

- ▶ IBM System p5
- ▶ IBM System i5
- ▶ IBM eServer p5
- ▶ BM eServer i5
- ▶ OpenPower
- ▶ iSeries 8xx and 270 models

The output of your dialogue with the SPT is a file, called the *system plan*, with a file type of .sysplan.

The SPT has extensive help text to assist you in creating and deploying a system plan.

In this book we provide an i5/OS-based example of SPT usage. We do not provide a complete example, but rather a subset of the partition configurations windows to give you a general understanding of the configuration options the SPT provides.

Three partitions, two i5/OS V5R3 partitions and one AIX 5L partition, are shown, along with an I/O tower for one of the i5/OS partitions.

Figure 11-10 shows two of the early configuration windows while using the SPT.

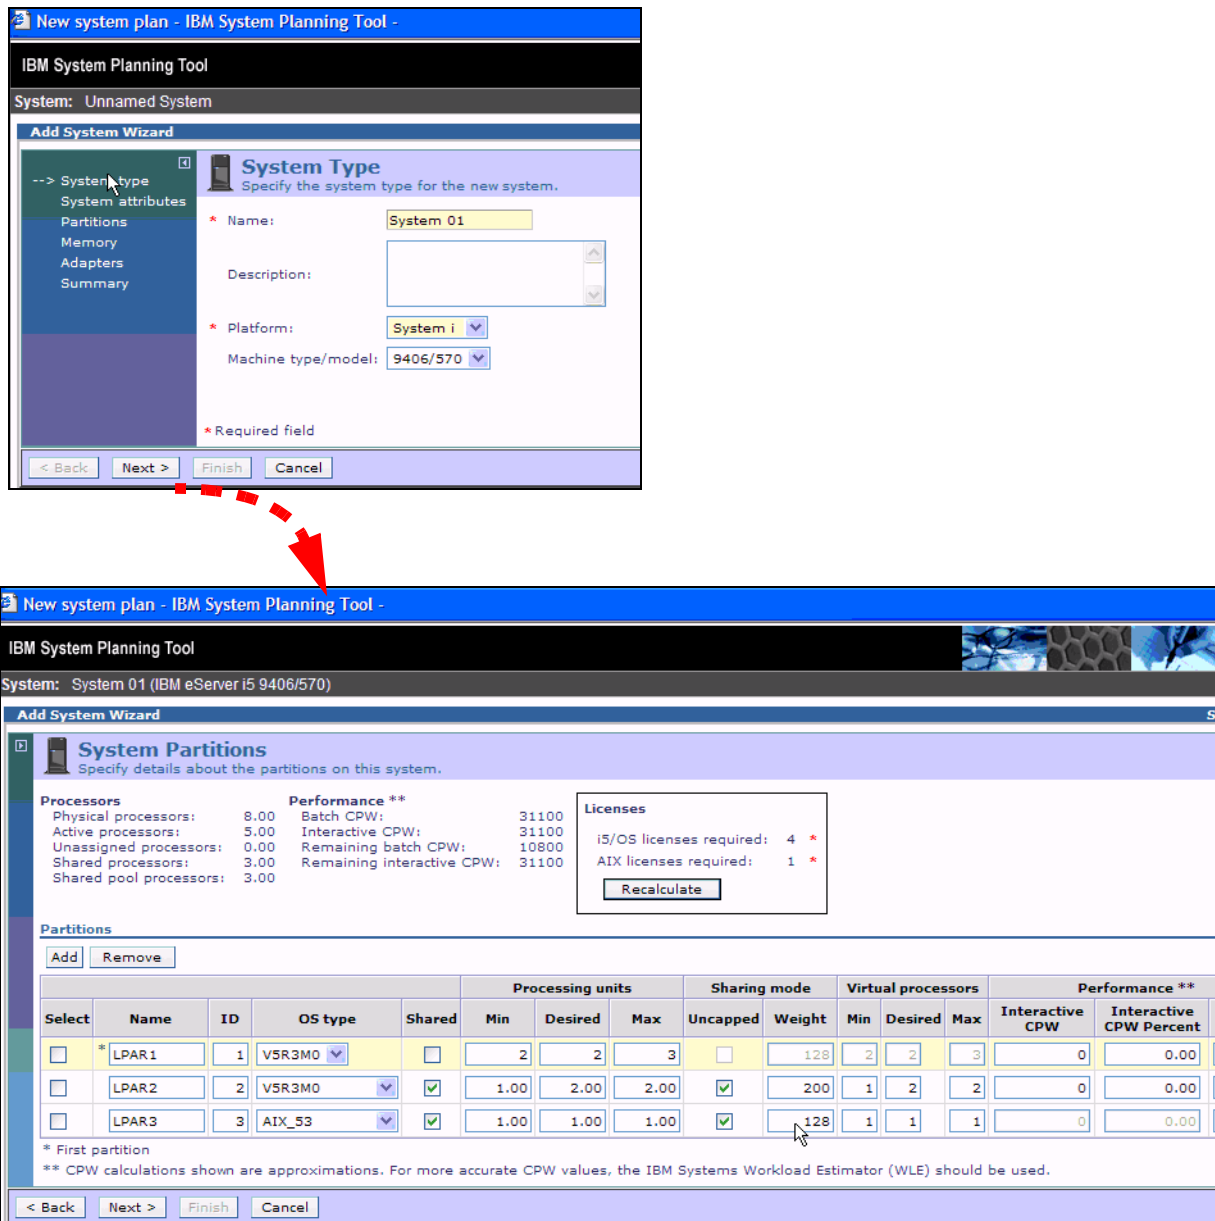


Figure 11-10 SPT example: starting partition configuration window for processors per partition

In the upper left window you see the name of your partition (System 01), selection of the System i platform, and selection of the System i5 Model 570. In the left navigation area of this window you have an indication of your progress through the major configuration steps through the SPT.

In the lower window we have already selected:

- ▶ Three partitions — two i5/OS partitions and one AIX 5L partition
- ▶ Partition processing unit assignments, including uncapped partition mode weight values and the number of virtual processors

Note that in the upper part of this window there is a summary of the running status of processor assignments, estimated CPW ratings, and number of i5/OS and AIX 5L licenses required.

When the .systemplan file is exported, these configuration details are included and can become input to the actual ordering (via eConfig) process. The details in this system plan can later (after your hardware configuration has been installed) be used as a starting base partition configuration for expediting the actual partition profile configuration on the HMC.

You can change many of these values from their original specification when the partition profile is configured. You assume responsibility of correlation with the hardware configuration within you system.

In the upper window example shown in Figure 11-11 we show an in-progress window of real memory (not virtual) assigned and virtual LAN and SCSI adapters. You can see the estimated amount of main memory the POWER5 Hypervisor requires.

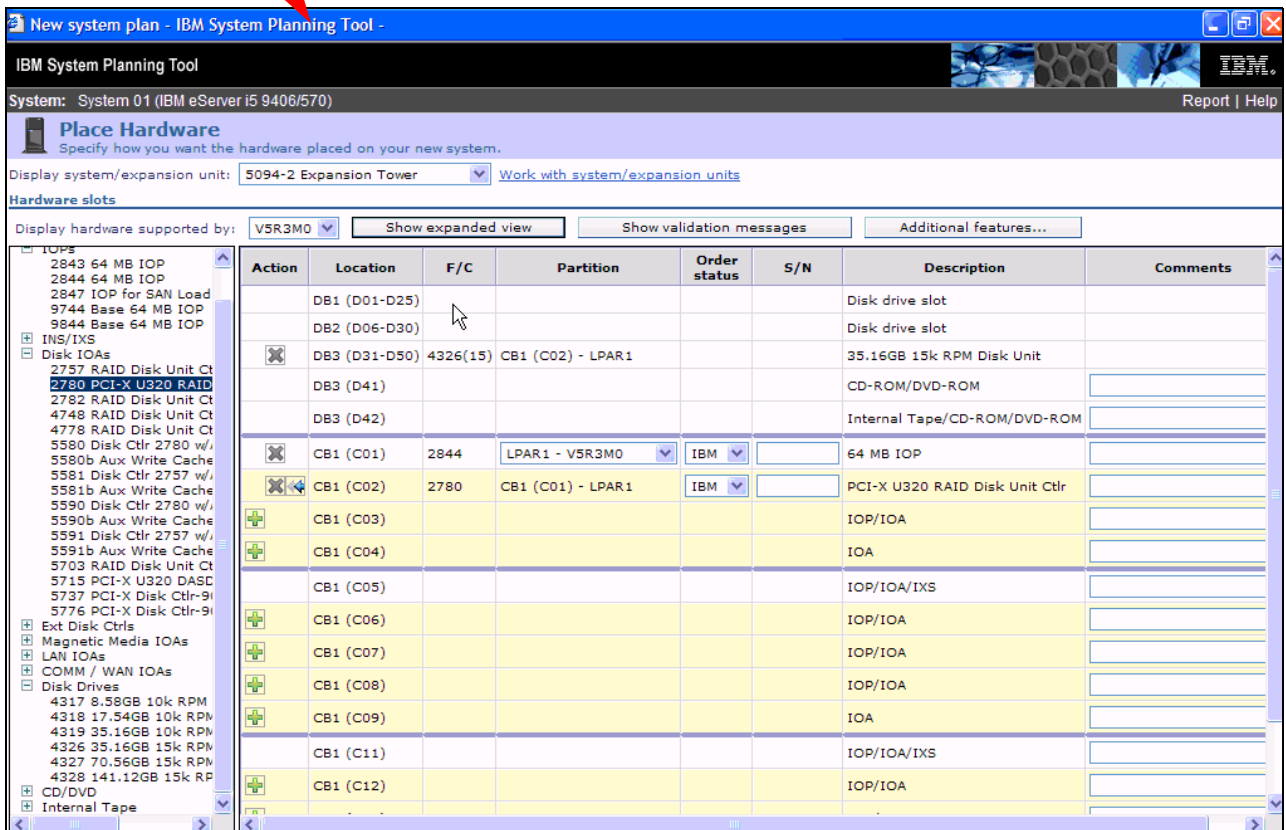
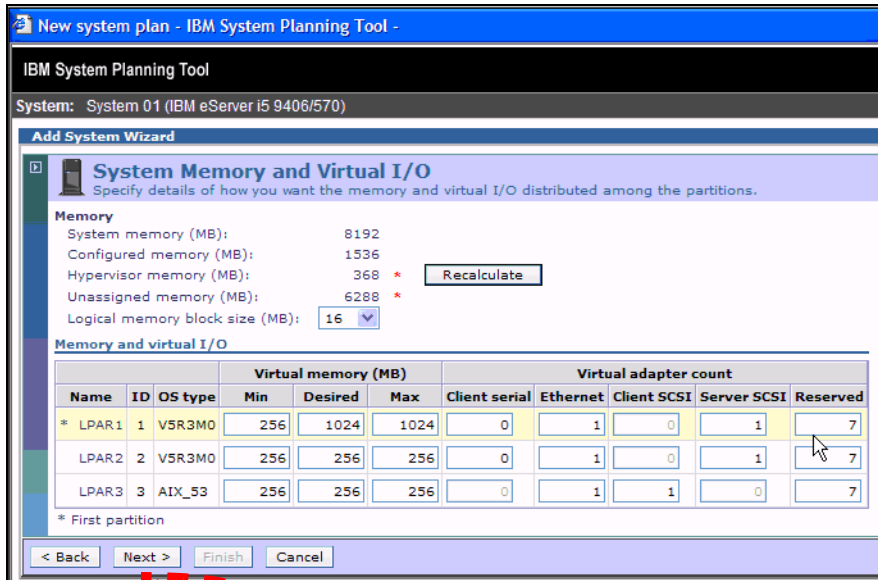


Figure 11-11 SPT example: memory, virtual I/O, and disk-based features

Again, these values can be left as specified when creating the partition profile through the HMC or changed at that time.

The lower window shows an in-progress configuration of an I/O tower (5094-2 Expansion Tower). It shows the initial card location assignments of the 2844 IOP, 2780 disk controller (IOA), and 15-K RPM (revolutions per minute) disk drives for partition LPAR1.

If you choose one of the options to have IBM configure your hardware and LPAR pre-load an operating system, you must be very careful if you change hardware locations from these settings when using the HMC to configure your partition profile.

See the following Web-based information that is available to assist you in using the SPT effectively:

- ▶ Web lectures:
 - System i: Configuring Logical Partitions with SPT (Web lecture and demo)
 - System p: Configuring Logical Partitions with SPT (Web lecture and demo)
 - Partition Deployment: Explore the system plan tasks of the HMC (Web lecture)
- ▶ Tutorials:
 - Deploy system plans: Information Center Web site (tutorial)
 - Planning and Deploying Virtual Servers (tutorial)

Integration with Workload Estimator and IBM Performance Management

The integration between the SPT and the Workload Estimator (WLE), not shown in our Figure 11-10 on page 261 and Figure 11-11 on page 263 examples, allows you to create a system that is based upon performance data from an existing system or that is based on new workloads that you specify through the WLE interface. If you want to create systems using the integration between the SPT and the WLE, either you must have the WLE installed locally on your PC or you must be able to access the online version of WLE through the Internet. For more information about how to use the online version of the WLE or download a copy of the WLE, see the WLE Web site:

<http://www-912.ibm.com/supporthome.nsf/document/16533356>

If you want to create a system using the SPT that is based upon performance data from an existing system, you must collect this performance data for at least one week using IBM Performance Management (PM) and have that summarized PM data sent to the IBM Performance Management location defined on your system.

For examples of how to activate and use PM on your server and of how PM and WLM work together, see the *Performance Management for IBM eServer iSeries and pSeries: A Systems Management Guide*, SG24-7122, which can be found at:

<http://www.redbooks.ibm.com/abstracts/sg247122.html>

Many WLE tutorials are available as Portable Document Format (.pdf) files, and the WLE allows you to print the final recommended system as a .pdf file. For more information about WLE, refer to the i5/OS Performance Management Web site at:

<http://www.ibm.com/eserver/iseries/perfmgmt>

Select the **Sizing Tools** link.

11.3 General LPAR updates

This section contains updates to miscellaneous LPAR-related support on an IBM System i5 model.

11.3.1 Reserve Capacity on Demand update

This section reviews how Reserve Capacity on Demand works. The PDF about System i and System p Capacity on Demand, which can be found at the IBM Systems Hardware Information Center

(<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>), contains more information. Use the search words *capacity on demand*.

Reserve Capacity on Demand (CoD) provides prepaid temporary processor capacity that is automatically used by the system when it determines that all fully activated processors are running at maximum capacity.

This enables you to help meet the demands of business peaks with automatic activation of the reserve processor capacity.

With Reserve Capacity on Demand, you purchase a reserve capacity prepaid feature that represents a number of processor days. You get an *activation code* you must enter into the system via the HMC.

Through this interface you must then assign the reserve processor capacity to the system's shared processor pool. A partition must be uncapped (can be changed dynamically between capped and uncapped) to be able to benefit of from the reserved processor capacity in a one processor day increment.

A reserve CoD processor day is consumed when the shared processor pool is 100% utilized (all active non-reserve processors available to the pool are being fully utilized) and 10% of a Reserve CoD processor is put into use for more than 30 consecutive seconds. Once a Reserve CoD processor day is consumed within a 24-hour period against a specific Reserve CoD processor, no additional Reserve CoD processor days will be charged against that Reserve CoD processor within the same 24-hour period.

When a Reserve CoD processor is used, one processor day is subtracted from the prepaid number of processor days. You pay only for the number of processor days that the reserve processors are actually used. The processors can remain active until all of the reserve processor days that you have paid for have been used.

You can also use the HMC to view all processor activation statistics, including reserve processor on demand utilization.

The following is extracted from the *System i and System p Capacity on Demand* PDF.

1. In the navigation area of the Hardware Management Console window, expand **Server** and **Partition**.
2. Select **Server Management**.
3. In the contents area, right-click the server on which you want to view the processor capacity settings.
4. Select **Manage On Demand Activations**.
5. Select **Capacity on Demand**.
6. Select **Processor** → **Capacity Settings**.
7. Select **Reserve CoD**. On the Reserve CoD page of the CoD Processor Capacity Settings, you can see information about the Reserve CoD state, activated Reserve CoD processors, inactive processors, prepaid processor days remaining for Reserve CoD, and hours remaining for activated Reserve CoD processors.

On demand memory and cross-partition movement

When you have ordered on demand memory, that memory is available in *memory day units*. Requested memory days equal the number of temporarily activated memory units multiplied by the number of days specified on a request for temporary capacity using On/Off Capacity on Demand.

After a request for temporary capacity is made on the server, the server records one memory day for each requested memory unit at the beginning of each 24-hour period that the request is running. One memory unit is 1 GB of memory.

You can start and stop requests for On/Off memory CoD, and you can also change the number of resources or days in a running On/Off request. This offering has contract requirements.

The HMC interface activates the capability to have on demand memory. It also provides the interface to add the newly or already available memory to a partition. Once assigned to an i5/OS partition, the added or moved memory units are, by default, assigned to the i5/OS *BASE memory pool.

Figure 11-12 shows the HMC interface window sequence to add (similar for move or remove) memory to an i5/OS partition. You can see the available memory units that can be added (or moved). You also see the Advanced memory window showing the time period in which the add (or move or remove) must complete. A message window appears when the add or movement has completed or the time period has expired before the requested action completes.

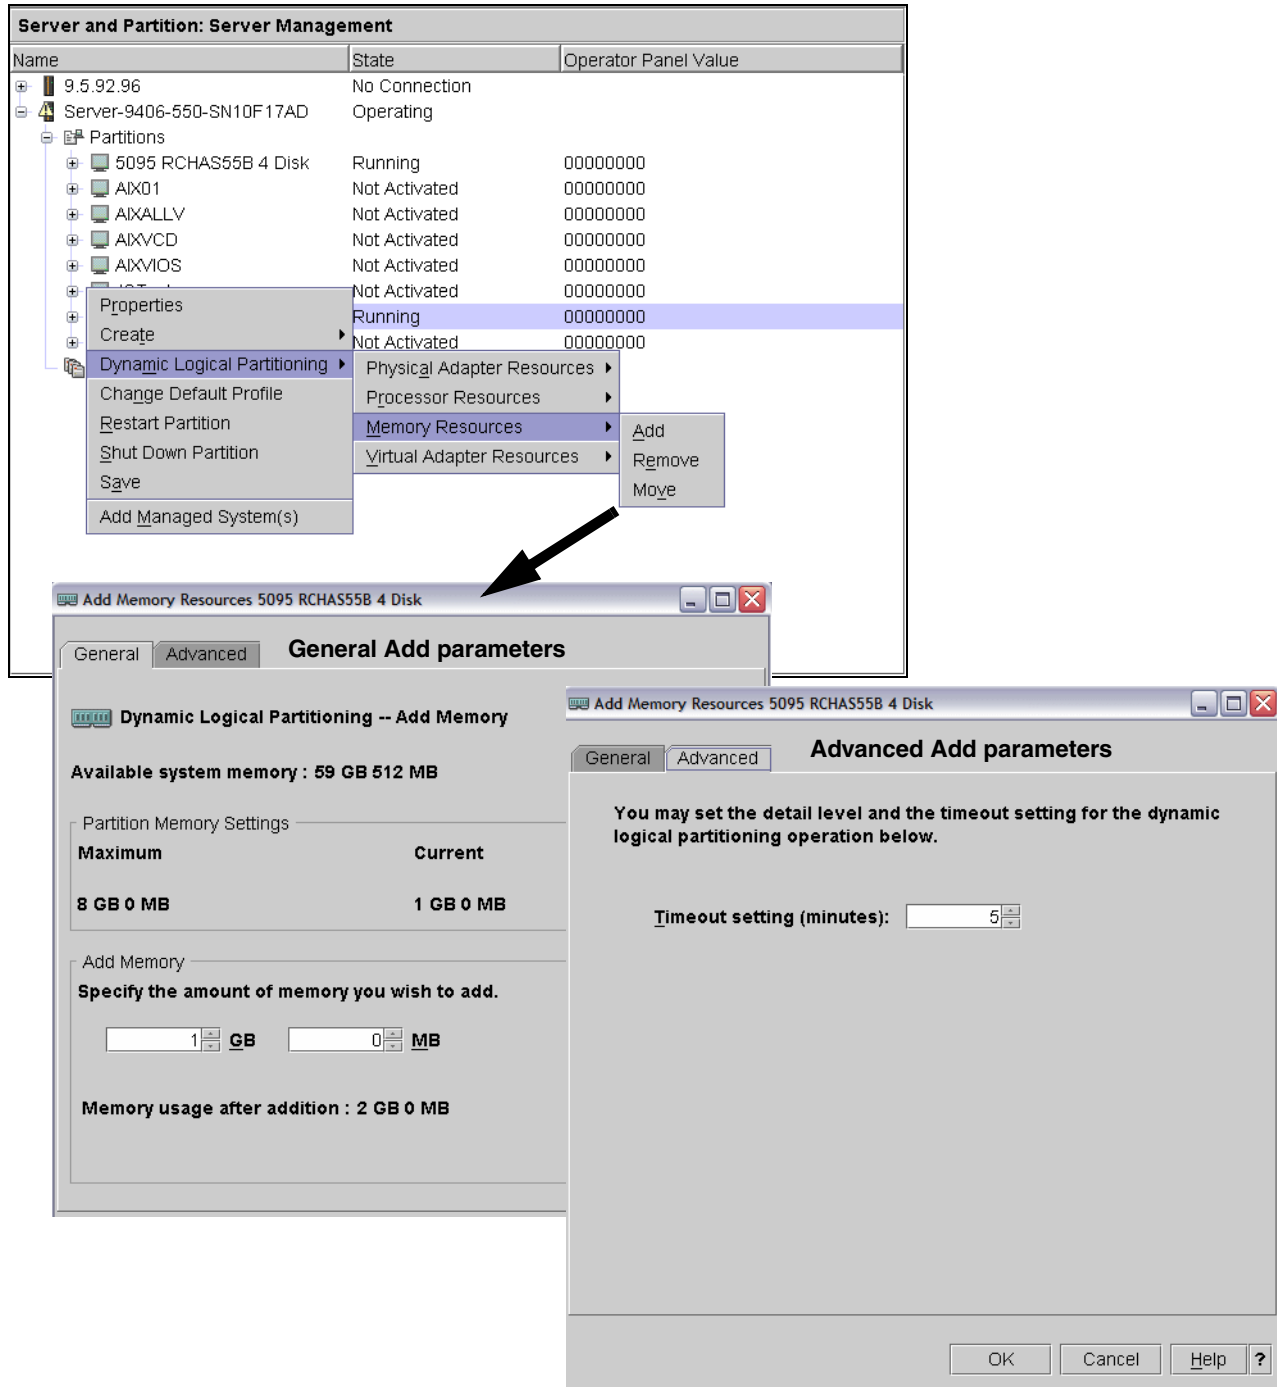


Figure 11-12 HMC interface to add memory to i5/OS partition

Once memory has been added or moved into the base (*BASE) storage pool of the i5/OS partition you must use a normal i5/OS command, such as Work with System Status

(WRKSYSSTS) or Work with Shared Pools (WRKSHRPOOL), to move the memory to another storage pool if you do not want all the added storage to remain in the *BASE pool.

Note: When memory is added or moved into a partition, the contents of that memory are first *zeroed* before they can be used by any software running in that partition.


Dynamic memory changes cannot cause the amount of memory in the base (*BASE) pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value).

If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

To prevent any data loss during dynamic memory movement, the system first writes any data from memory pages to the disks before making the memory pages available to another partition. Depending on the amount of memory you have requested to move, this might take some time.

Memory in each logical partition operates within its assigned minimum and maximum values. The full amount of memory that you assign to a logical partition might not be available for the logical partition to use. Static memory overhead that is required to support the assigned maximum memory affects the reserved or hidden memory amount. This static memory overhead also influences the minimum memory size of a partition.

For additional information refer to the PDF previously listed in this topic — *System i and System p Capacity on Demand*.



i5/OS-based application development languages and tools

This chapter covers V5R4 enhancements to various application development languages supported under i5/OS. This includes i5/OS CL, RPG, COBOL, Java, and C++.

12.1 Application developer roadmap

We start with a diagram of the overall application developers roadmap, from a programming languages and development environment viewpoint, for applications running under i5/OS. Before we dig into the V5R4 changes, we review this to provide a broader context.

Figure 12-1 shows the expected evolution of a development effort. Realize that in real life there is no requirement to take every application implementation all the way to the right with WebSphere Application Server servlets, portlets, eXtensible Markup Language (XML), and Web services. If you have existing programs, for example, that do business processing with no need to directly interface to Internet users or want them to be usable as Web-based services programs, leave them as currently implemented.

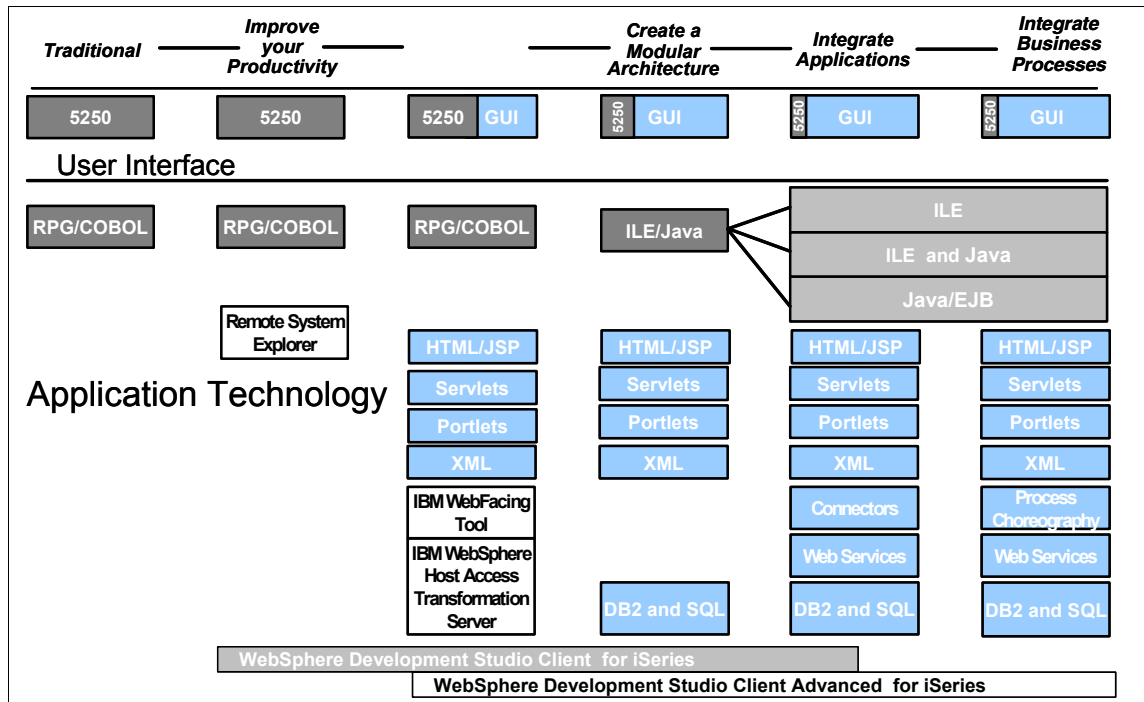


Figure 12-1 The IBM application development roadmap

On the other hand, often to remain competitive or have a competitive advantage, you need some use of advanced programming capabilities. This includes having your applications fully accessible to Web users, being able to run as a servlet or a portlet (multiple applications active on the same Web window), or being a Web Services program that other applications in the Internet can make use of.

Starting with traditional RPG and COBOL (and Control Language, or CL), we expect that you will find solid value in progressing at least through the Integrated Language Environment (ILE) standardization modular programming.

As you proceed to the right, you will be adding graphical user interfaces (GUIs). There are several ways to do this that include:

- ▶ Using IBM WebFacing and or IBM Host Access Transformation Services (HATS) that provide conversion between 5250 and HTML with little or no programming changes. These tools are commonly referred to as *refacing tools*.
- ▶ Some other IBM or non-IBM refacing tool. Simple examples of IBM tools would be IBM Host on Demand, which do not run under a WebSphere Application Server instance, but

uses Java applets on the Browser workstation and iSeries Access for Web's 5250 interface, which uses its basic 5250 servlets running under a WebSphere Application Server instance on the i5/OS system or partition. Non-IBM vendors providing tools include IBM business partners:

- Seagull Software LegaSuite
<http://www.seagullsw.com>
- Jacada Interface Server
<http://www.jacada.com>
- Look Software NewLook
<http://www.looksoftware.com>
- System Objects JACi400
<http://www.systemobjects.com>
- Arterial Software aXes
<http://www.arterialsoftware.com>
- Farabi Technologies Host Front
<http://www.farabi.com>
- Advanced BusinessLink
<http://www.businesslink.com>
- pks
<http://www.pks.de>

Further coverage of these business partner solutions are beyond the scope of this book. All of these products provide functions that should be considered, so we provide their Web addresses to assist you. In Chapter 13, “i5/OS-based WebSphere application support” on page 313, we provide updates about Webfacing and HATS, but thorough coverage of these tools is beyond the scope of this book.

- ▶ New programs written in programming language, such as Java (and C++) that have clear interfaces with HTML, XML, and JavaServer™ Pages™.

Through the use of HTML, XML, JavaServer Pages, and refacing (5250 <-> HTML/JSP™ translation) tools, you get your applications accessible via Web browsers. Along this path these portions of an application run as servlets and perhaps as portlets.

Towards the right you see the expansion of the ILE to include Java-based applications and use of the cross-operating system Internet-based components or concepts, such as Web Services. The theme here is collaboration among application developers.

As we stated earlier, in the real world mixtures of the various application technologies should be expected. For example, RPG programs can be doing the data processing, while Java and XML applications perform the SQL interfaces to the database, HTML-based user interfaces to the browser users, and cross-application data exchange (which typically uses XML for this interchange).

The following sub sections provide overview definitions of the common terms used with Web-based applications.

12.1.1 JavaServer Pages (JSPs)

JavaServer Page is a technology for controlling the content or appearance of Web pages through the use of servlets that are specified in the Web page and run on the Web server to modify the Web page before it is sent to the user who requested it.

A JSP is comparable to Microsoft's Active Server Page (ASP) technology. Whereas a JavaServer Page calls a Java program that is executed by the Web server, an Active Server Page contains a script that is interpreted by a script interpreter (such as VBScript or JScript®) before the page is sent to the user.

Think of the 5250-based DDS keywords that do things like provide a list of compare values that the data must first match before being processed, or provide a 5250-based coloring or other field highlighting function. This definition and processing is analogous to the processing of JSPs.

12.1.2 eXtensible Markup Language

XML stands for eXtensible Markup Language. Think of XML as the Internet way to define data type and length. For those application developers with a System 3x or AS/400 historical background, you can look at XML as the industry-standardized follow-on to AS/400 and i5/OS Data Description Specifications (DDS).

Though technically there are several implementations of XML, this format is well understood by nearly every provider of Internet application software. XML, created by the World Wide Web Consortium (W3C), is part of a large family of markup languages. It is defined as a metalanguage — that is, a language that describes other languages.

One of the goals of W3C was to make XML optionless so that it would remain pure, unlike HyperText Markup Language (HTML), which has many different conventions and as a result is rendered differently on various browsers. The difficulty that this presents in presenting data in a uniform fashion is planned to be alleviated by XML. Note how XML crosses many of the columns shown in Figure 12-1 on page 270.

12.1.3 Servlet

A servlet is a small program that runs on a server, commonly thought of in the context of the Java applet, a small program that is sent as a separate file along with a Web (HTML) page to a browser workstation client. The Java servlet runs on the server system in contrast to the Java applet, which runs on the workstation.

In a WebSphere-based application environment, the servlet typically provides interfaces to server operating system functions or application databases. The servlet could call many other programs written in any supported language, to actually perform the functions requested by the Web browser user. Consider the Java servlet as a different way of managing and implementing functions that also run on the host server, but performed in a Common Gateway Interface (CGI) application.

However, with Java running in the server, such programs can be implemented with the Java programming language. The advantage of a Java servlet on servers with lots of traffic is that they can execute more quickly than CGI applications. Rather than causing a separate program process to be created, each user request is invoked as a thread in a single daemon process, meaning that the amount of system overhead for each request is slight.

12.1.4 Portlet

A portlet is a component of a portal Web site that provides access to a specific information source or application, accounts payable, order entry, news updates, technical support, or an e-mail program, among many other possibilities. Portals aggregate different content into a single interface. Portlets connect the user to specific content within that interface. At the user interface level you have multiple applications active within the same window, in contrast to each application having its own window.

12.1.5 Web services

Web services are self-contained, modular applications that you can describe, publish, locate, and invoke over a network. Web services are intended to be used in the new, service-oriented approach to programming. This approach is based on the idea of building applications by discovering and implementing network services or by invoking available applications to accomplish some task.

This approach is independent of specific programming languages or operating systems. Web services delivers interoperability — the ability for components created in different programming languages to work together as if they were created using the same language.

To work productively there is an architecture for defining, cataloging, and using Web services together. Extensible Markup Language (XML) is the base to making this all work together.

Web services are described in Web Services Description Language (WSDL). A Web services description is registered in the Universal Description, Discovery and Integration (UDDI) directory.

UDDI is an industry initiative to create a platform-independent, open framework for describing services, discovering businesses, and integrating business services. The UDDI OASIS Standard specifies protocols for creating a registry for Web services, methods for controlling access to the registry, and a mechanism for distributing or delegating records to other registries.

OASIS is an international standards consortium involved with establishing standards for Web Services. One of the tasks it performs is to approve standards used in delivering Web Services, for example, approving UDDI Version 3.0.2 during 2005.

UDDI is commonly regarded as a cornerstone of Web services, defining a standard method for publishing and discovering network-based software components in a service-oriented architecture (SOA).

The Web Services Description Language can be used in several ways. In particular, WSDL can work with UDDI registries in several different ways, depending upon the application needs. WSDL is an XML language for describing Web services as a set of network endpoints that operate on messages.

The purpose of WSDL is to describe your Web services. Businesses exchange WSDL files to understand the other's services. SOAP comes in once you know your partners' services and wish to invoke them. You can think of services as objects that are accessed by SOAP.

Simple Object Access Protocol (SOAP) is based upon XML and is a lightweight (simple) protocol that provides a service-oriented architecture for applications on the Web. A client composes a request and sends the request in a SOAP envelop to a provider. The provider replies within a SOAP response.

12.2 IBM System i5 Initiative for Innovation program

The iSeries Initiative for Innovation (Figure 12-2) is about accelerating IBM efforts on behalf of thousands of ISVs and tools providers, extending our reach from the largest to the smallest, deepening our relationship with them, thereby, broadening our capabilities through new innovative solutions. There are currently:

- ▶ 590+ enhanced applications with over 400 in the pipeline
- ▶ 490+ new solutions through 266 ISVs
- ▶ 100+ tools vendors with over 200 applications
- ▶ 30+ ISVs on System i5 Advisory Board

IBM System i5 Initiative for Innovation		
Application Innovation	Tools Innovation	IBM System i5 Innovation
Free support	Open ecosystem	Technical consultants
Free virtual loaner program	IBM endorsement	Advisory board for System i5 roadmap
Open development roadmap	Technical reviews	Industry enablement
Free educational offerings	Committed partnership	Community building
Free conversion assistance	Competitive advantage	
Accelerate On Demand	Extend Capabilities	Redefine Solutions
IBM Charter for System i5 Innovation		
Investing in the future of System i5 Clients, ISVs, Business Partners		

Figure 12-2 Initiative for Innovation

12.2.1 Application innovation

The iSeries Initiative for Innovation will help accelerate the rate and quality of application innovation by opening up and providing IBM resources in a programmatic, secure, and virtual fashion to assist our ISVs and our clients.

iSeries ISVs will benefit from a step-by-step approach to accessing the latest iSeries server hardware and software and using IBM worldwide technical resources for assessments, education, assistance, and support.

Customized packages that provide step-by-step plans along with new developer roadmaps will help ISVs update and extend their applications without rewriting them. ISVs can then extend this benefit to customers by incorporating additional on demand capabilities into their offerings. Through these efforts, iSeries servers can become the quickest and simplest way to benefit from on demand.

12.2.2 Tools innovation

The IBM System i Developers Roadmap and Tools Innovation Program provides an open approach to the application and solutions ecosystem by providing customers with access to a broad a choice of high-quality tools from IBM and IBM tools partners. Customers can benefit from a broader view and support, as well as be assured that IBM is standing behind its partners to deliver innovative tools and solutions.

New tools and improvements to existing solutions will be driven by customer requirements and requests. IBM will help customer needs to be addressed effectively and efficiently by providing IBM Business Partners with of all types easy access to the Virtual Innovation Center

for Hardware, a one-stop shopping Web site for ideas, partners, software, and tools to meet new and changing business needs.

12.2.3 System i5 Innovation

The speed of innovation and integration for customers can be further accelerated through more pervasive and broader efforts and programs fostered by the IBM iSeries business partners and its ISVs. Through greater access to the technologies and capabilities in Rochester, Minnesota, and through partner-based Industry and Technical Advisory Councils, work on key marketplace insights and critical success factors will be shared in an effort to help meet customers' needs in faster and smarter ways. Technical experts will be assigned to Council members to foster long and productive relationships in support of our mutual customers.

One example in this area is the current activity among Rochester JAVA developers and Independent Software Vendors. Several ISVs are getting help with running their applications on the new with V5R4 32-bit Java JVM.

The sections that follow describe the high-level language-specific enhancements. The corresponding WebSphere-based runtime environment and application development tools are addressed in Chapter 13, "i5/OS-based WebSphere application support" on page 313.

12.3 Control Language

This section covers the enhancements in V5R4 in Control Language (CL).

12.3.1 Subroutine support

Subroutine support allows you to create and manage subroutines within a CL program. A subroutine is identified by the code residing between a SUBR and a ENDSUBR/RTNSUBR command. The subroutine is invoked by a CALLSUBR command.

This command may be placed anywhere within the procedure, including other subroutines, with the exception of a program-level Monitor Message (MONMSG) command. Each Call Subroutine (CALLSUBR) command, when run, places a return address onto the subroutine stack, and the size of the stack can be changed with the use of the Subroutine Stack (SUBRSTACK) parameter of the Declare Process Options (DCLPRCOPT) command. If a globally monitored message causes a GOTO command to be run, the subroutine stack will be reset by the next Call Subroutine (CALLSUBR) command that is run. As for all declare commands (DCL, COPYRIGHT, DCLF, and DCLPRCOPT), the DCLPRCOPT must follow the PGM (Program) command and must precede all other commands in the program. Only one DCLPRCOPT command is allowed by the CL compiler. If more than one are specified, message CPD0323 is sent and the compile fails. The SUBRSTACK parameter specifies how many entries you want to allow on the subroutine stack. Each time a CALLSUBR (Call Subroutine) command is run, an entry is added on the subroutine stack. The entry is removed when a RTNSUBR (Return from Subroutine) or a ENDSUBR (End Subroutine) command is run. The subroutine stack can have multiple entries when CALLSUBR commands are run from within a subroutine. A subroutine can invoke another subroutine or recursively invoke itself.

The name of the subroutine, as used by the CALLSUBR command, is defined by the SUBR parameter on the Subroutine (SUBR) command. The first Subroutine (SUBR) command that is encountered in a procedure also marks the end of the mainline of that procedure. All

commands from this point forward, with the exception of the ENDPGM command, must be contained within a subroutine (in other words, between Subroutine (SUBR) and End Subroutine (ENDSUBR) commands). The Subroutine (SUBR) and End Subroutine (ENDSUBR) commands must be matched pairs, and subroutines cannot be nested. In other words, no SUBR - ENDSUBR nesting of subroutines is allowed. Both the ENDSUBR and RTNSUBR commands can be used to exit a subroutine, and when processed, control is returned to the command immediately following the Call Subroutine (CALLSUBR) command that called the subroutine. Both commands have an optional RTNVAL parameter, which can be anything that can be stored in a CL variable of TYPE(*INT) and LEN(4). If no RTNVAL parameter is defined on the Return Subroutine (RTNSUBR) or End Subroutine (ENDSUBR) command, a value of zero is returned.

In Example 12-1, the Declare Processing Options (DCLPRCOPT) command was used to specify the size of the subroutine stack to be 25. The variable &RTNVAR is used to contain the return value from the subroutine. The Call Subroutine (CALLSUBR) command will transfer control to the subroutine SUBR1, as defined by the SUBR command. If the Return Subroutine (RTNSUBR) command is run, the value of &RTNVAR will be -1. If the End Subroutine (ENDSUBR) command is run, &RTNVAR will equal 0. If no RTNVAL parameter was defined on the Call Subroutine (CALLSUBR) command, the return value from the subroutine would be ignored.

Example 12-1 Subroutine example

```
PGM
DCLPRCOPT SUBRSTACK(25)
DCL VAR(&RTNVAR) TYPE(*INT) LEN(4)
:
CALLSUBR SUBR(SUBR1) RTNVAL(&RTNVAR)
:
SUBR SUBR(SUBR1)
:
RTNSUBR RTNVAL(-1)
:
ENDSUBR ENDPGM
```

12.3.2 Variables

Pointer variables have a fixed length of 16 bytes and contain an address. You define them as shown in Example 12-2, where the pointer variable is initialized to point to the &ANOTHER variable in the program's automatic storage.

Example 12-2 Defining pointer variables

```
DCL &POINTER TYPE(*PTR)
or
DCL &ANOTHER TYPE(*CHAR) LEN(20)
DCL &POINTER TYPE(*PTR) ADDRESS(&ANOTHER)
```

Based variables can be used to map variables passed to the program or manipulate arrays of values. The basing pointer must be set using the ADDRESS keyword on the DCL command (Example 12-3) or with the %ADDRESS built-in function before being used. After the basing pointer is set, the variables will work like local variables.

Example 12-3 First DCL command

```
DCL &POINTER *PTR
DCL &ANOTHER *CHAR LEN(10) STG(*BASED) BASPTR(&PTR)
```

The second DCL command (Example 12-4) declares a character variable that is found at the location addressed by the &POINTER variable.

Defined variables make it easy to manage complex data structures in CL by eliminating the need to substring values out of a large variable. Defined variables can be used to map different parts of the defined variable or the same part of a given variable in different ways.

Example 12-4 Second DCL command

```
DCL &ANOTHER *CHAR LEN(48)
DCL &POINTER *PTR STG(*DEFINED) DEFVAR(&CHAR 17)
```

The second DCL command declares a pointer variable in bytes 17 through 32 of the variable &ANOTHER.

12.3.3 Built-in functions

%ADDRESS or %ADDR can be used to change or test the memory address stored in a CL pointer variable. In a CHGVAR command, you can specify the %ADDRESS function to change the value of a pointer variable. On the IF command, the %ADDRESS function can be specified on the COND parameter to check the value stored in a pointer variable. In Example 12-5, pointer variable &P1 is initialized to the address of the first byte of character variable &C1. Later in the procedure, the pointer is checked using the %ADDRESS function to see if it still points to &C1 and, if not, resets &P1 to the first byte of CL variable &C1 using the %ADDRESS function. See Example 12-5.

Example 12-5 Example

```
PGM
DCL &C1 *CHAR 10
DCL &P1 *PTR ADDRESS(&C1)
:
IF COND(&P1 *NE %ADDRESS(&C1)) + THEN(CHGVAR &P1 %ADDRESS(&C1))
:
ENDPGM
```

%OFFSET or %OFS can be used to store or change the offset portion of a CL pointer variable. In a CHGVAR command, the %OFFSET function can be used to specify the %OFFSET function for the variable (VAR parameter) to set the offset portion of a pointer variable or to specify the %OFFSET function for the value (VALUE parameter) to which the variable is to be changed. In an IF command, the %OFFSET function can be specified in the expression (COND parameter) and is treated like a four-byte unsigned integer value.

In Example 12-6, the pointer variable &P1 has no initial value specified and therefore is initialized to null. The first CHGVAR will store the offset from the null pointer value in integer variable &I1. The command will run without an error, and the value of &I1 will probably be zero, though the offset for a null pointer is not defined. The second CHGVAR command sets &P1 to the address of local character variable &C1. Pointer &P1 now points to byte 1 of variable &C1. The third CHGVAR command will store the offset portion of &P1 into the integer variable &I2 (Example 12-6).

Example 12-6 Third command

```
DCL &C1 *CHAR 30000
DCL &P1 *PTR
DCL &P2 *PTR
DCL &I1 *UINT 4
DCL &I2 *UINT 4
DCL &I3 *UINT 4
CHGVAR &I1 %OFFSET(&P1) /* 1 */
CHGVAR &P1 %ADDRESS(&C1) /* 2 */
CHGVAR &I2 %OFFSET(&P1) /* 3 */
CHGVAR &I3 (%I2+100) /* 4 */
CHGVAR %OFFSET(&P1) &I3 /* 5 */
CHGVAR %OFFSET(&P1) (%OFFSET(&P1)-20) /* 6 */
```

Note that the offset is from the start of the storage for all automatic variables for the current thread and not from the start of automatic storage for the current program or from the start of variable &C1. The fourth CHGVAR command takes the integer value stored in &I2, adds one hundred, and stores the resulting number into the integer variable &I3. The fifth CHGVAR command changes the offset portion of pointer &P1 using integer variable &I3. Pointer &P1 now points to byte 101 of variable &C1. The sixth CHGVAR command will calculate the integer expression for the VALUE parameter by storing the offset portion of pointer &P1 into a temporary integer variable and subtracting 20, and use this calculated value to reset the offset portion of pointer &P1. Pointer &P1 now points to byte 81 of variable &C1.

12.4 ILE RPG programming

This section provides a short summary of enhancements.

12.4.1 Control specifications

The new parameters in DEBUG of *INPUT, *DUMP, and *XMLSAX give more options for debugging aids:

- ▶ *INPUT: Fields that appear only on input specifications are read into the program fields during input operations.
- ▶ *DUMP: DUMP operations without the (A) extender are performed.
- ▶ *XMLSAX: An array of Simple API for XML (SAX) event names is generated into the module to be used while debugging a SAX event.

12.4.2 File specifications

PREFIX allows you to remove characters from all field names in a file.

12.4.3 Definition specifications

OPTION accepts *NULLIND to define whether the null-byte-map should be passed with the parameter. PREFIX allows the removal of characters from a string.

When OPTIONS(*NULLIND) is specified for a parameter, the null-byte map is passed with the parameter, giving the called procedure direct access to the null-byte map of the caller's parameter. Note that the following rules apply when specifying OPTIONS(*NULLIND):

- ▶ ALWNULL(*USRCTL) must be in effect.
- ▶ OPTIONS(*NULLIND) is not valid for parameters passed by value.
- ▶ The only other options that can be specified with OPTIONS(*NULLIND) are *NOPASS and *OMIT.
- ▶ Only variables may be passed as the parameter when OPTIONS(*NULLIND) is specified, and the variable must be an exact match even when CONST is specified.
- ▶ If the parameter is a data structure, the passed parameter must be defined with the same parent LIKEDS or LIKEREC as the prototyped parameter. Furthermore, the null-capability of the prototyped parameter and passed parameter must match exactly.
- ▶ A prototyped data structure parameter can have OPTIONS(*NULLIND) specified whether or not there are any null-capable subfields.
- ▶ If a non-data-structure prototyped parameter is defined with OPTIONS(*NULLIND), the parameter in the procedure interface is defined as null-capable.

12.4.4 Operation codes

In this section we discuss operation codes.

EVAL-CORR

New operation code EVAL-CORR assigns data and null-indicators from the subfields of the source data structure to the subfields of the target data structure. The subfields that are assigned are the subfields that have the same name and compatible data type in both data structures:

```
EVAL-CORR{(EH)} ds1 = ds2
```

For example, if data structure DS1 has character subfields A, B, and C, and data structure DS2 has character subfields B, C, and D, statement EVAL-CORR DS1 = DS2; will assign data from subfields DS2.B and DS2.C to DS1.B and DS1.C. Null-capable subfields in the target data structure that are affected by the EVAL-CORR operation will also have their null-indicators assigned from the null-indicators of the source data structure's subfields, or set to *OFF, if the source subfield is not null-capable. See Example 12-7.

Example 12-7 DS1 subfields

```
// DS1 subfields      DS2 subfields
//  s1 character      s1 packed
//  s2 character      s2 character
//  s3 numeric
//  s4 date           s4 date
//                   s5 character
EVAL-CORR ds1 = ds2;
// This EVAL-CORR operation is equivalent to the following EVAL operations
//  EVAL ds1.s2 = ds2.s2
//  EVAL ds1.s4 = ds2.s4
// Other subfields either appear in only one data structure (S3 and S5)
// or have incompatible types (S1).
```

EVAL-CORR makes it easier to use result data structures for I/O operations to externally described files and record formats, allowing the automatic transfer of data between the data structures of different record formats, when the record formats have differences in layout or minor differences in the types of the subfields.

XML-SAX

The new operation code XML-SAX is shown in Example 12-8.

Example 12-8 New operation code

```
XML-SAX{ (e) } %HANDLER(eventHandler : commArea ) %XML(xmlDocument { : saxOptions
} );
```

XML-SAX initiates a SAX parse for the XML document specified by the %XML built-in function. The XML-SAX operation begins by calling an XML parser, which begins to parse the document. When the parser discovers an event such as finding the start of an element, finding an attribute name, finding the end of an element, and so on, the parser calls the eventHandler with parameters describing the event. The commArea operand is a variable that is passed as a parameter to the eventHandler, providing a way for the XML-SAX operation code to communicate with the handling procedure. When the eventHandler returns, the parser continues to parse until it finds the next event and calls the eventHandler again.

XML-INTO

Example 12-9 shows new operation code XML-INTO.

Example 12-9 XML-INTO

```
XML-INTO{ (EH) } variable %XML(xmlDoc { : options });
XML-INTO{ (EH) } %HANDLER(handler : commArea ) %XML(xmlDoc { : options });
```

XML-INTO reads the data from an XML document in one of two ways:

- ▶ Directly into a variable.
- ▶ Gradually into an array parameter that it passes to the procedure specified by %HANDLER. Various options may be specified to control the operation.

The first operand specifies the target of the parsed data. It can contain a variable name or the %HANDLER built-in function.

The second operand contains the %XML built-in function specifying the source of the XML document and any options to control how the document is parsed. It can contain XML data or it can contain the location of the XML data. The doc option is used to indicate what this operand specifies. See Example 12-10.

Example 12-10 Example

```
// Data structure "copyInfo" has two subfields, "from"
// and "to". Each of these subfields has two subfields
// "name" and "lib".
// File cpyA.xml contains the following XML document
// <copyinfo>
//   <from><name>MASTFILE</name><lib>CUSTLIB</lib></from>
//   <to><name>MYFILE</name><lib>*LIBL</lib>
// </copyinfo>
xml-into copyInfo %XML('cpyA.xml' : 'doc=file');
// After the XML-INTO operation, the following
// copyInfo.from .name = 'MASTFILE' .lib = 'CUSTLIB'
// copyInfo.to   .name = 'MYFILE'   .lib = '*LIBL'
```

12.4.5 Built-in functions

In this section we discuss built-in functions.

%HANDLER

%HANDLER is used to identify a procedure to handle an event or a series of events.

%HANDLER does not return a value, and it can only be specified as the first operand of XML-SAX and XML-INTO. The first operand, handlingProcedure, specifies the prototype of the handling procedure. The return value and parameters specified by the prototype must match the parameters required for the handling procedure. The requirements are determined by the operation that %HANDLER is specified for.

The second operand, communicationArea, specifies a variable to be passed as a parameter on every call to the handling procedure. The operand must be an exact match for the first prototyped parameter of the handling procedure, according to the same rules that are used for checking prototyped parameters passed by reference. The communication-area parameter can be any type, including arrays and data structures.

%XML

%XML is used as the second operand of the XML-SAX and XML-INTO operation codes to specify the XML document to be parsed, and the options to control how the document is parsed. %XML does not return a value, and it cannot be specified anywhere other than for the XML-SAX and XML-INTO operation codes. The first operand specifies the document to be parsed. It can be a constant or variable character or UCS-2 expression containing either an XML document or the name of a file containing an XML document. The second operand

specifies options that control how the XML document is to be interpreted and parsed. It can be a constant or variable character expression.

12.4.6 Free format SQL

Using free format for SQL statements in ILE RPG programs makes it much easier to code these statements. The only requirement is to start every SQL statement with EXEC SQL and to end it with a semicolon (;). An SQL statement can span over several lines.

Bracketed comments (*/*...*/*) are allowed within embedded SQL statements between positions 8 through 80 and whenever a blank is allowed, except between the keywords EXEC and SQL. Comments can span any number of lines. Single-line comments (*//*) can also be used.

SQL statements can be contained on one or more lines. To continue an SQL statement across multiple lines, the SQL statement can be split wherever a blank is allowed. The plus sign (+) can be used to indicate a continuation of a string constant. The literal continues with the first nonblank character on the next line.

Example 12-11 Fixed form

```
C/EXEC SQL UPDATE DEPARTMENT
C+ SET MANAGER = :MGRNUM
C+ WHERE DEPTNO = :INTDEP
C/END-EXEC
```

Example 12-12 Free form

```
EXEC SQL UPDATE DEPARTMENT
  SET MGRNO = :MGR_NUM
  WHERE DEPTNO = :INT_DEP;
```

12.4.7 Rules for naming in embedded SQL are updated

SQL statements can be contained on one or more lines. To continue an SQL statement across multiple lines, the SQL statement can be split wherever a blank is allowed. The plus sign (+) can be used to indicate a continuation of a string constant. The literal continues with the first nonblank character on the next line.

Naming restrictions:

- ▶ Do not use host variable names or external entry names that begin with the characters SQ, SQL, RDI, or DSN. These names are reserved for the database manager.
- ▶ The length of host variable names is limited to 64.
- ▶ The names of host variables must be unique within the program. The one exception is that if a stand-alone field, parameter, or both are defined exactly the same as another stand-alone field, parameter, or both, the duplicated name is accepted. (New in V5R4.)
- ▶ If a host variable is a duplicated name and does not belong to the exceptional category mentioned in the previous item, but does have the same type, the precompiler issues SQL0314 as a severity 11 error instead of its normal severity of 35. If you want to ignore these severity 11 errors, change the GENLVL parameter value on the CRTSQLRPGI command to be 11 or higher. (New in V5R4.)

12.5 ILE COBOL programming

XML support has been enhanced. A new statement, XML GENERATE, converts the content of COBOL data records to XML format. XML GENERATE creates XML documents encoded in Unicode UCS-2 or in one of several single-byte EBCDI or ASCII CCSIDs.

Nonnumeric literals can be null-terminated. They can be used anywhere a nonnumeric literal can be specified, except that null-terminated literals are not supported in ?ALL literal? figurative constants.

*NOCOMPRESSDBG/*COMPRESSDBG specifies whether listing view compression should be performed by the compiler when DBGVIEW option *LIST or *ALL is specified.

Table 12-1 lists the new intrinsic functions.

Table 12-1 New intrinsic functions

Function	
DISPLAY-OF	This function returns an alphanumeric character string consisting of the content of argument-1 converted to a specific code page representation. The type of the function is alphanumeric.
NATIONAL-OF	This function returns a national character string consisting of the UCS-2 representation of the characters in argument-1. The type of the function is national.
TRIM	This function returns the given string with any leading and trailing blanks removed, or the given string with any leading and trailing specified characters removed.
TRIML	This function returns the given string with any leading blanks removed, or the given string with any leading specified characters removed.
TRIMR	This function returns the given string with any trailing blanks removed, or the given string with any trailing specified characters removed.

12.6 Java

This section provides an in-depth look at new Java implementations in the Java Virtual Machine (JVM) and the Java Developer Kit (JDK™).

12.6.1 IBM technology for Java Virtual Machine

In OS/400 V4R2, shipped in 1998, a Java Virtual Machine (JVM) was included, running in 64-bit mode on 64-bit hardware. This JVM was implemented into the System Licensed Internal Code (SLIC). It had unique features such as Direct Execution. This allowed creation of a hidden object, associated with a .class or .jar file, with the CRTJVAPGM command. It also allowed automatic and concurrent garbage collection. Both of these functions boosted performance.

This 64-bit JVM has displayed constant leadership in scalability and is the better choice on high-end systems, where the combination of sufficient system memory and high application memory requirements exist. It also coexists with then new 32-bit JVM, which is better suited for small systems, delivering up to as much as 20% improved performance. The 32-bit JVM also allows you to standardize Java and WAS-based offerings across all eServer platforms.

32-bit versus 64-bit architecture

i5/OS V5R4 delivers two JVM and SDK environments:

- ▶ Java™ 2 Software Development Kit, Standard Edition
- ▶ Classic, running in 64-bit mode
- ▶ 32-bit IBM Technology for Java VM runs in 32-bit mode

Both of these include Java Platform, Enterprise Edition (Java EE) Version 5.0 (also commonly referred to as JDK 1.5) support. The 64-bit JVM also supports JDK levels 1.3 and 1.4.

Using V5R4 JVM support, most job logs will include a message indicating whether 32-bit JVM or 64-bit JVM is being used. For example, most i5/OS system (jobs) programs written in Java, such as jobs with prefix names of QJVACMDSRV, QYPSJSVR (Management Central server), and the IBM-supplied instance sysinst/admin (job shows as ADMIN/QEJBSVR) in subsystem QWAS6, will show a job log message indicating the 64-bit JVM is being used.

JVAB56D - Java Virtual Machine is Classic

WAS V6.0 does not formally support an option to select between the 32-bit or 64-bit JVM. Plans call for this option to be available with WAS V6.1.

Figure 12-3 depicts the i5/OS Java runtime environment when running either JAVA 32-bit JVM or Java 64-bit JVM programs.

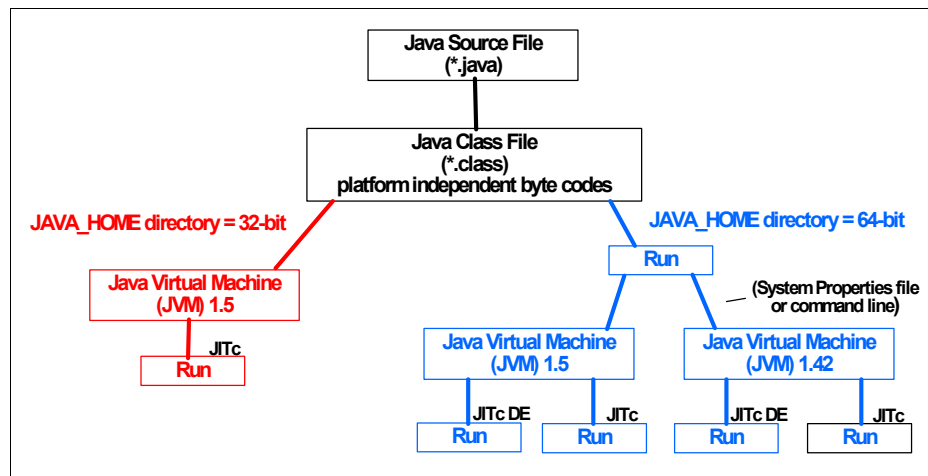


Figure 12-3 Runtime architecture: Java 32-bit, Java 64-bit

The first two blocks are represent Java source files and Java class files that are common across all Java environments on all supporting operating systems.

The red color is used to identify the 32-bit JIT only (JiTc) runtime environment available under V5R4 i5/OS. The blue color is used to identify the 64-bit JiTc or JiTc_DE (Direct Execution) environments available under V5R4 i5/OS (and earlier i5/OS releases).

You use either runtime Add Environment Variable statement or, for 64-bit Java, an available SystemDefault.properties file to specify which SDK level (1.5 or 1.42) to use. 32-bit JVM unconditionally uses JIT and JDK 1.5.

Through i5/OS V5R4, when using the 64-bit JVM, Direct Execution is the default.

32-bit IBM Technology for JVM

To switch to 32-bit you need to have the IBM Developer Kit for Java (5722-JV1) Option 8 (J2EE™ 5.0 32 bit) installed and run the following command:

```
ADDENVVAR ENVVAR(JAVA_HOME) VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit')
```

RUNJAVA *VERSION displays the output shown in Example 12-13.

Example 12-13 Output

```
java version "1.5.0"  
Java(TM) 2 Runtime Environment, Standard Edition (build jclap32dev)  
IBM J9 VM (build 2.3, J2RE 1.5.0 IBM J9 2.3 OS400 ppc-32 (JIT enabled)  
J9VM - 20051027_03723_bHdSMR  
JIT - 20051027_1437_r8  
GC - 20051020_AA  
JCL - jclap32dev
```

32-bit implements a standard IBM JVM 1.5 runtime environment and SDK. This is specifically focused on applications for smaller systems. It has less impact on memory utilization. This simple implementation can be used with minimum code changes for the majority of the applications.

This affects management of storage references. If you want to use IBM Technology for Java and have programs that use native methods, you must compile these programs with teraspace storage enabled. Since this is not the default setting, it is likely that you need to recompile.

This is necessary because the Java object is in i5/OS PASE storage, which is mapped on top of teraspace storage, and a teraspace storage pointer is returned. Also, the JNI function `NewDirectByteBuffer` has a pointer parameter, which is the address of the storage for the direct byte buffer. This pointer must point to the area in teraspace storage that is part of the i5/OS PASE address space.

The easiest way to obtain this storage is to call `Qp2malloc`. If you have not compiled your program with teraspace storage enabled and attempted to run the native method with IBM Technology for Java, you will receive the escape message MCH4443 (Invalid storage model for target program LOADLIB).

IBM Technology for Java does not support adopted authority for Java programs. Garbage collection is similar as on most other platforms for the 32-bit JVM, while it remains unchanged for the 64-bit JVM.

Comparing the 32-bit JVM with the 64-bit, one must include the non functional requirements such as scalability and predictability needs for the application. It is clear that a 32-bit JVM needs less space to address objects, but it is also evident that the addressability is smaller. So, as on any platform, one needs to closely examine the application architecture, its functionality, and the non functional requirements if asked to select one of the implementation ways.

Apart from the teraspace addresses needed by native methods when migrating to the 32-bit JVM, there are some other remarks regarding performance. ILE native method calls will typically be slower, but are currently a focus item within development. However, since i5/OS PASE native methods use the same address space as the 32-bit JVM, they will execute faster. As always, the overall performance results will highly depend on the application being used.

Just In Time (JIT) compiler

Direct Execution is not applicable to the 32-bit JVM — only JIT or interpreted mode. Starting with V5R3, JIT is always recommended for best performance. This makes the use of optimization levels on the CL commands associated with the JVM no longer relevant. Also, by default, WebSphere Application Server runs in JIT mode.

To set the compiler to use JIT:

1. From a command-line prompt on your iSeries server, add the environment variable by using the Add Environment Variable (ADDENVVAR) command. Then run your Java program using the Run Java (RUNJAVA) command or JAVA command. For example, use:

```
ADDENVVAR ENVVAR (JAVA_COMPILER) VALUE(jitc) JAVA CLASS(Test)
```
2. Set the java.compiler system property on the iSeries command line. For example, enter:

```
JAVA CLASS(Test) PROP((java.compiler jitc))
```
3. Set the java.compiler system property on the Qshell Interpreter command line. For example, enter:

```
java -Djava.compiler=jitc Test
```

A quick checklist: 32 bit or 64 bit

The question about which JVM to run is an important one. If you are in doubt, your IBM representative or Business Partner can help analyze your Java environment. However, we have noted two key items that are strong indicators of your needs:

- ▶ Does your application need a Java heap much greater than approximately 2 GB for higher performance?
- ▶ Does your application use computationally intensive algorithms for statistics, security, encryption, and so on, which can benefit from high precision computation support?

If you answered no to these questions, then 32-bit could be for you. If you answered yes to even one, we would advise use of the 64-bit JVM.

SPECjbb2000 benchmark testing results

Standard Performance Evaluation Corporation (SPEC) is a non-profit corporation. They were formed to establish, maintain, and endorse standardized benchmarks. Additionally, they review and publish submitted results.

The benchmark SPECjbb2000 evaluates the performance of servers running typical Java business applications by simulating an order processor for wholesale suppliers. It is implemented as a three-tier client/server system with emphasis on the middle tier.

SPEC and the benchmark name SPECjbb2000 are registered trademarks of the Standard Performance Evaluation Corp. (SPEC).

Additional information and current results showing Java on i5/OS can be found at:

<http://www.spec.org/jbb2000/results/>

12.6.2 What is new for IBM Developer Kit for Java

In this section we discuss what is new for IBM Developer Kit for Java.

Java Virtual Machine Tool Interface

The Java Platform Debugger Architecture (JPDA) consists of the JVM Debug Interface, the JVM Tool Interface, and the Java Debug Wire Protocol. All these parts of the JPDA enable any front end of a debugger that uses the JDWP to perform debugging operations. The debugger front end can either run remotely or run as an iSeries application.

JVMTI replaces JVMDI and JVMPI, and adds new functions. JVMTI was added as part of J2SE™ 5.0. In future releases, the JVMDI and JVMPI interfaces will no longer be offered, and JVMTI will be the only option available.

Previously delayed PTFs for the JVM can now be applied when no JVM is active. You should use the Display Java Virtual Machine Jobs (DSPJVMJOB) to find out if there are any active.

A service program called QJVAJVMTI, which resides in the QSYS library, supports the JVMTI functions.

Java Cryptography Extension (JCE) 1.2

The Java Cryptography Extension 1.2 is a standard extension to the Java 2 Software Development Kit (J2 SDK), Standard Edition. The JCE implementation on an iSeries server is compatible with the implementation of Sun Microsystems™, Inc.

In order to understand this information, you should be familiar with the general documentation for the JCE extensions. See the Sun JCE documentation for more information about JCE extensions.

Refer to the iSeries V5R4 Information Center under Programming → Java for more documentation that covers the unique aspects of the i5/OS implementation.

Java Platform Standard Edition

Support is now also provided for J2SE5.0, together with previous versions. See Table 12-2.

Table 12-2 Support for properties for J2SE Version 5.0

Option	JDK	Java.home	Java version
5	1.3	/QIBM/ProdData/Java400/jdk13/	1.3
6	1.4	/QIBM/ProdData/Java400/jdk14/	1.4
7	1.5 - J2SE 5.0	/QIBM/ProdData/Java400/jdk15/	1.5

New Java tools

Some new Java tools are available only with JDK 1.5 and subsequent versions. They are available using the Qshell Interpreter, and come with this release:

- ▶ The apt tool processes program annotations (JSR 175).
- ▶ The pack200 and unpack200 tools compresses or decompresses a jar file into or from a pack200 file. These are highly compressed files that can be directly deployed, saving bandwidth and reducing download time.

New Java API libraries

Some new APIs are available with JDK 5.0. They are meaningful to Java experts. For more information search <http://www.sun.com> or <http://www.google.com> or <http://www.ask.com> or <http://www.whatis.com>. We list most of the APIs here. JSR is Java Specification Requirement.

- ▶ JSR 003: The JMX™ 1.0 specification. Packages: javax.mangement.
Java Management Extensions (JMX) is a set of specifications for application and network management in the J2EE development and application environment. JMX defines a method for Java developers to integrate their applications with existing network management software by dynamically assigning Java objects with management attributes and operations.
JMX facilitates the centralized management of managed objects (called MBeans), which act as Java wrappers for applications, services, components, or devices in a distributed network.
- ▶ JSR 013: Additions to Java math class that include improved arithmetic operations using BigDecimal.
- ▶ JSR028: Java SASL packages: java.security.sasl.
Simple Authentication and Security Layer (SASL) is an Internet standard (RFC 2222) that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged but does not itself specify the contents of that data. It is a framework into which specific authentication mechanisms that specify the contents and semantics of the authentication data can fit.
- ▶ JSR114; Java Rowset implementations. This more completely specifies rowset functions.
Rowset is a Java-Bean Compliant Object that Encapsulates Database Access. Using Rowset you use a single object rather than the typical three objects: Connection (for creating connection), Statement (for query execution), ResultSet (to get the result).
- ▶ JSR 160: JMX remote API,V1.0.
- ▶ JSR 163: Java Platform Profiling Architecture JVMTI. This is a replacement for JVMPPI.
The JVMTM Tool Interface (JVM TI) replaces Java Virtual Machine Profiler Interface (JVMPPI) and the Java Virtual Machine Debug Interface (JVMDI) interfaces. JVMTI is a new native programming interface for use by tools. It provides both a way to inspect the state and to control the execution of applications running in the Java virtual machine (JVM).
- ▶ JSR 166: Concurrency utilities. Packages: java.util.concurrent.
- ▶ JSR 174: Monitoring and management specification for the JVM.
- ▶ JSR 200: Network transfer format for Java Archives.
- ▶ JSR 206: Java API for XML Processing (JAXP) 1.3.
- ▶ JSF 204: Unicode Supplementary Character support.

12.6.3 What is new for the IBM Toolbox for Java

Toolbox no longer ships the x4j400 (IBM XML parser).

The new toolbox can be installed on i5/OS V5R2 systems and later and includes enhancements listed in the following sections.

Java Database Connectivity (JDBC) support

JDBC support includes several *increased maximums* as well as other functional enhancements.

2 MB SQL statement support

SQL statement length has been increased from 65,535 to 2,087,152 (1,048,576 character) bytes. Prior to V5R4, the limit on SQL statement size was 65,535 bytes. This corresponds to 65,535 characters when the statement text is represented using a single-byte CCSID, and 32,767 characters when the statement text is represented using a double-byte CCSID. Some customers, particularly those using applications that automatically generate SQL statements, were affected by this limit.

In V5R4, the i5/OS statement size limit has been increased to two megabytes, or 2,097,152 bytes. The IBM Toolbox for Java JDBC driver always sends statement text in 2 byte Unicode. Therefore, the maximum statement length in characters will be one megabyte or 1,048,576 characters.

128 byte column name support

Starting with V5R4, the database will support column names up to 128 bytes for SQL tables. Prior to V5R4, column names up to 30 bytes were supported. The IBM Toolbox for Java JDBC driver will provide these possibly longer names to its users.

There is one exception where 128-byte column names will not be returned. When local package caching is used and column names exceed 30 characters, the server will return the column names as the system column name.

Database host server trace support

A new option was added to the Toolbox for Java JDBC driver to turn on database host server tracing. To support this feature, option 64 was added to the server trace connection property.

eWLM Correlator support

The IBM Toolbox for Java will accept an IBM Enterprise Workload Manager (eWLM) correlator and pass it on to the host as a connection attribute correlator for use with the Application Response Measurement (ARM) APIs. This correlator can be sent to the host at any time after a connection is made using the `setDB2eWLMCorrelator` method in the `AS400JDBCConnection` class. eWLM correlators require i5/OS V5R3 or later servers. This request is ignored when running on OS/400 V5R2 or earlier servers.

New and enhanced classes

New and enhanced classes are:

- ▶ `BidiConversionProperties`
Provides a set of properties that can be used to control the conversion of character set data
- ▶ `CallStackEntry`
Represents an entry in the call stack of a specific thread of a server job
- ▶ `IFSFileReader` and `IFSFileWriter` reads
Writes character files in the IFS and replaces `IFSTextFileOutputStream`
- ▶ `IFSSystemView`
Provides a gateway to the i5/OS IFS, for use when constructing `javax.swing.JFileChooser` objects

- ▶ iSeriesNetServer
Represents the NetServer service on a server and replaces the NetServer class
- ▶ SaveFile
- ▶ SignonHandler (interface)
- ▶ Subsystem
- ▶ Enhancements to the classes IFSFile, the JarMaker and AS400ToolboxJarmaker, the FTP and AS400FTP

12.7 Structured Query Language (SQL)

There have been a number of enhancements to SQL.

12.7.1 ISO timestamp format

Support for the ANSI Timestamp format will be a welcome addition for SQL developers, especially those porting applications to the iSeries. In past releases, programmers would have had to change to a different format. Supported formats are:

- ▶ ANSI/ISO SQL Standard ('2006-01-31 08:30:00.010000')
- ▶ IBM SQL ('2006-01-31-08.30.00.010000')
- ▶ 14- character form ('20060131083000')
- ▶ ISO ('2006-01-31 08:30:00.010000')

The ANSI/ISO SQL Standard and ISO formats show dash separators for the date fields, a blank separating the date and time, and two colons separating the hour from minutes and minutes from seconds. Many years ago the original ISO specification timestamp standard was three hyphens separating the three date fields and the hours position, with three periods separating hours, minutes, seconds, and nanoseconds. There was no blank separating date from time. This original standard, implemented by most databases, was changed after that original three hyphens and three periods specification.

The 5250 interactive Start SQL (STRSQL) command, iSeries Navigator Run SQL scripts, and Display Physical File Member *DSPFM) command interfaces may exhibit some differences in the formats viewed or actually entered into SQL run statements using the two different ISO formats.

The Run SQL Scripts interface uses the newer ISO format. The 5250 interactive STRSQL interface has intentionally not been updated to this newer ISO format.

12.7.2 Special register

Many distributed and Web-based applications utilize connection pooling to improve performance. Often pooled connections are set up to use a shared, generic user ID to set up the connection. The result is that when monitoring data access, most diagnostic tools will only report the shared user instead of the real user's identity.

With V5R4, the new SET SESSION AUTHORIZATION statement can be executed after accessing a pooled connection to identify the actual identity of the end user. While this makes database auditing more accurate, the downside is that many DB2 resources such as Open Data Paths (ODPs) cannot be reused when the user ID for a connection is changed. Thus, executing this new statement can negate some of the performance benefits of using a pooled connection. New special registers were also added to allow an application to retrieve the

current session user as well as the shared user ID used to set up the connection with the system user register.

SET SESSION AUTHORIZATION is available on V5R3 via the latest DB Group PTF.

The SESSION_USER special register specifies the runtime authorization ID at the current server. The data type of the special register is VARCHAR(128). The initial value of SESSION_USER for a new connection is the same as the value of the SYSTEM_USER special register.

The SYSTEM_USER special register specifies the authorization ID that connected to the current server. The data type of the special register is also VARCHAR(128). Its value can be changed by invoking the SET SESSION AUTHORIZATION statement, which also changes the name of the user profile associated with the current thread. If the authorization name specified on the statement is different from the value in the SYSTEM_USER special register, the privileges held by the authorization ID of the statement must include the system authority of *ALLOBJ. No authorization is required to execute this statement if the authorization name specified on the statement is the same as the SYSTEM_USER special register.

Full select in a subquery and scalar-fullselect

A scalar-fullselect is a fullselect, enclosed in parentheses, that returns a single result row and a single result column. If the result of the fullselect is no rows, then the null value is returned. An error is returned if there is more than one row in the result.

12.7.3 Recursive common table expressions and recursive views

Recursive common table expressions and common views allow the use of hierarchical relations stored in a single table (for example, Bill of Materials structure). See Figure 12-4.

Level	Parent	Description	Component	Description	Quantity in Assembly
1	11L5441	Power Assembly ...	09H4112	Power Supply ...	1
1	11L5441	Power Assembly ...	54K6942	Shipping Instruction...	1
2	09H4112	Power Supply ...	1789225	Fuse 12 AMP ...	2
2	09H4112	Power Supply ...	1586237	Power Cord 6 ft. ...	2
2	09H4112	Power Supply ...	18H1588	Support Bracket ...	1
2	09H4112	Power Supply ...	15H1200	CIF Tool	1
2	09H4112	Power Supply ...	578K211	Slider Long sleeve ...	2
2	09H4112	Power Supply ...	57G3289	Rail	4
3	1586237	Power Cord 6 ft. ...	1584451	Power Cord Retaine...	2
3	18H1588	Support Bracket ...	47P5281	Screw 8 mm ...	12
3	18H1588	Support Bracket ...	47J5213	Nut 8 mm	12
3	15H1200	CIF Tool	15H2900	Language Selector ...	1
3	578K211	Slider Long sleeve...	15H1200	CIF Tool	1
3	578K211	Slider Long sleeve...	09H2278	Filler Panel	1
3	578K211	Slider Long sleeve...	09K2557	Filler Panel Spring ...	4
4	15H1200	CIF Tool	15H2900	Language Selector ...	1
4	09H2278	Filler Panel	21L5554	Front Cover	1
4	09H2278	Filler Panel	29L5454	Front Cover Clip ...	4

Figure 12-4 Recursive common table expression

In the example in Figure 12-4 on page 291 we use two tables. The first is ITEMMAST, which defines the item. The second is PRODSTRUCT, which defines the product structure on a parent-child basis. We use these to create a view, called EXPLPROD. This view combines the parts structure with the description, and optionally other fields.

Example 12-14 Table definition

```

CREATE TABLE PRODMAST.ITEMMAST (
  PARTNBR CHAR(12)NOT NULL, PARTDESC CHAR(30) NOT NULL, PARTSHORT CHAR(10) NOT
  NULL,
  PARTBR CHAR(4) NOT NULL, PARTCLASS CHAR(3) NOT NULL, PARTSCLAS CHAR(3) NOT NULL,
  PARTREF CHAR(12) NOT NULL, PARTPACK CHAR(8) NOT NULL, PARTOWNER CHAR(6) NOT
  NULL)
CREATE TABLE PRODMAST.PRODSTRUCT (
  PARENT CHAR(12) NOT NULL, CHILD CHAR(12) NOT NULL, CHILDQTY CHAR(5) NOT NULL)
The view looks as follows:
CREATE VIEW PRODMAST.EXPLPROD (
  PARENT , PDESC , CHILD , CDESC , CHILDQTY )
AS
  SELECT
    STRUCT.PARENT, PARTP.PARTDESC AS PDESC, STRUCT.CHILD, PARTC.PARTDESC AS CDESC,
    STRUCT.CHILDQTY
  FROM
    PRODMAST.ITEMMAST PARTP, PRODMAST.ITEMMAST PARTC, PRODMAST.PRODSTRUCT STRUCT
  WHERE
    PARTP.PARTNBR = STRUCT.PARENT AND
    PARTC.PARTNBR = STRUCT.CHILD
WITH rpl (level, parent, pdesc, child, cdesc,childqty) AS
  ( SELECT 1, root.parent, root.pdesc, root.child, root.cdesc, root.childqty
    FROM explprod as root
    WHERE root.parent = '11L5441'
    UNION ALL
    SELECT level+1, low.parent, low.pdesc, low.child, low.cdesc, low.childqty
    FROM rpl high, explprod low
    WHERE high.child = low.parent
  )
SELECT level as "Level", parent as "Parent", pdesc as "Description", child as
"Component", cdesc as "Description", childqty as "Quantity in Assembly"
  FROM rpl;

```

The net result of this support is that you now can use a single select command to view an exploded BOM without having to go through creating result tables that need to be combined in order to give the same view, thus limiting complexity and possible inconsistency with tables in a live environment. It also should allow you to spend fewer system resources needed to build to individual result tables, thus increasing overall system throughput.

In-depth example

Common table expressions have been supported by DB2 UDB for iSeries since V4R4. With V5R4, CTEs have become an even more useful and powerful SQL feature, with the ability to perform recursive processing. A common table expression that includes a recursive reference is known as a recursive common table expression (RCTE). Database tables can contain rows that have an inherent relationship to other rows in the same table such as organizational hierarchies, bill-of-materials, or travel connections that are naturally navigated with recursive algorithms.

Recursive queries are useful for navigating tables where rows are inherently related to other rows in the same table, as in a bill of materials, or a bill of lading, or as in organizational hierarchies.

Table 12-3 RCTE

Annotated example of a recursive common table expression	
Initialize the query	WITH emp_list (level, empid, name) AS (SELECT 1, empid, name FROM emp WHERE name = 'Carfino'
Recursive reference to the next level	UNION ALL SELECT o.level + 1, next_layer.empid, next_layer.name FROM emp as next_layer, emp_list o WHERE o.empid = next_layer.mgrid)
Start query, return results	SELECT level, name FROM emp_list

From what we show in Table 12-3, you can see that an RCTE is composed of three different phases. We cover each of these phases in more detail to explain the capabilities of Recursive Expressions.

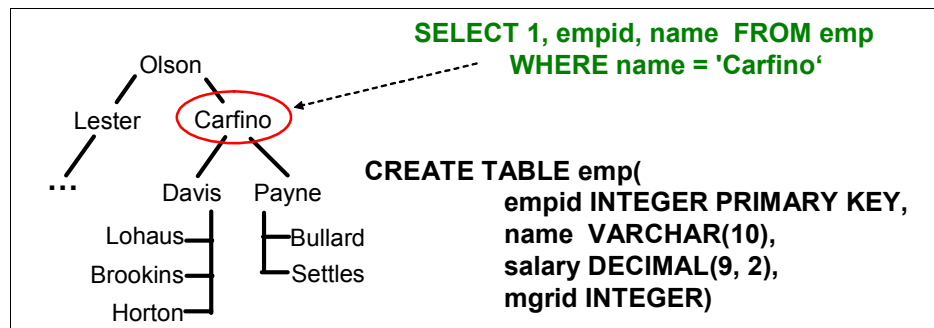


Figure 12-5 Initialization phase

In this RCTE example, the initialization phase (Figure 12-5) needs to help identify all of the employees that work directly for Carfino. Thus, the first step is retrieving Carfino’s employee ID so that the next phase can find all the employees that have a manager ID value that references Carfino’s employee ID.

A constant value of 1 is returned on this seeding SELECT statement to help show in the final report the organization level or depth of an employee in Carfino’s organization chain. The initialization statement will only be run once by DB2 UDB and not executed on any recursive executions of the emp_list table expression. DB2 UDB requires that the initialization statements be coded at the beginning of the RCTE.

The blue SELECT statement (Figure 12-6) is the recursive portion of the RCTE since this statement contains a recursive reference to the emp_list table expression on the From clause.

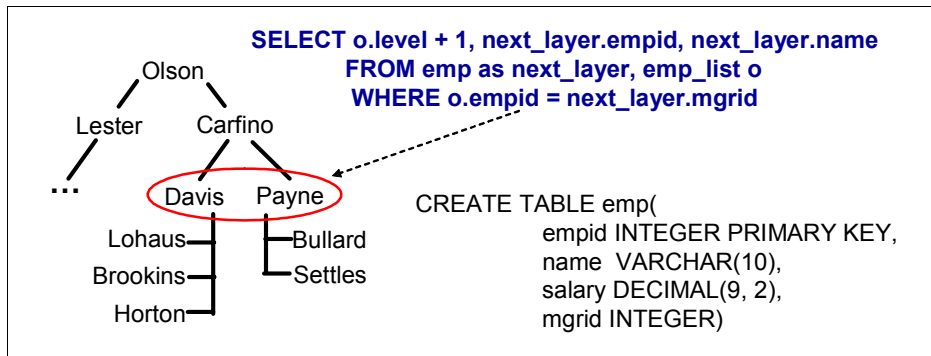


Figure 12-6 Recursive pass number one

The first time through the recursive phase of the query, the SELECT statement will produce a list of the employees that have Carfino as a manager, Davis & Payne, since o.empid will contain Carfino's employee number, and that value is being joined and compared to the manager ID for all employees.

The join column on the recursive select statement always needs to join to the seed value produced by the initialization statement (that is, the empid column). This join condition is how the RCTE feeds itself with a list of employees to investigate on the next recursive execution. Also notice that the level column is incremented to track the recursion depth within the organization on each execution (see Figure 12-7). The join column on the recursive select always needs to join to the seed value produced by prior execution.

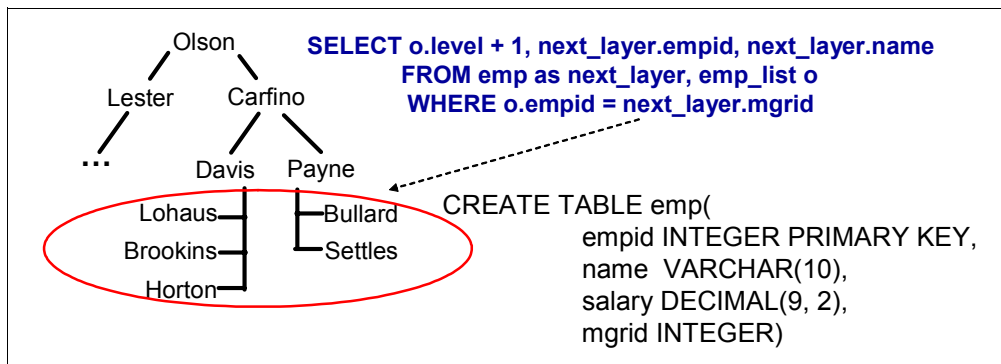


Figure 12-7 Recursive pass number two

The last execution of this recursive expression (Figure 12-8) will then produce a list of the employees that have Carfino's direct reports, Davis & Payne, as managers. The recursive traversals of the RCTE are stopped when the statement encounters employees that are not managers — meaning their employee numbers are never referenced in the mgrid column.

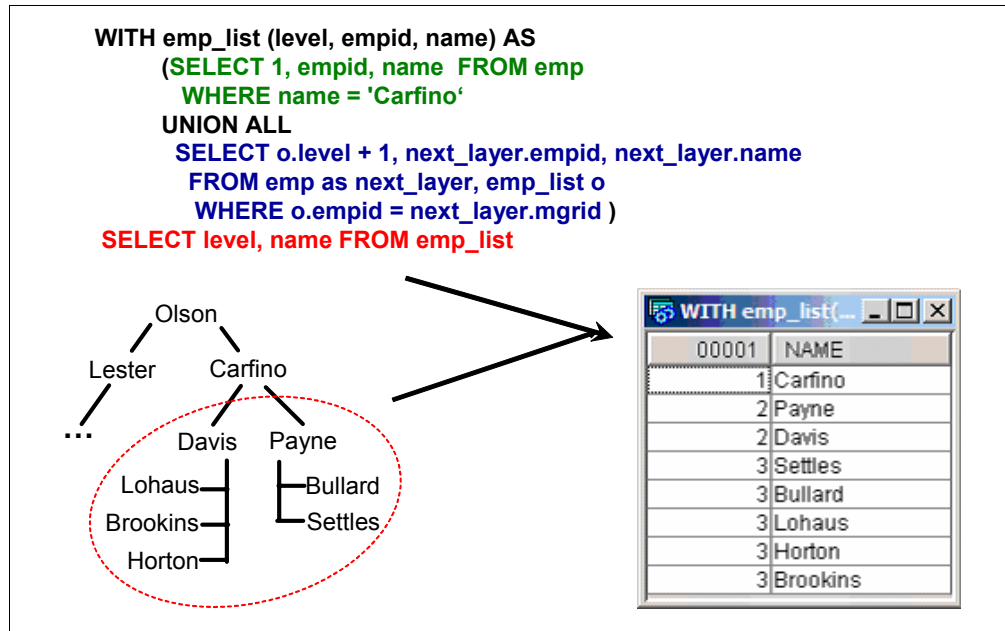


Figure 12-8 Final results

To run the RCTE and produce a result set, the main SELECT statement (in red) is needed to start the recursive processing of the organizational hierarchy.

This SELECT statement references the emp_list common table expression. This reference causes the seeding SELECT statement to be executed first, and then control is passed onto the recursive phase, which then constructs the final result set displayed on the chart. The UNION ALL operation is used to combine the results from the initialization and recursion phases.

In the final result set, you will notice that employees are not listed under their manager. By default, DB2 UDB will always search and process the data in breadth first order. This means that DB2 UDB traverses across one layer of the organization before going onto the next layer.

Recursive considerations

A depth first search (Figure 12-9) is available with the DB2 UDB for iSeries RCTE support in V5R4 because it is common for applications to want the data in this order. In a depth first search, DB2 UDB would take one of Carfino's immediate employees and first try to process all of the employees that work for that manager before processing another one of Carfino's direct reports. The processing order is specified with the SEARCH clause demonstrated in the query and results contained on the chart. It should be obvious that the depth first result set on this chart is ordered in a way that naturally matches the relationships in an organization chart.

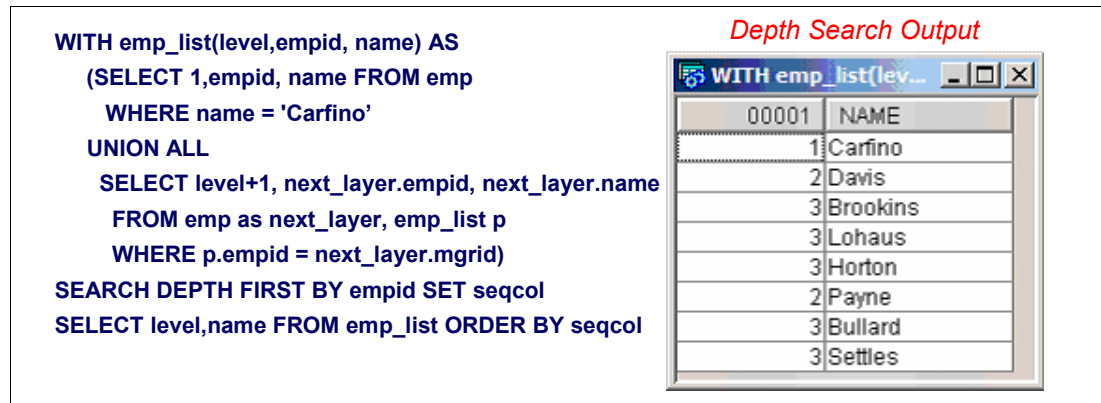


Figure 12-9 Considerations

The first thing to notice in this depth example is that the SEARCH clause needs to use the same column that is specified on the recursive join reference — in this case, empid. In addition to this navigation column, the SEARCH clause also contains a set keyword with another column name, SeqCol. DB2 UDB uses this column to keep track of the order in which it visited the rows on this SQL request. The sequencing column specified on the SET is very important because this same column name must also be specified on the ORDER BY clause. The column name itself is not important, so you are free to choose any name.

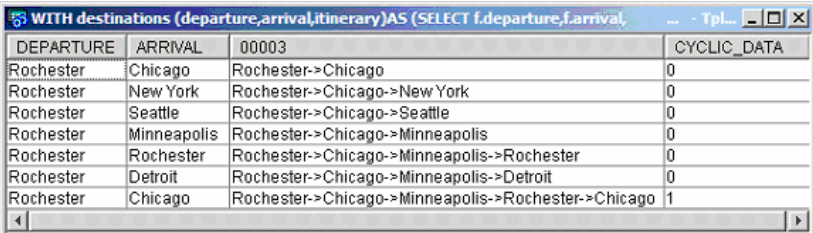
If this sequencing column is not referenced on the ORDER BY clause, then the entire SEARCH clause is ignored by DB2 UDB. Remember that with SQL, if you want your data return in a specific order, then an ORDER BY clause is the only way to guarantee that ordering.

There is one part of a recursive common table expression that has not been discussed yet, and that is the cycle clause. There are some cases where your tables may contain recursive data relationships that are cyclic or never ending. In an organization hierarchy, you do not have to worry about getting into a never-ending cycle because you do not have many company presidents that report to entry-level employees at the bottom of the organization. On the other hand, let us consider a table such as a flights table used by an airline that contains the air flight combinations for a set of cities. See Figure 12-10.

```

WITH destinations (departure, arrival, itinerary) AS
(SELECT f.departure, f.arrival, CAST (f.departure ||'-'|| f.arrival AS VARCHAR(200))
 FROM flights f
 WHERE f.departure ='Rochester'
 UNION ALL
 SELECT r.departure, b.arrival, CAST (r.itinerary ||'-'|| b.arrival AS VARCHAR(200))
 FROM destinations r,flights b
 WHERE r.arrival =b.departure )
 CYCLE arrival SET cyclic_data TO '1' DEFAULT '0 '
 SELECT departure, arrival, itinerary, cyclic_data FROM destinations

```



DEPARTURE	ARRIVAL	itinerary	CYCLIC_DATA
Rochester	Chicago	Rochester->Chicago	0
Rochester	New York	Rochester->Chicago->New York	0
Rochester	Seattle	Rochester->Chicago->Seattle	0
Rochester	Minneapolis	Rochester->Chicago->Minneapolis	0
Rochester	Rochester	Rochester->Chicago->Minneapolis->Rochester	0
Rochester	Detroit	Rochester->Chicago->Minneapolis->Detroit	0
Rochester	Chicago	Rochester->Chicago->Minneapolis->Rochester->Chicago	1

Figure 12-10 Cycle clause

Right after the cycle keyword on this example query is the name of the column that DB2 UDB should watch to make sure that the RCTE does not enter into a never-ending cycle. On the set keyword, you supply a column so that DB2 UDB can flag those cyclic paths that it found during execution of the RCTE. In this example, the cyclic_data column will be set to the string value of 0 for non-cyclic results and a value of 1 for any cyclic paths found by DB2 UDB. The output of this recursive SQL request also demonstrates that you can include the cyclic_data column in the final result set.

In the result set for this example, the last row in the result set has been marked as cyclic, and processing ended due to the fact that Chicago was visited a second time in this set of flight connections. DB2 UDB was able to detect this unending cycle since the column specified on the cycle clause, arrival, caused the database engine to make sure that no value in the arrival column is processed twice on a single set of flight connections.

Note: You may be wondering why the fifth row in the result set (Rochester → Chicago → Minneapolis → Rochester) was not marked as cyclic since Rochester appears twice in this connection list. This non-cyclic output reflects the fact that the cycle checking is only done on the recursive leg of the RCTE (blue text starting at the SELECT r.departure ... statement). Remember that the first Rochester value was supplied by the initialization leg (that is, green text).

12.7.4 OLAP expressions

In this section we discuss OLAP expressions.

ROW_NUMBER

ROW_NUMBER computes a sequential number for the rows in the final result set, as shown in Figure 12-11.

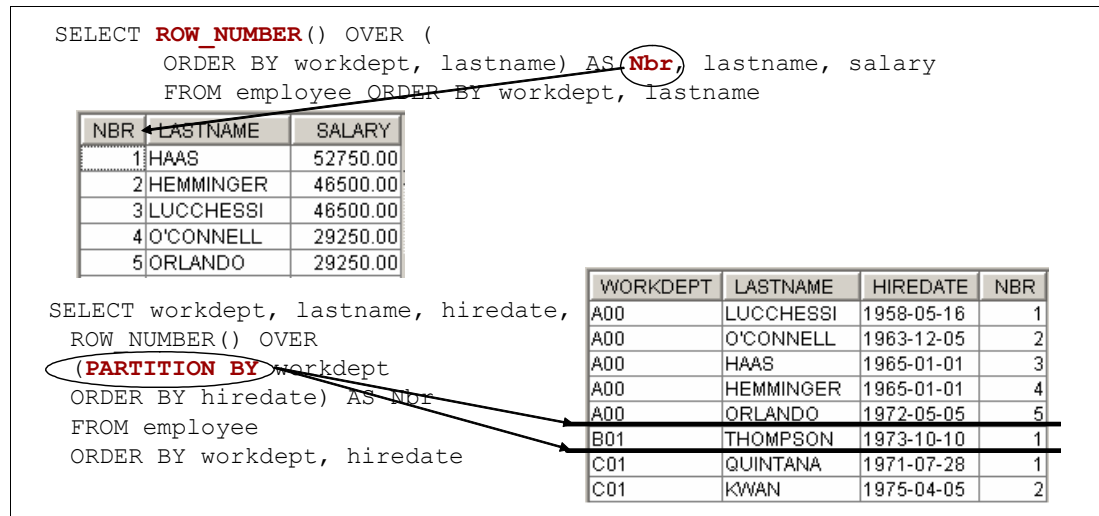


Figure 12-11 ROW_NUMBER

IS/OS programmers have often wanted the ability to arbitrarily number the rows returned by a query to have an easy way of identifying the rows in a result set.

The new ROW_NUMBER expression in V5R4 provides exactly that capability by computing that a sequential row number is to be computed for the row within the window defined by the ordering, starting with 1 for the first row. If the ORDER BY clause is not specified in the window, the row numbers are assigned to the rows in arbitrary order, as returned by the subselect (not according to any ORDER BY clause in the (select-statement)). The top example on this chart demonstrates this capability.

If your application needs the ability to restart this arbitrary row number for a subset or sub-grouping of the results, the ROW_NUMBER expression has a partitioning clause that fills this need, as shown in the second example. The row numbering is restarted each time that a new department number (that is, a partition) is encountered.

A partitioning-expression is an expression used in defining the partitioning of the result set. Each column name referenced in a partitioning-expression must unambiguously reference a result set column of the OLAP specification subselect statement. A partitioning-expression cannot include a scalar-fullselect or any function that is not deterministic or has an external action.

RANK and DENSE_RANK

RANK and DENSE_RANK are used for highlighting data attributes, independent of the result set sorting. For instance, the output captured in Figure 12-12 was created through Example 12-15.

Example 12-15 SELECT command

```
SELECT empno, lastname, salary+bonus AS TOTAL_SALARY,  
       RANK() OVER (ORDER BY salary+bonus DESC) AS Salary_Rank  
FROM employee WHERE salary+bonus > 30000 ORDER BY lastname
```

EMPNO	LASTNAME	TOTAL_SALARY	SALARY_RANK
000050	GEYER	40975.00	5
000010	HAAS	53750.00	1
200010	HEMMINGER	47500.00	2
000090	HENDERSON	30350.00	11
200220	JOHN	30440.00	9
000030	KWAN	39050.00	6
000110	LUCCHESI	47400.00	3
000220	LUTZ	30440.00	9
000070	PULASKI	36870.00	7
000060	STERN	32750.00	8
000020	THOMPSON	42050.00	4

SALARY_RANK
5
1
2
10
9
6
3
9
7
8
4

Figure 12-12 RANK and DENSE_RANK

When ordering the results of a query, it often would be nice to apply a ranking to highlight another data attribute, for example, returning a list of sales persons ordered by their last name and including their ranking in terms of total revenue. In the example on this chart, the returned employees are sorted by last name and then ranked according to the total salary with the RANK OLAP expression. This specifies that the ordinal rank of a row within the window is to be computed. Rows that are not distinct with respect to the ordering within their window are assigned the same rank. The results of ranking may be defined with or without gaps in the numbers resulting from duplicate values.

There is a closely related ranking expression known as *dense rank*. RANK and DENSE_RANK expressions behave in an almost identical fashion except when the values being ranked contain duplicate values. Whereas RANK specifies that the rank of a row is defined as one plus the number of rows that strictly precede the row (this is, if two or more rows are not distinct with respect to the ordering, then there will be one or more gaps in the sequential rank numbering), DENSE_RANK specifies that the rank of a row is defined as one plus the number of preceding rows that are distinct with respect to the ordering. Therefore, there will be no gaps in the sequential rank numbering.

If you examine the results of the RANK output, you will see that employees John and Lutz had the same salary and shared a ranking of 9, and the next employee got a ranking of 11. If you look at the DENSE_RANK output for the same query on the right, it does not have a gap in its ranking. Employees John and Lutz again share a ranking of 9, but the last employee in the salary ranking has a ranking of 10 instead of 11.

12.7.5 SELECT enhancements

In this section we discuss SELECT enhancements.

Subquery and Scalar Fullselect

The iSeries compliance with the SQL standard was finished off with the addition of an SQL standard syntax flagger and support for subquery and scalar fullselect. Fullselect is just the SQL term for multiple SELECT statements combined together with a UNION, INTERSECT, or EXCEPT set operator. While fullselects have been supported for many DB2 UDB for iSeries releases, they were not allowed on subqueries or scalar subselects.

The new scalar fullselect support is demonstrated with the UPDATE statement shown in Example 12-16.

Example 12-16 UPDATE statement example

```
UPDATE MyExchangeRates m SET conversion_rate=
    (SELECT rate FROM EuropeRates e WHERE e.ctryid = m.ctry_id
 UNION
    SELECT rate FROM AsiaRates a WHERE a.ctrycode = m.ctry_id)/100
```

In this example, your company has a table that contains currency exchange rates for all the countries that do you business with. On a daily basis, your company receives updated rates in tables from two different sources. Exchange rates for European countries come from one source, and rates for Asian countries from another. V5R4 simplifies applications by allowing the rates to be refreshed with a single SQL update by combining the sources together with the union operator. In prior releases, this same rate refresh task would have required two separate update statements.

Multi-column predicates

In addition to the fullselect support, subqueries are also easier to code as a result of the addition of multi-column predicate support. See Example 12-17.

Example 12-17 Multi-column predicate support

```
WHERE (c1,c2) IN (SELECT a,b FROM...)
WHERE (c1,c2) = (SELECT a,b FROM...)
```

Notice in the two predicate examples how multiple columns can be compared with the *in* and *equal* predicates instead of having to code separate subqueries for each column being compared.

Exclusive lock with SELECT

KEEP EXCLUSIVE LOCKS allows an application to acquire exclusive locks on rows that meet the selection criteria instead of read locks. See Example 12-18.

Example 12-18 KEEP EXCLUSIVE LOCKS

```
SELECT * FROM orders
    WHERE order_id = 'E5100'
    WITH RS USE AND KEEP EXCLUSIVE LOCKS
```

This can be useful when reading a row that you want to change, but not wanting to take the chance that the row can be locked for update after the column changes have been made. This new *keep exclusive locks* clause should be used with caution since the exclusive row locks that are acquired will prevent concurrent access by other users to those rows until the end of the transaction.

12.7.6 Richer toolbox of SQL functions

A richer and deeper set of SQL built-in functions also simplifies life for SQL programmers trying to address new business requirements.

Encrypting data stored in DB2 UDB for iSeries is another common business problem many customers are trying to solve. V5R4 includes a new built-in encryption function, ENCRYPT_TDES, to complement the ENCRYPT_RC2 function delivered in the prior release. This new encryption function provides SQL developers with the ability to secure sensitive data with the Triple DES encryption algorithm. Check also with the Data Encryption White Paper:

http://www-03.ibm.com/servers/enable/site/education/abstracts/4682_abs.html

Other new SQL functions in V5R4 include the LAST_DAY, NEXT_DAY, VARCHAR_FORMAT, and ADDMONTHS to help with the processing of date and time values.

GENERATE_UNIQUE is a new SQL function that can be used to return a unique value in the form of a bit data character string 13 bytes long. Each successive value is greater than the value returned on any previous execution of the function. This can be useful when the application requires a unique identifier for internal processing and does not need that value to be understandable to an end user.

The result of the function is a unique value that includes the internal form of the Universal Time, Coordinated (UTC) and the iSeries system serial number. The result cannot be null. The result of this function can be used to provide unique values in a table. Each successive value will be greater than the previous value, providing a sequence that can be used within a table. The sequence is based on the time when the function was executed. This function differs from using the special register CURRENT_TIMESTAMP in that a unique value is generated for each row of a multiple row insert statement or an insert statement with a fullselect. The timestamp value that is part of the result of this function can be determined using the TIMESTAMP function, with the result of GENERATE_UNIQUE as an argument.

The RAISE_ERROR function gives you the ability to signal an exception from an SQL statement based on a specific calculation or condition found in order to alert the application that an unexpected event has been detected. Example 12-19 shows that a query will terminate and signal SQLSTATE 70001 if an invalid job_type is encountered during execution of the Select statement, as shown “Exclusive lock with SELECT” on page 300.

Example 12-19 Query terminating

```
SELECT emp_name,  
       CASE job_type  
         WHEN 1 THEN 'Programmer'  
         WHEN 2 THEN 'Administrator'  
         WHEN 3 THEN 'Project Manager'  
         WHEN 4 THEN 'Manager'  
         ELSE RAISE_ERROR('70001', 'Invalid JobType') END  
  
FROM employee
```

There are also a couple of new statistical functions to provide standard deviation (STDDEV_SAMP) and variance (VARIANCE_SAMP) calculations based on the sampling of a data set instead of processing every row in a table or result set.

12.7.7 Instead of Triggers (IOTs)

A new trigger type can be used to change the semantics of INSERTs, UPDATEs, and DELETE operations against a view. This can only be defined over an SQL view. Many views are read-only due to derivations, joins, grouping, and so on.

IOTs are useful in setting up encryption to be semi-transparent. V5R4 support builds on base IOT support delivered via V5R3 PTF by providing support for join views, and so on. Refer to the following Web site:

<http://www.ibm.com/series/db2/iot.html>

Example 12-20 Example

```
CREATE VIEW empdept AS
    SELECT empno, firstname, lastname, deptname FROM employee, department
    WHERE workdept=deptno
```

```
CREATE TRIGGER UpdateJoin
    INSTEAD OF UPDATE ON empdept
    REFERENCING OLD ROW AS o NEW ROW AS n
    FOR EACH ROW
    BEGIN
        UPDATE employee
            SET empno=n.empno, firstname=n.firstname, lastname=n.lastname
            WHERE empno=o.empno;
        UPDATE department SET deptname = n.deptname WHERE deptname=o.deptname;
    END
```

An Instead of Triggers is a new type of trigger for SQL Views that was introduced on V5R3 via PTF. This V5R3 support delivered mid-release had a number of restrictions that have now been eliminated in V5R4. Unlike the existing trigger support, Instead of Triggers can only be created over an SQL view. They can never reference a table or physical file.

This example of an Instead of Trigger created over an SQL view containing a join was not allowed with the restrictive V5R3 support. This example demonstrates how Instead of Triggers can be used to provide updateable join support on the empdept view.

Whenever the empdept SQL view is updated, the Instead of Trigger (UpdateJoin) allows the developer to have the update succeed instead of DB2 returning an error that the view is read only. The UpdateJoin trigger transforms the original update request into two separate update statements against the two underlying tables, employee and department, that are joined together in the empdept SQL view.

12.7.8 Miscellaneous enhancements

In this section we discuss miscellaneous enhancements.

Blanks are allowed in hex literals

A constant (also called a literal) specifies a value. The HEX function returns the hexadecimal encoding of an integer expression.

Support was added for LABEL ON INDEX

You can now perform *label on index* functions.

Optional journaling for SQL tables

NOT LOGGED INITIALLY is a new option available on the Create & Alter Table statements. This new clause provides an SQL interface to temporarily end journaling. For example, if a work table is going to be populated with data to get started and there is no requirement to commit or roll back the new data, the NOT LOGGED clause can be used to speed performance during this initial data population by temporarily ending journaling. At the end of the transaction, journaling is automatically restarted on the table.

Example 12-21 CREATE TABLE

```
CREATE TABLE new_materials
  AS (SELECT * FROM materials WHERE YEAR(item_date)=2006)
  WITH DATA NOT LOGGED INITIALLY
```

Simpler SQL and CL integration

You will find that you have simpler SQL and CL integration with the new qcmdexc stored procedure:

```
CALL qcmdexc('ADDLIBLE DBLIB2', 15)
```

instead of:

```
CALL qcmdexc ('ADDLIBLE DBLIB2', 0000000015.00000)
```

Syntax flagger for SQL is a core standard that allows for those SQL applications that may need to interact with the i5/OS operating system using system CL commands. V5R4 simplifies that interaction with an improved QCMDExc stored procedure that is easier to code for CL command integration. Syntax flagger is only available with green screen interfaces.

12.8 XML Toolkit

To ensure the availability of the latest versions of the XML parsers, the XML Toolkit for iSeries (5733-XT1) was developed.

The XML Toolkit replaces, both in functionality and support, the integrated XML parsers introduced in i5/OS in V5R1. This offering consists of two new parsers: XML4C 5.5 for C++ and XML4PR 5.5 for procedural languages (ILE RPG, ILE C, and ILE COBOL), as well as Versions 4.0, 5.0, 5.2, and 5.3.1 and their supporting API documentation, sample programs, and include files.

This offering also includes the new Version 1.9 of the XSL transformer, as well as Version 1.6 and their supporting API documentation, sample programs, and include files. The Java version of the XML parser and XSL transformer are available in the IBM Toolbox for Java (5722 JV1).

Older versions of the Java version of the XML parser can be found in the Integrated File System directory /QIBM/ProdData/OS400/xml/lib.

New components added to the XML Toolkit include XML Scripting for iSeries and Web Services Client for C++.

For more information see the iSeries (i5/OS) XML Web site at:

<http://www-03.ibm.com/servers/eserver/iseries/software/xml/>

Also see the iSeries Information Center for V5R3 at:

<http://www.ibm.com/eserver/series/infocenter>

Select **iSeries Information Center, Version 5 Release 3** → **Programming** → **XML Toolkit for iSeries**.

Or:

<http://publib.boulder.ibm.com/infocenter/series/v5r3/index.jsp?topic=/rzamj/rzamjmain.htm>

For V5R4 go to:

<http://publib.boulder.ibm.com/infocenter/series/v5r4>

Select **iSeries Information Center, Version 5 Release 4** → **Programming** → **XML Toolkit for iSeries**.

12.9 WebSphere Development Studio and Client

Over the years, a number of developmental tools have been implemented on System i. Some examples include:

- ▶ PDM - Program Development Manager
- ▶ SDA - Screen Design Aid
- ▶ SEU - Source Entry Utility
- ▶ RLU - Report Layout Utility
- ▶ DFU - Data File Utility
- ▶ AFP - Advanced Function Presentation™
- ▶ CGU - Character Generator Utility

As of V4R5, there is only one application development product sold by IBM for System i. This is WebSphere Development Studio, often called Development Studio. This tool includes all four host compilers. That is, the Application Development ToolSet (ADTS) incorporates all traditional tools (ADTS = PDM+SEU+SDA+RLU+DFU+AFP+CGU). Additionally, there are unlimited licenses of the workstation-based toolset named WebSphere Development Studio Client (formerly WebSphere Development Tools).

If you are an existing customer who has a subscription, you can upgrade to Development Studio free of charge. Without a software subscription, there is an upgrade fee. New licenses of Development Studio are priced very competitive compared to the combined prices of all constituent products. As of V5R1, there is no way to purchase the compilers or tools individually. So if you have RPG at V5R1 or later, you must have Development Studio, and hence are entitled to Development Studio Client.

ILE RPG enhancements include the ability to access data in XML documents. Also with i5/OS V5R4, RPG developers will now be able to use free format SQL database access in their applications. With improved flexibility to recognize and analyze SQL statements, this enhancement should make it much easier for RPG developers to incorporate SQL into their current applications. ILE COBOL enhancements include the ability to output COBOL data records into XML format. Application Development ToolSet supports network server configuration objects in PDM and changing PDM defaults.

A fourth generation development tool, Enterprise Generation Language (EGL), has been added to the workstation tool. This tool allows developers to rapidly generate Java applications, as well as iSeries COBOL green screen applications in the Advanced Edition,

using the same fourth generation language. Enterprise Generation Language is the follow-on product to VisualAge® Generator.

The RPG enhancements are described in 12.4, “ILE RPG programming” on page 278, while COBOL enhancements are described in 12.5, “ILE COBOL programming” on page 283.

Development Studio has been a huge success, with over 80,000 licenses sold. Just as every development machine used to have PDM and SEU, every development machine will now have all the modern Application Development tools from IBM. This ubiquity is especially important for business partners who build and sell software. These Business Partners are now free to build software using any of the technologies or tools in Development Studio, and can assume that their customers will have the tools required to tailor everything from RPG to Java and Web user interfaces. This effectively raises the lowest common denominator to a level unparalleled by any other operating system.

For consultants who do not have an iSeries of their own, but still wish to have the client tools, Development Studio Client is also made available as a passport advantage product so it can be purchased off the shelf from IBM Direct.

The WebSphere Development Studio Client is based on WebSphere Studio Site Developer, while Development Studio Client Advanced is based on WebSphere Studio Application Developer.

WebSphere Development Studio Client for iSeries has a new extension to the Web service wizard to create Web Services from ILE programs in one step, making it easier to create SOA applications.

The WebSphere Development Studio Client Studio Client Advanced is based on WebSphere Studio Application Developer. It includes everything in the Developer, and:

- ▶ EGL COBOL Generation
- ▶ Enterprise Java Beans (EJB™)
- ▶ Java 2 Enterprise Edition (J2EE)
- ▶ Test cases
- ▶ Portal Tool Kit

WDSC V6.0 packaging

The Standard Edition comes with:

- ▶ RWD (Rational® Web Developer) V6.0
- ▶ WDSC V6.0 plug-ins
- ▶ Media package contains
 - DVD: RWD and WDSC
 - CD1: CD7 RWD and WDSC
 - Seven other CDs
 - CODE and VisualAge RPG V6.0
 - IBM Agent Controller
 - WebSphere Test Environment Version 5.x (CD1, CD2)
 - WebSphere Application Server for Developers V6.0
 - WebSphere Host Access Transformation Services Toolkit (HATS)
 - DB2 (for the workstation)

The Advanced Edition contains all of the above, plus 18 other CDs:

- ▶ CODE and VisualAge RPG V6.0
- ▶ IBM Agent Controller
- ▶ WebSphere Test Environment Version 5.x (CD1, CD2)

- ▶ WebSphere Application Server for Developers V6.0
- ▶ Embedded messaging client and server
- ▶ WebSphere Portal V5.0 Test Environment
- ▶ WebSphere Host Access Transformation Services Toolkit (HATS)
- ▶ DB2 (for the workstation)
- ▶ Crystal Enterprise Professional Edition V10 (CD1, CD2)
- ▶ Crystal Enterprise Embedded Edition V10
- ▶ WebSphere Portal for Multiplatforms (6 CDs)

12.9.1 WDS V6.0.1

Key enhancements for V6.0.1 begin with its base. Using the Rational Web Developer for WebSphere Software (RWD) refresh pack V6.0.1 (part of the IBM Rational Software Development Platform), it also includes other Eclipse-based development tools that allow you to adapt and extend your development environments with IBM and partner tools.

IBM WebFacing Tool delivers extensions to allow Web-enabled applications created with the WebFacing Tool to interoperate with 5250 applications in a browser. (Deployment for applications using these extensions requires IBM WebFacing Deployment Tool for WebSphere Development Studio with HATS technology, V6.0.1, 5724-N52.)

HATS Toolkit V6.0.4 delivers dynamic runtime transformation of the 5250 datastream with no requirement for Online Transaction Processing (OLTP) capacity starting with i5/OS V5R4.

A new WDS Lite technology preview provides a light-weight edit/compile/debug environment for developing native i5/OS applications (256 MB of RAM is recommended to operate effectively).

It provides V5R4 support for RPG, COBOL, and CL.

It has a new extension to the Web Service wizard to create Web Services from ILE RPG/COBOL programs in one step, making it easier to create SOA applications.

The WebFacing Tool and Web Tools generate Common Base Event (CBE) logging data that can be used by the Log and Trace Analyzer, provided in WebSphere Development Studio Client Advanced Edition for iSeries.

Eclipse is an Integrated Development Environment (IDE), an Open Source extensible platform for developing Java applications and integrating them. IBM Tools (Rational and WebSphere Studio) are built on the Eclipse framework. Eclipse was originally developed by IBM and donated to the Open Source community. Eclipse IDE is extended by building plug-ins, which are applications that run in Eclipse IDE.

WDS V6.0.1 packaging

In this section we discuss WDS V6.0.1 packaging.

Server tools

The server tools are:

- ▶ ILE RPG
- ▶ LE COBOL
- ▶ LE C/C++
- ▶ ADTS

Workstation tools

These are received via the entitlement of unlimited workstation licenses of WDSC V6.0 per System i5 with WDS installed:

- ▶ WDSC refresh pack V6.0.1
- ▶ RWD refresh pack V6.0.1
- ▶ Remote System Explorer - a powerful integrated development environment for i5/OS application development.
- ▶ i5/OS plug-ins to help Java, Web, and Web services developers access i5/OS data and applications:
 - i5/OS re-facing tools
 - IBM WebFacing Tool
- ▶ Host Access Transformation Services Toolkit V6.0.4
- ▶ IBM integrated i5/OS debugger (Eclipse-based) for all languages
- ▶ Separate installation of CODE and VARPG

12.10 WebFacing and HATS

The IBM WebFacing Tool of WebSphere Development Studio now delivers extensions to allow Web-enabled applications created with the WebFacing Tool to interoperate with the 5250 application in a browser using IBM Host Access Transformation Server (HATS). The new IBM WebFacing Deployment Tool for WebSphere Development Studio with HATS Technology (WDT) provides an integrated runtime for applications that have been converted with the IBM WebFacing Tool and 5250 applications that are converted on-the-fly by HATS. The HATS Toolkit V6.0.4 that is included in WDS now delivers dynamic runtime transformation of a 5250 datastream application with no requirement for 5250 OLTP capacity starting with i5/OS V5R4.

Figure 12-13 shows an example of a 5250 screen and a corresponding browser screen processed under WebFacing support.

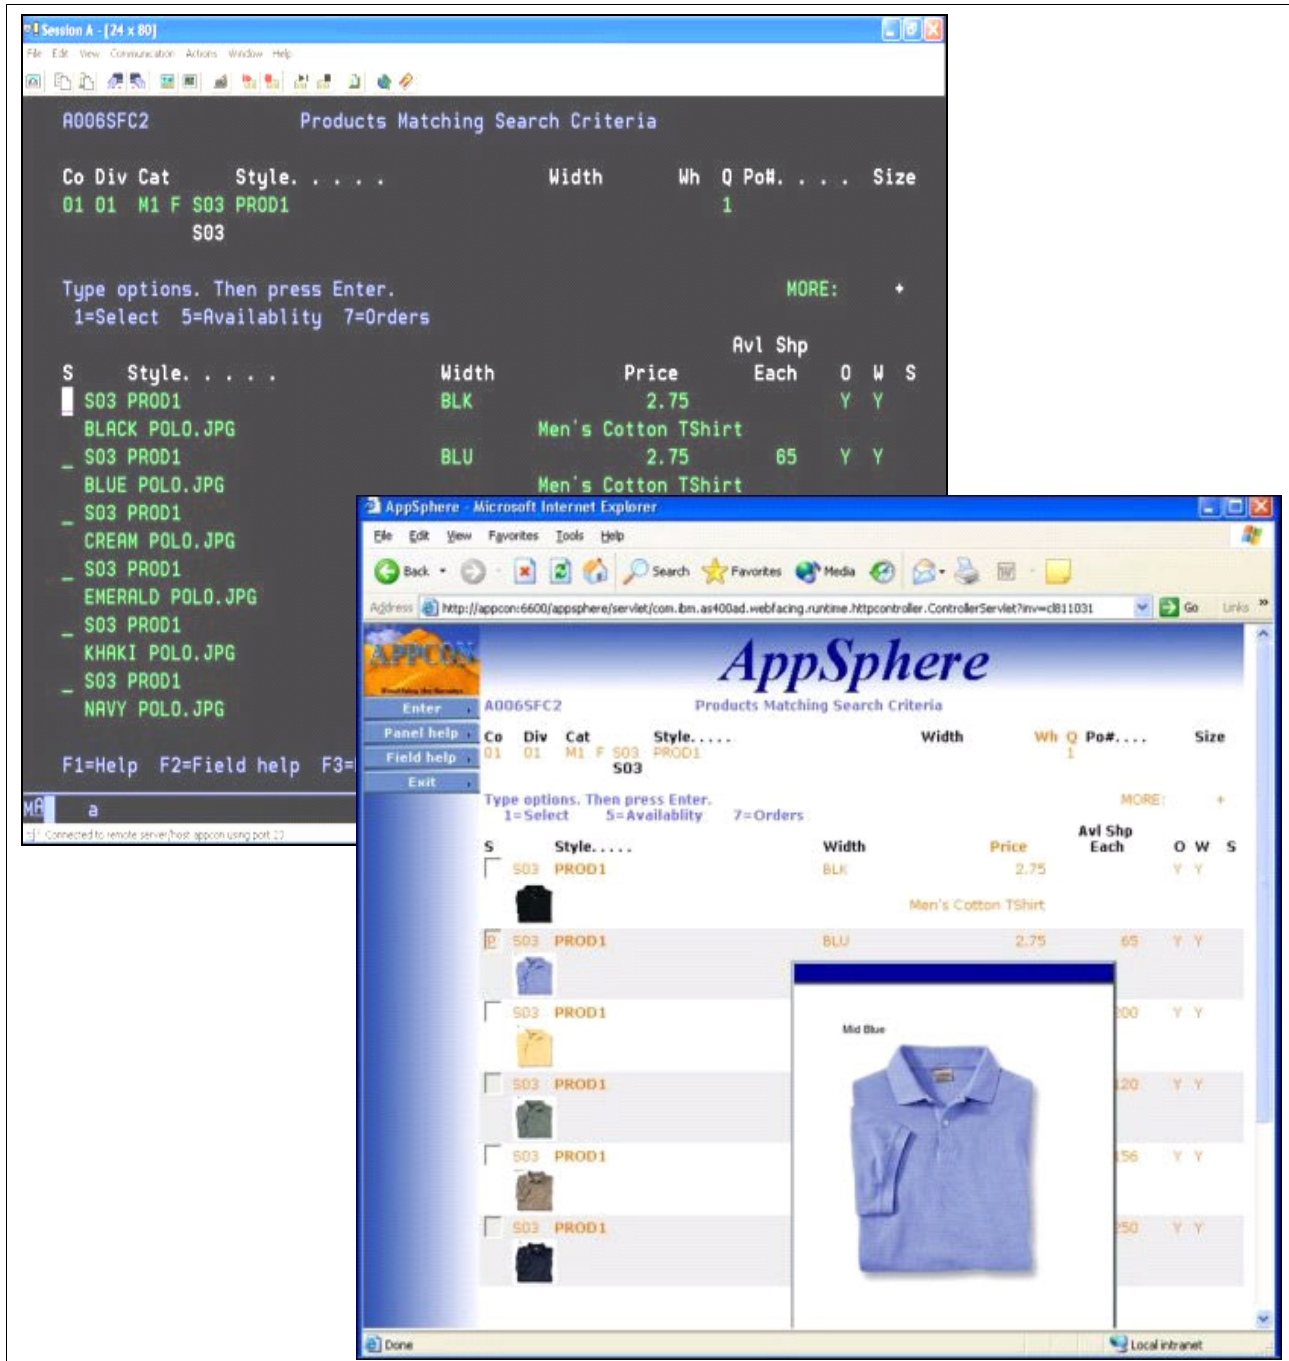


Figure 12-13 5250 and Webfaced example screens

The IBM WebFacing Tool generates JavaServer Page-based Web interface to 5250 applications based on Data Description Specifications (DDS) source. HATS, on the other hand, does not require any application source code, as it generates a Web interface on-the-fly or dynamically based the 5250 datastream.

Figure 12-14 shows an example of a 5250 flight reservation screen and the HATS version of the same screen on a browser window.

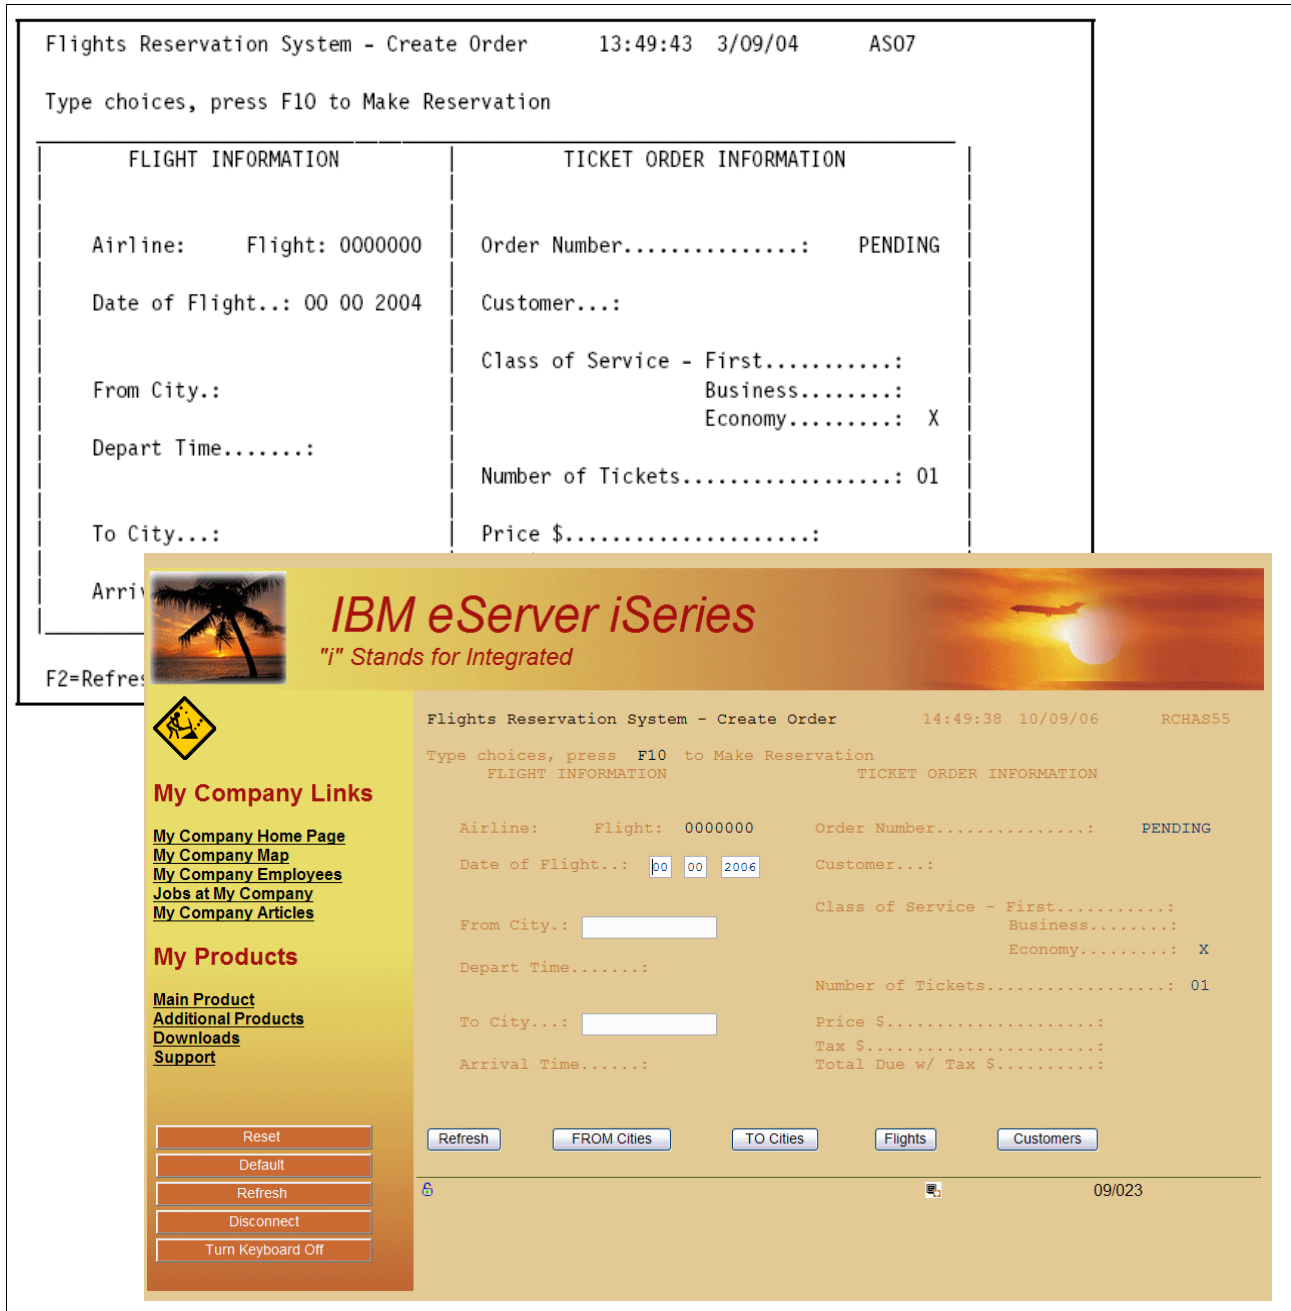


Figure 12-14 5250 and HATS example screens

Note: The colors used for text in this HATS browser screen example are legible in real life. However, the text shown in this figure becomes unclear when copied into this book because of the color choices used in this example. You can see enough of the text to get an idea of the HATS screen relative to the 5250 screen where we used a text version of the 5250 screen.

Using either WebFacing or HATS, you can leave the 5250 screen relatively unchanged on the browser interface. However, typically you would use the features of either WebFacing or HATS to improve the user interface or add usability functions to the interface.

The WebFacing Deployment Tool with HATS Technology (WDHT) delivers a runtime enablement that exploits the capabilities of both HATS and WebFacing. It allows refacing projects to combine the strengths of both products using WDS. It provides a system-wide solution for multiple applications, with or without source code.

No interactive (OLTP) workload is required for HATS or WebFacing with i5/OS V5R4 and WDS editions at V5.1 or V6.0.1.

In the past, WDS Advanced Edition was required if a WebFaced application called a system screen or included User Interface Manager (UIM) panels. But if the application called other DDS applications that were not WebFaced, then the application would fail. Now, with the integration of WebFacing and HATS, the WDHT will use HATS technology to automatically convert the 5250 interface, preserving the same look and feel as the user saw in the WebFaced application with WDS or WDS AE.

The effect of these changes is to help customers and ISVs that liked the WebFacing Tool but could not convert all of their applications in this manner.

IBM WebFacing Deployment Tool for WebSphere Development Studio V6.0.1 with Host Access Transformation Services (HATS) Technology delivers an integrated modernization solution for the IBM System i5 marketplace. It is a runtime enabler allowing the WebFacing Tool and HATS technologies to work together in an integrated fashion. Both WebFacing Tool and HATS Toolkit customers can work seamlessly with their existing projects, while taking advantage of the new capabilities when combined with i5/OS V5R4.

Deployment enablement for new IBM WebFacing Tool and HATS functions includes:

- ▶ *Out of the box* integration of WebFaced and non-WebFaced applications in a browser session
- ▶ New ability for HATS to Web-enable 5250 applications without requiring Online Transaction Processing (OLTP) capacity in i5/OS
- ▶ Additional, attractive deployment-based pricing options

IBM WebFacing Deployment Tool for WebSphere Development Studio V6.0.1 with HATS Technology is now available for electronic delivery. Media and documentation are also now available.

The IBM WebFacing Tool and the HATS Toolkit are included with WebSphere Development Studio Client for iSeries and WebSphere Development Studio Client Advanced Edition for iSeries. WebFacing Tool customers who do not take advantage of the new capabilities can continue to create and deploy applications from the WebFacing Tool.

However, customers who take advantage of the new WebFacing Tool extensions will find that deployment for the new functions is limited to two user connections for testing purposes only. To enable the new functions for deployment to production, customers must purchase WDHT. Web-enabled applications created using the HATS Toolkit are limited to two user connections for testing purposes only. WDHT is required for deployment of production mode HATS applications.

Though there will be pricing changes under this new version of WDS (go to the WDS Web site for the latest information), the introduction of WDHT makes HATS more affordable by eliminating the requirement for OLTP capacity for HATS applications working through the WebFacing server. Without the OLTP requirement for deploying a HATS application to i5/OS

starting with V5R4, the overall cost of HATS solutions is significantly reduced. HATS applications can be deployed to i5/OS Standard Edition configurations.

Attention:

You can find more information about HATS and WebFacing in the following sources:

- ▶ HATS information center
<http://publib.boulder.ibm.com/infocenter/hatshelp/v60/index.jsp>
- ▶ WebSphere Development Studio Client for iSeries information center
<http://publib.boulder.ibm.com/infocenter/iadthelp/v6r0/index.jsp>
Search with *webfacing*.
- ▶ WebSphere Development Studio Client for iSeries support
<http://www-306.ibm.com/software/awdtools/wdt400/support/>
- ▶ IBM Redbooks
 - WebFacing: *Mastering the IBM WebFacing Tool*, SG24-6331
 - WebFacing and HATS: *IBM System i Application Modernization: Building a New Interface to Legacy Applications*, SG24-6671

During October 2006, i5/OS announced a new IBM Express Runtime Web Environments for i5/OS product. This product installs HTTP Server for i5/OS, WebSphere Application Server for i5/OS V6.0, and Web applications that run under the application server created by this product. Installing this configuration and applications can help speed up your learning of the Web application serving environment without having to create your own applications. Running the provided applications can give you early insights on starting, stopping, and managing these workloads without spending any resources on application development and deployment to the Web application server. The applications provided include:

- ▶ iSeries Access for Web
- ▶ An example RPG 5250 application called flight400
- ▶ A Webfaced version of flight400
- ▶ A HATS-ed version of flight400

For more information in this book refer to Chapter 14, “i5/OS-based Web enablement enhancements” on page 337.

WDS Lite Technology

WDS Lite is a large file set. In response to concerns about this, a lighter version is being developed. Capabilities to limit workbench functionality to E/C/D for iSeries were added. This would be in both Standard and Advanced.

You must install the whole product first. However, you will have the option to select the level of functionality, using the Rational Product Updater. This would add a 500 MB of disk space requirement. Memory (RAM) required at runtime will be 256 MB for E/C/D.

This is a fully supported technology preview. In a future release we will have more componentization, such as install on demand.



i5/OS-based WebSphere application support

In this chapter we cover WebSphere Application Server Versions 6.0 and 6.1 for i5/OS products.

13.1 WebSphere Application Server V6.0

WebSphere Application Server V6.0 for i5/OS became available several months after i5/OS V5R3 was introduced. This major section summarizes the primary new or enhanced areas within WebSphere Application Server V6.0 for i5/OS.

Figure 13-1 represents a history of primary functions, Java Development Kit (JDK) levels and processes available from WebSphere Application Server from Version 4 through Version 6.1, with an indication of potential areas of enhancements in future versions.

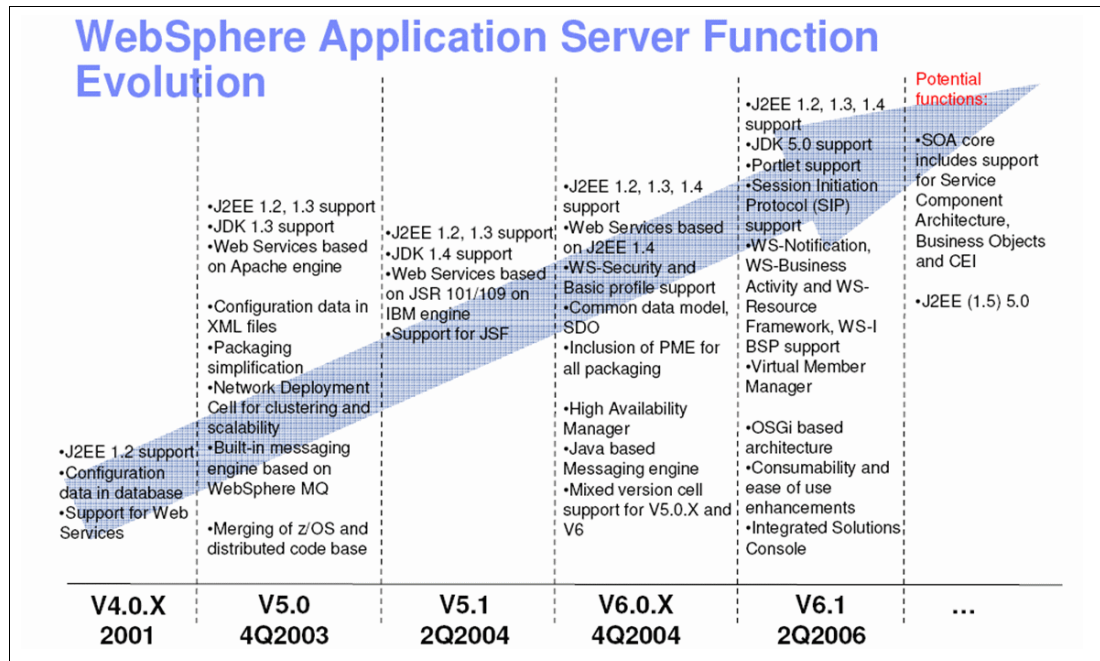


Figure 13-1 WebSphere Application Server: evolution of major functions and processes

In this section we focus on WebSphere Application Server V6.0, after first reviewing some basic Java-based application specifications at the J2EE 1.4 level.

For information in this book on Java J2EE 1.5 level capabilities refer to 12.6, “Java” on page 283. For more information about WebSphere Application Server V6.1, refer to 13.2, “WebSphere Application Server Version 6.1 on i5/OS” on page 327.

13.1.1 WebSphere Application Server programming model

We now support multiple Java 2 Platform, Enterprise Edition (J2EE) versions in Version 6 Application Server. You can use J2EE 1.2, 1.3, and 1.4. This simplifies migration by allowing your existing J2EE 1.2 or 1.3 applications to run in Version 6 Application Server. This also allows your WebSphere 4 applications to run in WebSphere 6 without modifications.

Similar to Version 5 Application Server, you will have Version 4 Data Source for your J2EE 1.2 applications needing access to the backend relational database.

J2EE 1.4 overview

The Java 2 Platform, Enterprise Edition (J2EE) 1.4 specification is a standardized architecture for Web services enabled applications. This specification is maintained by Sun, with input from a large audience of users and developers.

In this section we highlight a number of what we feel are important changes to the J2EE specifications only. If you would like to read, download, or buy a copy, the entire J2EE1.4 specification is available at:

<http://java.sun.com/j2ee/1.4/index.jsp>

Enterprise Java Beans (EJB) 2.1

With support of J2EE 1.4 comes the addition of typed message beans. These are used for any inbound Java Connection Architecture (JCA) including pluggable Java Message Service (JMS) providers.

A container-managed timer service is now available. The timer service provides coarse-grained, transactional, time-based event notifications to enable enterprise beans to model and manage higher-level business processes.

There is now support allowing all enterprise beans to utilize external Web services. Also, you now have the ability to have a stateless session bean to implement a Web service endpoint.

The architecture has been generalized to allow message-driven beans to support the use of messaging types in addition to JMS.

The Enterprise JavaBeans™ Query Language (EJB QL) has been enhanced to include support for aggregate functions, ordering of results, and additional scalar functions, and have clarified the rules for allowing null values to be returned from finder and select methods.

Deployment descriptor elements have been added to enable the specification of the flow of messages.

We have added deployment descriptor elements to enable the specification of the flow of messages from producers to consumers by means of common message destinations.

Java Message Service (JMS) 1.1

There has been a unification of the point-to-point (PTP) and publish/subscribe (pub/sub) interfaces. PTP service is built on the concept of message queues, senders, and receivers. In a pub/sub, an author publishes a topic, to which consumers can subscribe. The topic functions somewhat like a bulletin board. See a good overview at:

<http://java.sun.com/j2ee/1.4/docs/tutorial/doc/JMS3.html>

J2EE Connection Architecture (J2CA) 1.5

Some key points are enhancements to contracts, which are specified communication formats.

The Lifecycle Management Contract allows the application server to manage the life cycle of the resource adapter. This is the startup and shutdown functionality.

Work Management Contract is what lets the resource adapter do work by submitting it to an application server for execution. This removes thread management requirements from the resource adapter, as the application server manages this aspect. This is recommended, as application servers are more efficient.

Transaction Inflow Contract allows a resource adapter to pass a imported transaction to the application server. It also allows additional controls to completion and crash recovery coming from a Enterprise Information System (EIS).

Message Inflow Contract allows the resource adapter to deliver messages regardless of message format.

An in-depth overview of the 1.5 specification is available at:

http://java.sun.com/developer/technicalArticles/J2EE/connectorarch1_5/

Miscellaneous

Some miscellaneous notes are:

- ▶ Servlet 2.4 includes extensible deployment descriptors and request/response listeners.
- ▶ JavaServer Pages (JSP) 2.0 added expression language and simple tag extension.
- ▶ Java Database Connectivity (JDBC) 3.0 added metadata and cursor support.
- ▶ JavaMail™ 1.3 updates, which allow for platform-independent mail and messaging.

Changes in Web Services with V6.0

Web Services and XML support have been enhanced:

- ▶ Standards/portability - XML Schema definitions for all deployment descriptors
- ▶ JAX-P 1.2 - new properties for XML parsers
- ▶ JAX-R - XML registry API
- ▶ JAX-RPC - APIs for representing WSDL-based services as RPCs in Java (and vice-versa)
- ▶ JSR 109 - Web services programming and deployment model
- ▶ SAAJ 1.1 - SOAP Attachments API for Java

In Table 13-1 we show the changes in Web services by WebSphere release.

Table 13-1 Web Services changes by release

WebSphere Application Server 4.0 & 5.0	Apache SOAP The programming model, deployment model, and engine.
	Proprietary APIs Because Java standards for Web services did not exist. Not WS-I compliant.
WebSphere Application Server 5.02/5.1	JAX-RPC (JSR-101) 1.0 New standard API for programming Web services in Java
	JSR-109 1.0 New J2EE deployment model for Java Web services.
	SAAJ 1.1
	WS-Security Extensions added.
	WS-I Basic Profile 1.0 Profile compliance.
	UDDI4J Version 2.0 (client)
	Apache Soap 2.3 enhancements The engine is a new high-performance SOAP engine supporting both HTTP and JMS.

WebSphere Application Server 6.0	JAX-RPC (JSR-101) 1.1 <ul style="list-style-type: none"> ▶ Additional type support. ▶ xsd:list. ▶ Fault support. ▶ Name collision rules. ▶ New APIs for creating services. ▶ isUserInRole().
	JSR-109 - WSEE <ul style="list-style-type: none"> ▶ Moved to J2EE 1.4 schema types. ▶ Migration of Web services client DD moving to appropriate container DDs. ▶ Handlers support for EJBs. ▶ Service endpoint interface (SEI) is a peer to LI/RI.
	SAAJ 1.2 <ul style="list-style-type: none"> ▶ APIs for manipulating SOAP XML messages. ▶ SAAJ infrastructure now extends DOM (easy to cast to DOM and use).
	WS-Security <ul style="list-style-type: none"> ▶ WSS 1.0. ▶ Username Token Profile 1.0. ▶ X.509 Token Profile 1.0.
	WS-I Basic Profile 1.1 Attachments support.
	JAXR support
	UDDI v3 support <ul style="list-style-type: none"> ▶ Includes both the registry implementation and the client API library. ▶ Client UDDI v3 API different from JAXR (exposes more native UDDI v3 functionality).

Programming model extensions

Programming model extensions (PMEs) are IBM-developed extensions to the J2EE model. We develop these ahead of the specification to meet customer-reported needs. All of our extensions are submitted to the J2EE committee. Once accepted, PMEs are removed from this status.

Formerly, these were available only in the Enterprise Edition. Now the core extensions are included in all versions, even Express. They are:

- ▶ Last participant support
- ▶ Internationalization service
- ▶ WorkArea service
- ▶ ActivitySession service
- ▶ Extended JTA support
- ▶ Startup beans
- ▶ Asynchronous beans (now called WorkManager)
- ▶ Scheduler service
- ▶ Object pools
- ▶ Dynamic query
- ▶ Web Services Gateway Filter Programming Model (with migration support)
- ▶ DistributedMap
- ▶ Application profiling

See a good overview at:

<http://websphere.sys-con.com/read/48334.htm>

13.1.2 WebSphere Application Server V6.0 system management features

Changes extend the Version 5 System Management Model with the configuration files in XML. This reduces the learning curve for managing WebSphere Application Server Version 6 environments.

Support for the J2EE 1.4 specification includes the Java Management Extensions (JMX) 1.2.

Two additional Java Specification Request (JSR) items are supported. JSR 077 covers J2EE Management, which provides Application Program Interfaces (APIs). This allows tool vendors to create tools to control WebSphere. These APIs include notification emitters and standard patterns. There is an information model representing J2EE application server concepts.

JSR-088, which covers J2EE Deployment, is analogous to JSR-077. In this case, it allows you to deploy your own tools and applications in WebSphere. It utilizes XML-based deployment interfaces for this.

Fine-grained Application Update adds the ability to introduce small delta changes to installed applications. That is, you have the ability to add, update, or remove parts of the installed application and restart the changed part. This is done through a wizard in the administration console.

Support for extensible server types for Web server and generic server have been added.

We also have the introduction of node groups. This allows z/OS and distributed nodes within same cell.

WebSphere Application Server V6.0 renamed an *instance* running a Web application as a *profile*. The architectural changes are significant enough to create a new term. New WebSphere profile enhancements have the WebSphere Application Server instance architecture now being called server profiles. There are many new options to manage instances. Default Application Server and Deployment Manager are also an instance.

The improved administrative console (Figure 13-2) look and feel provide consistent cross-IBM product look and feel. The screen changes were designed to better organize the links and options in order to group them relative to the work you are doing. Changes to console views were based on context, version, platform, and installed capabilities.

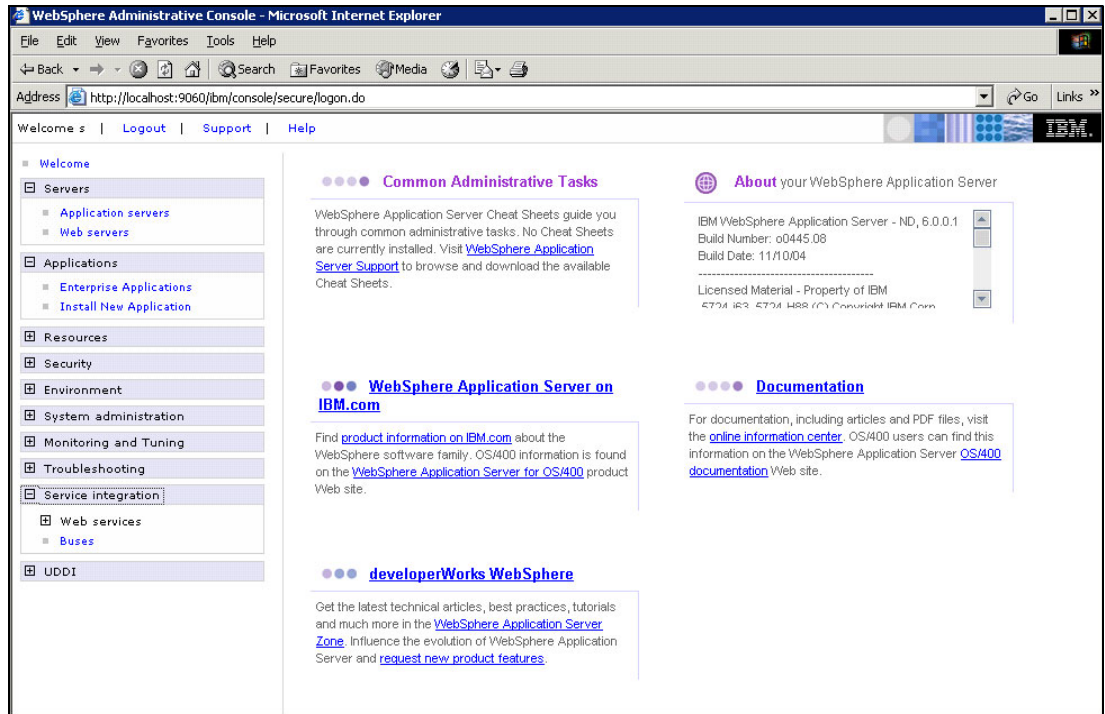


Figure 13-2 Improved administration console

It also now includes the integration of Tivoli Performance Viewer and the IBM HTTP Server Version 6 management. In the case of the Tivoli Performance Viewer (Figure 13-3), there is no need to install the client on a workstation, as this is displayed on any browser capable attached device. There is only a small piece of code that must be installed the first time a display accesses the viewer. This install is automated.

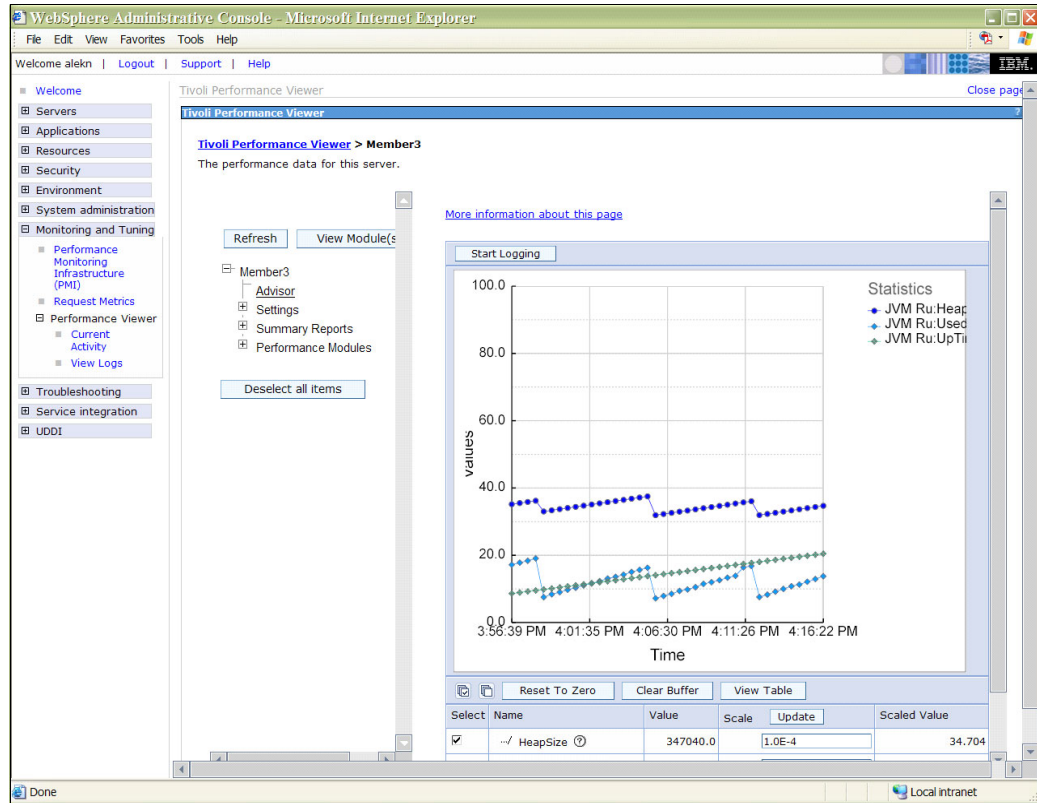


Figure 13-3 Integrated Tivoli Performance Viewer

System applications include the administrative console and file synchronization.

It should be noted that there is now a new URI:

`http://<hostName>:<port>/ibm/console`

Service integration technologies in Version 6

The service integration functionality within WebSphere Application Server provides a highly flexible messaging fabric. This supports a service-oriented architecture with a wide spectrum of quality of service options, supported protocols, and messaging patterns. It supports both message-oriented and service-oriented applications.

With platform messaging now fully integrated within WebSphere System Management, the Admin Console provides MQ-Explorer type management. All Java implementation is within the server process. There are no external processes required. So, while we can co-exist with WebSphere MQ, this is no longer a requirement.

Performance monitoring, trace, and problem determination are all included, as is integration with WebSphere Security.

Workload management high availability

High Availability is a function available under the Network Deployment version of WAS. It was completely redesigned and rewritten in Version 6.

Unified clustering

Workload Management (WLM) provides consistency for clustering of different resources. Operational ease of use has been improved. The view and use of clusters is administered in a unified and consistent manner for all protocols (HTTP, EJB, JMS, and so on). New WLM functions are implemented once for all protocols.

The replication of services, HTTP, and session beans ensures that there is no single point of failure in your system. For instance, if a Java Virtual Machine (JVM) fails, the WLM opens another instance of it. This makes WLM a highly available service, which makes cluster and routing information always available.

WebSphere Version 6 HA terminology

The Core Group is a component of the HA manager function that monitors the application server environment and provides peer-to-peer failover of application server components. A Network Deployment cell has at least one core group. A DefaultCoreGroup is automatically created.

A core group consists of the following:

- ▶ Core group members
 - Node agents, deployment manager, and application servers that join a view.
- ▶ Coordinator (or governing service)
 - An active core group process that manages the availability of singletons within the core group. There is a coordinator per core group.
 - Tracks which HA groups are joined or left by which core group members.
 - Receives notification when core group members fail or are unavailable.
 - If the coordinator fails, another active core group process is chosen as the coordinator, which enables it to be highly available.
 - Coordinator function can be partitioned over multiple processes.
 - Can configure a specific process or a set of processes as the preferred location for the coordinator to run. If the preferred processes are not available, an active process will be chosen. The Deployment Manager and Node Agent processes are good candidates to be the preferred processes.

By default, the core group includes all services, so modification is simplified. Some setup is required. Once set up, the failover happens without interaction, and should be invisible to the user. Again, this is not a product, it is a group of related functions.

Data replication service enhancements

Version 6.0 of the WebSphere Application Server incorporates an improved version of the Data Replication Service that has now been rebased on top of a multicast-based transport, the high availability services, and the channel framework.

As a result of this change, the Data Replication Service will be simpler to configure and use, and will perform and scale better. The product will also be simpler to configure with scripting, since a lot of the complicated tuning options are now unnecessary.

Changes include an optimized communication stack, and the allowance for use of both unicast and multicast IP. This allows improvements in the range of 4x to 8x.

Integration with WLM will allow WLM to intelligently pick a backup application server to route subsequent http requests to, such that future requests will be sent to an application server where a backup session exists.

This leverages the failure detection provided by high-availability services. Along with the WLM/Unified Clustering integration, this allows for active failure recovery. For instance, if your customer has a shopping cart, and their current session fails, you do not want them to have to start from the beginning. With HttpSession replication, if the affinity server for a HttpSession goes down, WLM can route to another server that has a backup copy ready to use.

WLM improves usability through leverage of group services to simplify partitioning. You now have n-replica, where you simply define the number of backup copies you want for data.

Added users of this component include Stateful Session Beans failover and the Entity Bean Persistence manager. Enhanced features and improved performance will be available to all the users of DistributedMap as well.

What used to happen in V5.0, the Application Server shipped with a JMS-based transport. This did not have as high of performance capabilities. In addition, configuring this component was not completely intuitive and there was a fair amount of complexity involved in setting it up.

What are the reasons you care about this? First, it is much simpler and more intuitive for you to configure Data Replication Service. More sample scripts make it easier to use wsadmin to set up the component. You have improved performance and scalability and active failure recovery.

13.1.3 WebSphere Application Server V6.0 security

We now have Java Authorization Contract with Containers (JACC) 1.0 support. These are APIs for registering J2EE component authorization policies, which allows plug-in of your authorization servers. So, if you have an existing solution regarding access to applications and resources, you can continue to use it as a single point of control.

JACC compliant Tivoli Authorization Module (TAM) is shipped with Version 6.0. Version 6.0 continues to support the non-JACC native authorization (similar to Version 5).

Security attributes can be propagated from WebSphere Application Server Version 5.1.1.

We have also implemented the WS-Security 1.0 WebServices Specification. A specification such as this is actually a number of specifications grouped together. This allows for simplification when speaking about complex specifications. It is important to remember that, in order to implement across platforms, you need to know the specification support level. You can read much more about WebServices Security, including linkage to the complete WS-Security 1.0 specification, at developerWorks®:

<http://www-128.ibm.com/developerworks/library/specification/ws-secure/>

13.1.4 WebSphere Application Server V6.0 supplied applications

Some of the default applications shipped with V6.0 are optionally installed, and some are installed by default. One of those installed by default is the WebSphere Application Server SamplesGallery. This is a set of small, generic samples that shows you how to perform common enterprise application tasks. For example, it demonstrates the use of session and entity enterprise beans, Java Database Connectivity (JDBC) access, connection pooling, Java Mail, message-driven beans, and other Web techniques. The gallery also contains descriptions of where to find additional samples and coding examples.

The SamplesGallery application includes:

- ▶ **Plants by WebSphere application**

This application demonstrates several Java 2 Platform, Enterprise Edition (J2EE) functions, using an online store that specializes in plant and garden tool sales.

The Greenhouse by WebSphere application can be used with this application to process its back orders.
- ▶ **WebSphere Bank application**

This application demonstrates multiple J2EE 1.4 technologies in a banking scenario.
- ▶ **Greenhouse by WebSphere application**

This application demonstrates multiple J2EE 1.4 technologies in a business-to-business (B2B) supplier scenario. It can process backorders from the Plants by WebSphere application.
- ▶ **Faces Client Tutorial - Sample Portfolio**

Sample Portfolio is a sample application that demonstrates the use of faces client components. The Hello world sample demonstrates how the faces client framework keeps a data model consistent in the browser.
- ▶ **Technology samples**

These samples demonstrate various core components in J2EE applications.
- ▶ **Web services samples**

These samples demonstrate J2EE beans and JavaBeans components that are available as Web services.
- ▶ **Service Data Objects (SDO) sample**

This sample demonstrates data access to a relational database through Service Data Objects (SDOs) and Java DataBase Connectivity (JDBC) Mediator technologies.
- ▶ **Java Adventure Builder application**

This application demonstrates J2EE 1.4 technology using an online travel shopping store.
- ▶ **JACL scripts**

These scripts enable you to configure resources and install the sample applications.
- ▶ **Programming model extensions samples in the Samples Gallery**

These samples demonstrate WebSphere programming model extension features such as dynamic query service, work area service, internationalization service, ActivitySessions service, application profiling, Java Transaction API (JTA) extensions, asynchronous beans, and scheduler.

13.1.5 WebSphere Application Server V6.0 on i5/OS packaging

WebSphere Application Server is available in three unique packages that are designed to meet a wide range of customer requirements. See Table 13-2. At the heart of each package is a WebSphere Application Server that provides the runtime environment for enterprise applications.

Table 13-2 Options

5733-W60	Description
Option 1	WebSphere Application Server V6 for OS/400 Express
Option 2	WebSphere Application Server V6 for OS/400
Option 3	WebSphere Application Server V6 Network Deployment for OS/400

The Express product is also delivered with i5/OS release V5R3M0 as part of the Web Enablement product (5722-WE2). This can be deselected from the order. WE2 includes WAS-Express V5.1 and WAS-Express V6.0. WE2 is a follow-on to 5722-WE1, which includes WAS-Express V5.0 and V5.1. Unlike on other platforms, Express is not limited to two CPUs on i5/OS.

These are covered in some depth in the Redpaper *WebSphere Application Server V6 Technical Overview*, available online at:

<http://publib-b.boulder.ibm.com/abstracts/redp3918.html?Open>

Version 6 for i5/OS does not install any additional products. There is a change from V5.x, which included WebSphere MQ product. The install features for all three editions are the same, including the Core product (application server runtime), samples, Javadoc™, and Web server plug-ins (for IBM HTTP Server for System i and Lotus Domino for System i), and the application client.

Version 6 for OS/400 media pack includes WebSphere Application Server Version 6 for OS/400, with the appropriate edition as listed in Table 13-2. We include IBM Business Solutions, which are an iSeries exclusive product. This collection of applications is designed to demonstrate the value of integrated solutions. They are easy to understand and simple to use. While they provide valuable services, their intent is to exemplify the use and integration of Web technologies on iSeries.

For whichever edition was ordered, you receive one CD per supported platform (for example, Windows, AIX, HP, Linux, Linux PPC). This CD contains the core product, IBM HTTP Server, Web server plug-ins, App client, and DataDirect JDBC drivers.

Application Server Toolkit (ASTK) has been dramatically improved in Version 6.0. It now provides the breadth of support and operation to allow it to be your single tool for development. It runs on Windows and Linux.

The Express edition also receives tools in the form of the Rational Web Developer and the WebSphere Developer Studio Client for iSeries 5.1.2. Be aware that this is only with Passport Advantage® orders.

The Network Deployment edition order also receives Edge components. These are features that enable load balancing, enhanced caching, and centralized security. Load balancing is enhanced through the employment of custom advisors and consultants. These on-board tools allow traffic monitoring and control, in the case of the advisor, and server performance in Nortel and Cisco environments with the consultant. A good overview is found at:

<http://www-306.ibm.com/software/webservers/appserv/was/network/edge.html>

We also include the DB2 Enterprise Edition, IBM Tivoli Directory Server, and a trial version of the Rational Application Developer.

Please note that the 32-bit JAM is not supported under Release 6.0. We have been told that support is expected in Release 6.1. We advise you to watch for an announcement, if this interests you.

WebSphere Application Server V6.0 product packaging

All three editions (Express, Base, and ND) are contained within the same product ID 5733-W60. The product is no longer packaged or installed as an i5/OS licensed program product (LPP).

Previous releases used RSTLICPGM to install the product under the covers (or explicitly in the Express case). In 6.0, we use the same packaging and installation mechanism as with other platforms. This utilizes the Macrovision InstallShield for MultiPlatforms (ISMP) 5.0.3. We include strict guidelines for product packaging so that the installer knows what to do.

Running WebSphere Application Server V6.0 prerequisites

Remember, times change, WebSphere changes, and your needs change. These prerequisites are minimums, and we strongly advise you to involve your IBM representative or Business Partner in sizing your system.

Install

To install, we recommend the following:

- ▶ 420 CPW or higher (minimum 300)
- ▶ 1 G or higher memory (minimum 768 M)
- ▶ OS/400 V5R2M0 or later
- ▶ IBM Developer Kit for Java 1.4.2 (5722JV1 option 6)
- ▶ Latest Java group PTF
- ▶ Qshell interpreter (5722SS1 option 30)
- ▶ Host servers (5722SS1 option 12), host servers started
- ▶ User profile executing the install must have *ALLOBJ and *SECADM authority

Runtime

For the runtime:

- ▶ Cryptographic Access Provider (5722-AC3) with the latest PTFs applied. Note that with i5/OS V5R4, 5722AC3 is no longer a separate product. Its capabilities have been integrated into the base i5/OS software. See 6.4, “Cryptographic support” on page 141, for related information.
- ▶ If you plan to use SSL with external Web server *or* with security and LDAP, you need OS/400 Digital Certificate Manager (5722-SS1 option 34).

Tip: You can find the latest WAS PTFs at:

<http://www-03.ibm.com/servers/eserver/series/software/websphere/wsappserver/services/service.html>

Runtime i5/OS subsystem

All WebSphere Application Server V6.0 *profiles* (called instances prior to V6.0) run in the i5/OS-supplied subsystem QWAS6 contained in library QWAS6.

Saving the product, configuration, and data

Starting with Version 6.0.2, SAVLICPGM and RSTLICPGM are supported for WebSphere Version 6. Do not try to uninstall using DLTLICPGM, as there is an uninstall routine in the IFS.

Besides saving the product, you must consider backing up the servlets and JSP-related data, such as your class/source files and HTTP configuration. The profile configuration and properties folders in your profile are important. You also need to get the EJB-related data, databases, and EJB classes, which are installed in the Apps subdirectory.

Backing up security data like your key files, validation lists, and so on, is vital.

WebSphere Application Server for i5/OS V6.0 additional documentation

Use the following Web Site as a starting place for specific i5/OS-based documentation for WebSphere Application Server V6.0:

<http://www-03.ibm.com/servers/eserver/series/software/websphere/wsappserver/docs/docws60.html>

Figure 13-4 shows an excerpt of the main Web page listing additional documentation sources.

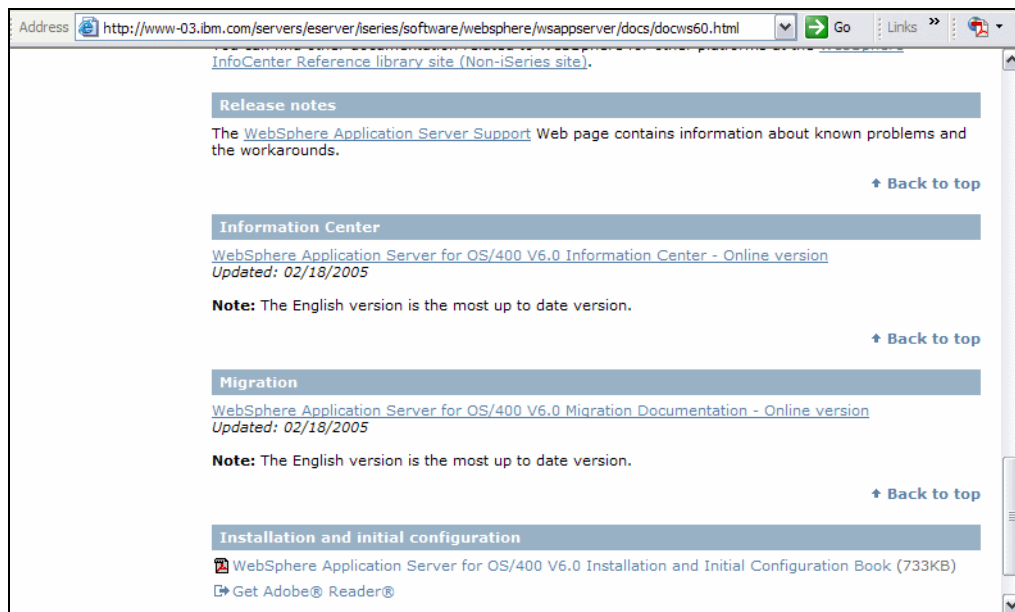


Figure 13-4 WebSphere Application Server V6.0 documentation sources

The following are more direct URLs for important Web sites for additional information about WebSphere Application Server V6.0 functions across the various V6.0 products and supported operating systems:

- ▶ WebSphere Information Center

<http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp>

- ▶ WebSphere function comparison

http://www-1.ibm.com/servers/eserver/series/software/websphere/wsappserver//product/was_60_comparison_table.html

13.2 WebSphere Application Server Version 6.1 on i5/OS

On April 11, 2006, IBM announced Version 6.1. There are a number of enhancements in Version 6.1, which became supported under i5/OS during June 2006. With V 6.1 the process of unifying WebSphere Application Server code and processes across z/OS, i5/OS, and distributed platforms has been completed. Figure 13-5 represents a history of primary functions, Java Development Kit (JDK) levels and processes available from WebSphere Application Server from Version 4 through Version 6.1 with an indication of potential areas of enhancements in future versions.

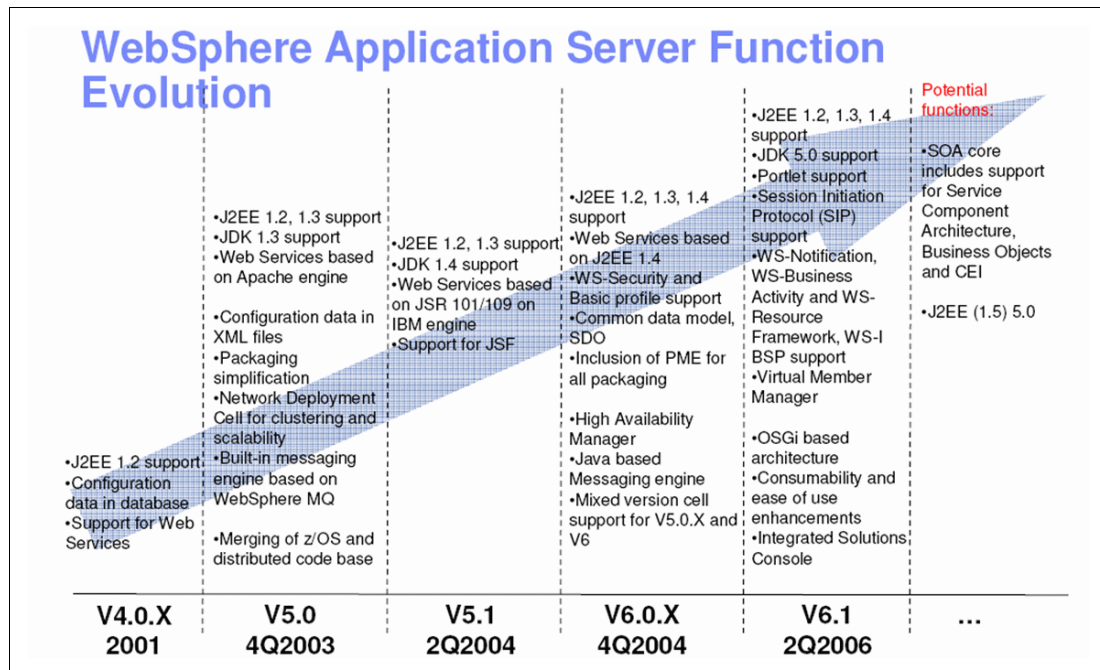


Figure 13-5 WebSphere Application Server: evolution of major functions and processes

In the V6.1 column you see the first mention of OSGi, an open system architected programming interface continuously evolving from the OSGi Alliance. The OSGi Alliance is an independent non-profit corporation comprised of technology innovators and developers and focused on the interoperability of applications and services across the Internet.

The initial group of companies that formed the initiative were Alcatel, Cable & Wireless, Electricite de France, Enron Communications, Ericsson, IBM, Lucent Technologies, Motorola, NCI, Nortel Networks, Oracle, Philips Electronics, Sun Microsystems, Sybase, and Toshiba.

Open Service Gateway Initiative (OSGi) is an evolving industry plan that is Java based. The plan is primarily focused on setting up a standard way to connect devices such as home appliances and security systems to the Internet. With such a standard, home users could, for example, install a security system and be able to change from one monitoring service to another without having to install a new system of wires and devices.

A home or small business's network of devices would connect over the Internet to a *service gateway* application server located in a computer. The OSGi architecture specifies the Java-based application program interface (API) for programmers to use to allow communication and control between service providers and the devices within the home or small business network.

OSGi is an open standard programming interface. Changes evolve through the normal Java community process.

WebSphere Application Server V6.1 embraces OSGi-based architecture interfaces, but application developers must code to this interface. For more OSGi information search <http://www.google.com> or <http://www.ask.com> or <http://www.what-is.com>, or the OSGi Alliance Web site at:

<http://www.osgi.org>

The following topics highlight the key enhancements included with WebSphere Application Server V6.1 on i5.OS. For more complete information about Version 6.1 running under i5/OS than presented here, see the book *WebSphere Application Server for i5/OS Handbook: Version 6.1*, SG24-7221.

13.2.1 Multiple WebSphere Application Server versions coexistence

All WebSphere Application Server versions and editions can co-exist on the same i5/OS, as long as the product is supported on that version of i5/OS. This allows for easier migration and upgrading between versions and editions.

The only exception to this is that WebSphere Application Server Version 6 Base edition and Version 6 Express edition cannot coexist on the same i5/OS.

WebSphere Application Server Express compliance with J2EE

Starting with Version 6, the WebSphere Application Server - Express edition is fully J2EE compliant (with Java JDK 1.5 support).

13.2.2 JDK 5.0 support

Applications running under WebSphere Application Server V6.1 can be using Java 2 Platform, Standard Edition (J2SE) 5.0 supported interfaces and functions. For more information about Java support see 12.6, "Java" on page 283.

13.2.3 WebSphere Application Server installation changes

With V6.1 there are several changes and enhancements in the installation process. The primary changes are listed below:

- ▶ You can now install WebSphere Application Server Version 6.1 for i5/OS using CD or DVD media.
- ▶ Security is now enabled by default at installation time. For a silent or graphical installation to succeed, you must either provide a user ID and password or disable security.

- ▶ A launchpad is now available when installing remotely onto an i5/OS from a Windows workstation.
- ▶ Using newly available installation routines for the base WebSphere Application Server Version 6.1 products, Web Server plug-ins, and application clients enables you to set up installation procedures to install only what you need into an i5/OS partition.

The Web server plug-ins are still included with the application server product. If you are running your application server and Web server on the same i5/OS, you do not need to install the Web server plug-ins.

- ▶ The Installation Factory is new with V6.1. The Installation Factory is on the Tools CD shipped with WebSphere Application Server or you can download it separately from the IBM Installation Factory Web site. With this toolkit you can create a customized, distributable installation image tailored to your specific needs. You can build an image of the base Websphere Application Server, related fixes, and applications unique to your WebSphere environment. In connected mode, you can create the installation packages on one operating system and distribute them to another system.

See “Application Server Toolkit” on page 332 for more information about the installation factory.

- ▶ The standard ISMP option format (-W option=value and -P option=value) is replaced with an -OPT option that takes sub-options whose names are simpler and more meaningful, and less subject to change from one release to another. These same options are used in the command line invocation of the wizard and in the silent installation parameters in the response files.
- ▶ Removal and reinstallation of application server environments is simplified. Uninstalling an application server profile does not change the shared system files of the product. Now you can leave profiles installed when uninstalling the product and its core product files.
- ▶ WebSphere Application Server Version 6.1 no longer uses MQSeries® to support the embedded Java Message Services (JMS) function. You can use either a separate WebSphere MQ license or the new service integration bus architecture for messaging services. If you had previously set up the integrated JMS under Version 5 or Version 5.1, you can continue to use any JMS servers on nodes that have not yet migrated to Version 6.1.

13.2.4 WebSphere Application Server administration

You can use the administrative console to display and change your WebSphere Application Server configurations and to manage your WebSphere Application Server resources.

The Version 6.1 administrative console has simplified administration with simpler panels and more wizards. It has been integrated into the Integrated Solutions Console (ISC), an IBM sponsored open systems management console architecture. Using the ISC architecture enables a different look and feel to the administration console user. The updated Version 6.1 administrative console also has new features that allow you to customize your window layout. This includes a My Tasks page that allows you to create and edit a list of tasks you defined to be viewed via the console navigation.

The port number for the administrative console has changed. It is now as follows:

<http://your.server.name:9060/ibm/console>

Command assistance in the administrative console maps your administrative activities to wsadmin scripting commands. This way you can capture your console dialogue and apply it to wsadmin. Using command assistance, you can view wsadmin scripting commands in the Jython language for the last action run in the administrative console. Jython ships with

WebSphere Application Server V6.1 and is the preferred scripting language for all new work in this area.

Version 6.1 also introduces an administrative console wizard that can update deployed applications or modules in the following ways:

- ▶ Replace an entire application (.ear file).
- ▶ Replace, add, or remove a single module (.war, EJB .jar, or connector .rar file).
- ▶ Replace, add, or remove a single file.
- ▶ Replace, add, or remove multiple files by uploading a compressed file.

If the application is updated while it is running, WebSphere Application Server automatically stops the application or only its affected components as required, updates the application logic, and restarts the stopped application or its components.

Previous versions of WebSphere Application Server only supported the replacement of an entire application and always stopped and restarted the entire application for any change.

For more information updates and customization of the Version 6.1 administrative console, see the following topic in the WebSphere Application Server Version 6.1 for i5/OS Information Center:

http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/topic/com.ibm.websphere.base.iseries.doc/info/series/ae/tcon_console.html

13.2.5 WebSphere Application Server security enhancements

This describes most of the security functions and ease-of-use enhancements included with V6.1:

- ▶ By default, administrative security can be enabled *out of the box* during installation. You can also enable administrative security when you create a profile with the manageProfiles Qshell script.
- ▶ Server IDs are automatically generated.
- ▶ WebSphere key and certificate management have been simplified.
- ▶ You can federate various repositories so that you can manage them as one entity.
- ▶ The Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) protocol is now supported, which allows flowing Kerberos tokens from Web browsers such as Mozilla, FireFox, or Internet Explorer.
- ▶ The WS-I Basic Security Profile 1.0 is now supported to improve interoperability with other vendors.
- ▶ Web authentication can now be performed with or without Web authorization. A Web client's authenticated identity is available whether or not Web authorization is required. There is enhanced control over Web authentication behavior.
- ▶ More administrative roles have been defined to provide degrees of authority that are needed to perform certain administrative functions.
- ▶ Administrative access can be granted to each user per resource instance.
- ▶ There is now hardware cryptographic device support for Web services security, which can be used to accelerate the cryptographic operations. Cryptographic keys can be stored on the hardware cryptographic device and never leave the device.

13.2.6 Runtime i5/OS subsystem

All WebSphere Application Server V6.1 profiles (called instances prior to V6.0) run in the i5/OS-supplied subsystem QWAS61 contained in library QWAS61.

13.2.7 Application development enhancements

New in Version 6.1 is support for the JSR 168 Portlet programming model, the JSR 116 Session Initiation Protocol (SIP) Servlets specification, and the Java Server Faces (JSF) Widget Library, supported through JSF 1.1.

The Portlet programming model is supported through wizards to create portlet projects and portlets, an editor to manage the portlet deployment descriptor (portlet.xml), and the ability to import a portlet .war file.

Portlets are now treated similar to servlets, and you can enable direct access to portlet URLs (just like servlets).

WebSphere Application Server Version 6.1 for i5/OS includes JSR 116 SIP Servlets built in a converged servlet container. Servlets can share a single application session across protocols. HTTP Servlets and Portlets can create SIP messages and tie state together with a single application. SIP tooling is also included.

It is predicted that the next generation of Web applications will use SIP. Applications that might be included in this next generation include:

- ▶ Voice over IP (VoIP)
- ▶ Video conferencing with the video over IP
- ▶ Collaborative applications
- ▶ Gaming applications

Refer to Appendix E, “System i IP Telephony powered by 3Com” on page 529. This product provides several of the applications listed above and uses SIP.

The JSF Widget Library (JWL) is a set of functions that are pre-built for reuse within a user interface and integrates widgets from a number of sources. It includes the JSF components from Rational Web Developer (RWD), with the exception of the base JSF components, which are already included in the Application Server runtime. This includes the IBM extended JSF components and the extended Odyssey components.

Essentially, JWL functions are intended to help create interactive Web pages and can improve overall transaction performance by reducing round-trip interactions with the host server.

AST enhancements

In previous versions of WebSphere Application Server, you deployed an application and set up its required configuration in two separate steps. In Version 6.1, AST enables you to define the required configuration (such as a data source) as a part of the application. At deployment, you can choose to process the embedded configuration data, which automatically sets up the required configuration for the application.

In Version 6.0, AST used the Eclipse Update Manager to retrieve and apply maintenance. In Version 6.1, it uses the Rational Product Updater (RPU), which is built on top of the Update Manager. RPU is able to update the JRE™ that is bundled with AST, whereas Update Manager knew only how to update Eclipse plug-ins. RPU runs as a separate process from the code it updates, whereas Update Manager runs within AST itself.

Application Server Toolkit

The Application Server Toolkit (WATK) in prior versions only allowed you to assemble applications. It has been significantly improved, and is now a fully tooled package.

The WATK provides you with basic assembly and deployment tooling for publishing to your application server, such as WebSphere Application Server Version 6.0. You can also use the tool to perform basic unit testing, debugging, and profiling functions. It supports new features that are defined in the EJB 2.1 and J2EE 1.4 specifications.

This was built using Eclipse v3.0 technology with features for team programming, debugging, J2EE application deployment, and more. Eclipse is an open and industry-supported platform for development tools. The WATK also provides a familiar interface for developers experienced with the Rational Developer family of products. Those products extend the capabilities of the Application Server Toolkit, such as supporting the creation of new applications.

There are new options in the areas of administration, serviceability, messaging, and installation, to ensure more productive and less error-prone environments. One tool in particular is the Installation Factory. The Installation Factory combines the installation image for a version or release of a WebSphere software product with applicable maintenance packages, a configuration archive, one or more enterprise archive files, customization scripts, and other files, to create a customized installation package.

For example, the Installation Factory can combine Version 6 of a WebSphere Application Server product with Refresh Pack 2 to create a customized installation package that installs V6.0.2.

Installing and configuring a WebSphere software product is usually a multiple step process:

1. Install the shipped version of the product.
2. Install the current refresh pack.
3. Install the current fix pack.
4. Install a Java 2 Software Development Kit (SDK) fix pack.
5. Install one or more interim fixes as needed.
6. Create and configure application servers and other artifacts.
7. Deploy applications.

The Installation Factory simplifies the process by creating a single installation image, the customized installation package (CIP) (see Figure 13-6). This package incorporates all of the above steps not as an aggregate, but as an integrated image.

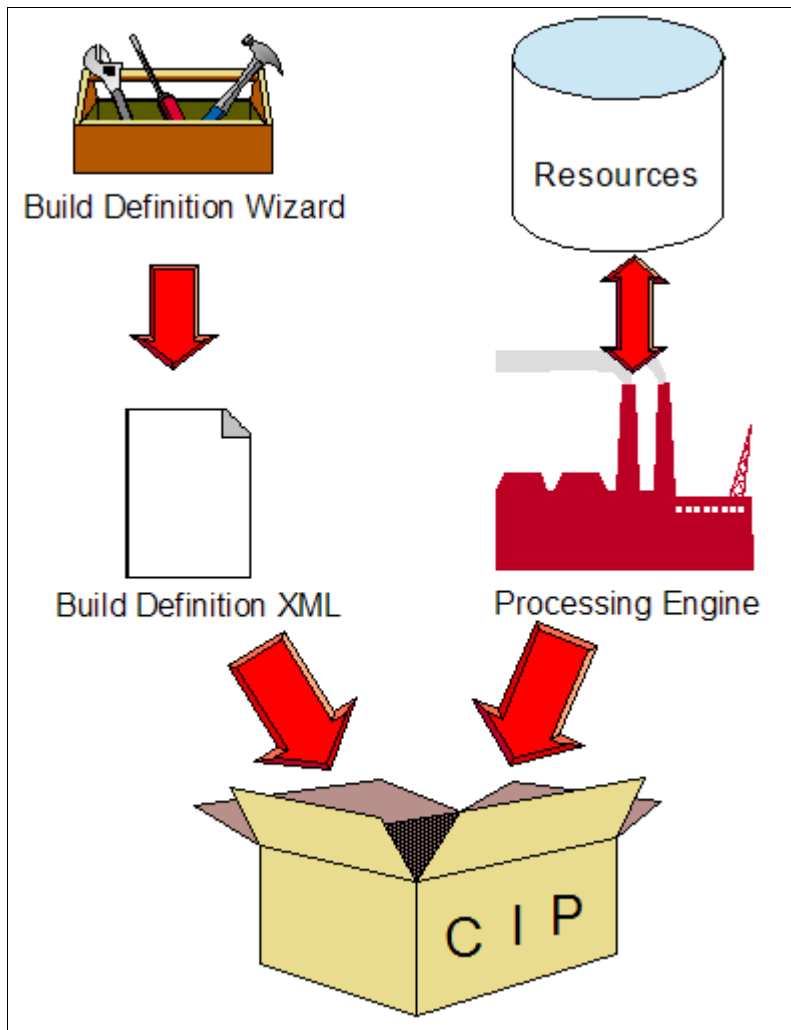


Figure 13-6 Custom installation package

After creating the customized installation package, installing and configuring a WebSphere software product is a one-step process: Install the customized installation package. This makes an easy-to-compile process, which allows you to create a fix level updater for all of your systems.

WebSphere Application Server for i5/OS V6.1 additional documentation

Use the following Web Site as a starting place for specific i5/OS-based documentation for WebSphere Application Server V6.1:

<http://www-03.ibm.com/servers/eserver/iseries/software/websphere/wsappserver/docs/docws61.html>

Figure 13-7 shows an excerpt of this Web page, listing additional documentation sources.

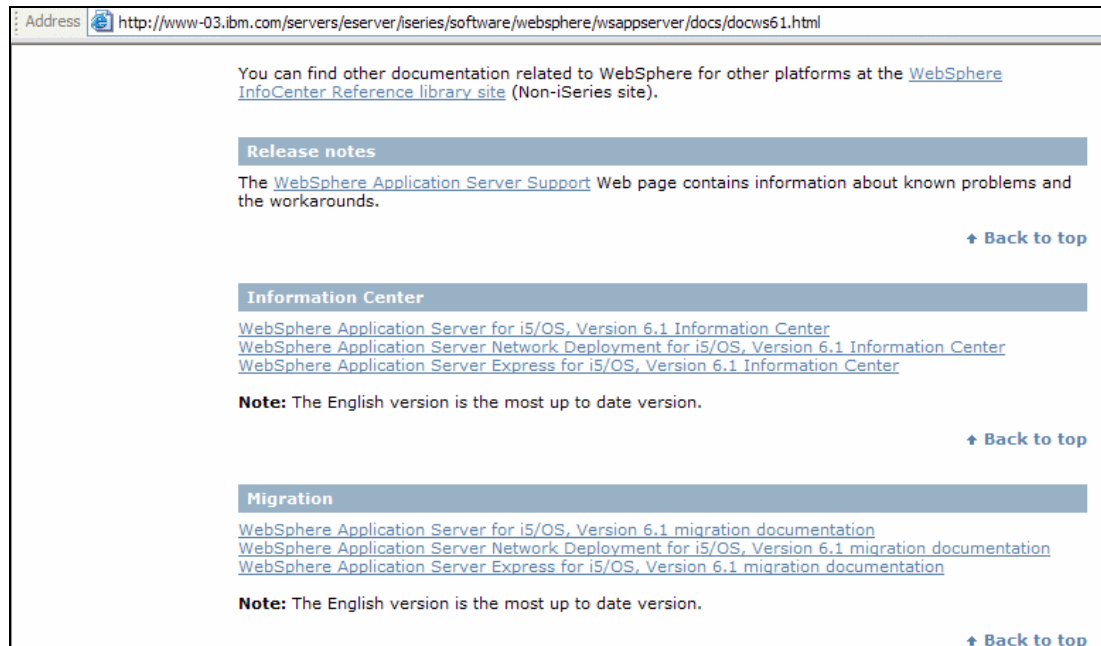


Figure 13-7 WebSphere Application Server V6.1 documentation sources

The following are more direct URLs for important Web sites for additional information about WebSphere Application Server V6.1 functions across the various V6.1 products and supported operating systems:

- ▶ WebSphere Information Center
<http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>
- ▶ WebSphere function comparison
http://www-1.ibm.com/servers/eserver/iseries/software/websphere/wsappserver//product/was_61_comparison_table.html

13.3 IBM Express Runtime

Express Runtime enables you to integrate your WebSphere Application Server-based application into an install process that first installs the following IBM middleware and then deploys your application solution, which uses the just installed middleware.

For information beyond what is contained in this section see IBM Redbook *Express Runtime V2.1.1*, SG24-7193.

13.3.1 Middleware included in Express Runtime V2.1.1

The following software is included in Express Runtime V2.1.1 (latest release as of 4Q 2006):

- ▶ Choice of two Database Management Systems (DBDMs):
 - IBM DB2 Universal Database Express Edition for Windows and Linux 8.1
 - Informix® Dynamic Server Express (IDS) 10.00

When using Express Runtime to install onto i5/OS or any other IBM operating system with IBM DB2 Universal Database already installed, this installation step is bypassed.

- ▶ IBM WebSphere Application Server Express 6.0.2

When using Express Runtime to install onto i5/OS or any other IBM operating system with IBM WebSphere Application Server Express 6.0.2 already installed, this installation step is bypassed. WebSphere Application Server for i5/OS Express is included as a no-charge option when you order i5/OS V5R3 or V5R4.

- ▶ IBM HTTP Server 6.0.2.

When using Express Runtime to install onto i5/OS or any other IBM operating system with IBM HTTP Server 6.0.2 already installed, this installation step is bypassed. WebSphere Application Server for i5/OS Express is included as a no-charge option when you order i5/OS V5R3 or V5R4.

- ▶ IBM Web server plug-in for deploying your application into the appropriate Web application server.

13.3.2 Express Runtime V2.1.1 development and deployment tools

In addition to the middleware previously listed, Express Runtime also includes the following tools for software solution development and deployment:

- ▶ Express Runtime Developer: An Eclipse-based plug-in for solution development, including installation and configuration through custom editors, Express Runtime helps you offer a complete solution.

- ▶ IBM Rational Web Developer: An Eclipse-based full featured Integrated Development Environment (IDE). It is designed for building, testing, and deploying Java and Java 2 Platform, Enterprise Edition (J2EE) applications. It provides integrated development support for building J2EE applications with HTML pages, servlets, and JavaServer Pages (JSPs).

- ▶ Deployment Wizard: The deployment wizard installs a solution on one or more target systems. The deployment wizard offers an interface that helps you deploy a solution task-by-task. You can customize the application solution installation process to display your own product titles, headings, logo, and so forth to demonstrate the integration of your solution with IBM middleware.

- ▶ Console for Express Runtime: This stand-alone Web-based console provides the administrative interface to all middleware components that are parts of IBM Express Runtime V2.1.1. This support is available across multiple machines and platforms. The console now has integrated common logging in its console capabilities, and the ability to automatically find a managed server to assist in application management and troubleshooting.



i5/OS-based Web enablement enhancements

In this chapter we cover the set of Web enablement capabilities announced during October 2006. Functions included in this chapter are delivered through the already available no-charge product Web Enablement for i5/OS, 5722-WE2. 5722-WE2 is a no-charge set of functions that can be ordered with any order for i5/OS V5R3 or V5R4. New support starting in October 2006 includes the following. Some of these are delivered via PTFs.

- ▶ iSeries Access for Web and iSeries Navigator tasks on the WebTasks on the Web can now perform their functions using an integrated Web application server, which is included with the PTFs that provide the new Web Enablement for i5/OS enhancements.
- ▶ A Web performance advisor that can be accessed from the i5/OS IBM Web Administration for i5/OS Web pages. This tool looks at selected system values and any active WebSphere Application Server parameters and identifies recommended changes, if any, based upon your current system workload.
- ▶ A Web Performance Monitor that can be accessed from the i5/OS IBM Web Administration for i5/OS Web pages. While activated, this monitor accounts Web application transaction counts, average transaction response time, and average CPU utilization.
- ▶ A new IBM Express Runtime Web Environments for i5/OS product (5733-SO1). This product expedites installation of an IBM Web Enablement Environment, thereby reducing the complexity of moving to and using a Web environment for your applications. This product expedites the install of HTTP Server for i5/OS and WebSphere Application Server V6.0 for i5/OS, which serve the iSeries Access for Web application and an example airline flight reservation application that is modernized using HATS and WebFacing and is accessible as a Web service.
- ▶ New WebSphere Application Server performance data collected by i5/OS Collection services.
- ▶ A simplified way to configure Secure Sockets Layer (SSL) for communicating between a Web browser and the IBM Web Administration for i5/OS server.

14.1 iSeries Access for Web, iSeries Navigator tasks on Web, and new integrated Web application server

With the PTFs available October 2006 through February 2007, i5/OS iSeries Navigator tasks on the Web and Series Access for Web are enabled to use a new Web application server integrated with i5/OS. Once the appropriate PTFs are installed, this Web application server is started when the IBM-supplied HTTP Administration server, IBM Web Administration for i5/OS, is started. This HTTP server is the one started via the Start TCP Server (STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)) command.

This application server becomes available during October 2006 for V5R3 and V5R4. It supports two IBM-supplied i5/OS Web applications as follows:

- ▶ iSeries Navigator tasks on the Web, V5R4 only. PTF SI24255 is required and available November 2006. Once this PTF has been installed, iSeries Navigator tasks on the Web will always run under this application server. Without this PTF installed, iSeries Navigator tasks on the Web can run only using the i5/OS-supplied WebSphere Application Server Version 6.0 Express profile/instance named SYSINST/ADMIN. For more information about SYSINST/ADMIN refer to 9.1, "iSeries Navigator tasks on the Web" on page 190.

This means that PTF SI24255 installed acts like a switch for iSeries Navigator tasks on the Web. Once installed, iSeries Navigator tasks on the Web uses the integrated Web application server. Uninstalled, iSeries Navigator tasks on the Web uses the SYSINST/ADMIN WebSphere Application Server profile supplied with i5/OS V5R3 or V5R4.

- ▶ iSeries Access for Web V5R4. This application has a simple one-step configuration parameter to enable it to run with the new integrated Web application server. You can choose to run Series Access for Web under this new Web application server or select a WebSphere Application Server instance or profile that you have previously created for a supported WebSphere Application Server version installed on your partition.

The servlet environment supports the integrated Web application server environment. New configuration examples have been added to the *Configuring iSeries Access for Web in a Web application server environment* topic under iSeries Access for Web on V5R4 iSeries Information center:

<http://publib.boulder.ibm.com/infocenter/iseriew/v5r4/index.jsp>

Select **Connecting to iSeries** → **Choose your connection** → **iSeries Access** → **iSeries Access for Web**.

A V5R4 PTF is required. To find out the latest PTFs refer to:

<http://www.ibm.com/eserver/iseriew/access/web/servicepacks.htm>

This support of the new integrated Web application server gives you the ability to be productive using the iSeries Access for Web functions and the iSeries Navigator tasks on the Web functions without having to install and configure a WebSphere Application Server product or an Apache server on i5/OS.

The integrated Web application server runs as a job under the i5/OS-supplied HTTP Administration server (*ADMIN). We used the Windows-based iSeries Navigator **Work Management** → **Subsystems** → **Active subsystems** → **qhttpsvr** to identify this job under the Function column, as shown in Figure 14-1.

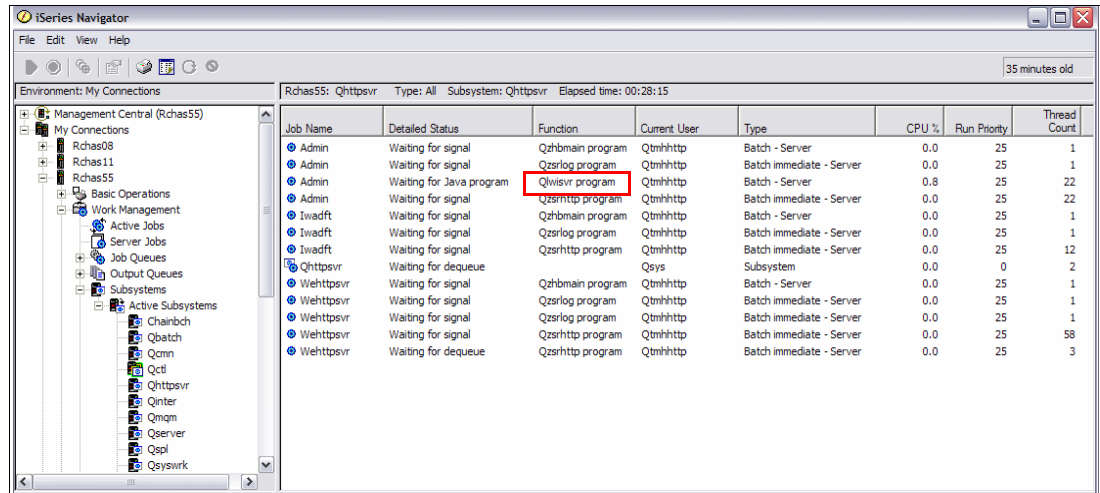


Figure 14-1 Manage all HTTP servers with the integrated Web application servers

For more information about V5R4 enhancements for iSeries Access for Web refer to 9.2, “iSeries Access for Web enhancements” on page 207.

14.2 IBM Express Runtime Web Environments for i5/OS

IBM Express Runtime Web Environments for i5/OS (5733-SO1) provides an integrated package of software products that you install and configure to your i5/OS, according to your needs, using a wizard. The main advantage of IBM Express Runtime Web Environments for i5/OS is that it is an easy way to install all of the components necessary for a working Web environment. It is ideal to get a customer who is considering or just starting out running a WebSphere Application Server-based application quickly up and running with an example environment.

IBM Express Runtime Web Environments for i5/OS is offered as a feature of IBM Web Enablement for i5/OS (5733-WE2), which is shipped by default when you order i5/OS V5R4. All WebSphere Application Server products can be ordered through the 5733-WE2 product.

IBM Express Runtime Web Environments for i5/OS is distributed by via CD or DVD. These CDs or DVDs can only be run on a Windows workstation. The Windows workstation must be connected to the installation target i5/OS partition.

Alternatively, IBM Express Runtime Web Environments for i5/OS can be ordered electronically through the IBM Electronic Software Delivery Web site. For more information see:

<http://www.ibm.com/servers/eserver/ess>

IBM Express Runtime Web Environments for i5/OS is available free of charge, because you are already entitled to V5R4 embedded products at no cost.

Note: In this chapter, as in other chapters, we use the term *i5/OS partition* to mean both:

- ▶ A single partition within an IBM System i configuration running i5/OS when more than one partition may be configured.
- ▶ The entire IBM System i configuration running i5/OS when there are no specifically configured partitions. That is, the system is one i5/OS partition that owns all hardware.

By running one or more of the supplied software products and sample applications you can:

- ▶ Quickly become familiar with the Web application server environment.
- ▶ Gain initial knowledge in the use of IBM 5250 refacing products (Host Access Transformation Services (HATS), WebFacing, and iSeries Access for Web) and see how Web Services can work.

A sample 5250-based flight400 application is installed that you can run as originally written to a 5250 workstation as well as run as customized HATS or WebFaced versions from a browser. Running this provided application can help you better understand ways to customize your own 5250 applications with either HATS, WebFacing, or the extended capabilities available through the WebSphere Development Studio for iSeries Client application that installs on a PC workstation.

- ▶ Gain initial knowledge in the use of IBM iSeries Access for Web product, which provides a browser-based 5250 refacing interface and a powerful subset of iSeries Access for Windows functions via a browser-based interface. See 9.2, “iSeries Access for Web enhancements” on page 207, for V5R4 level information about iSeries Access for Web.

5733-SO1 provides the following:

- ▶ Installs V5R4 iSeries Access for Web, 5722-XH2. This includes installing PTFs.
- ▶ Installs WebSphere Application Server - Express V6.0 for OS/400, 5733-W60. This includes fix pack level 6.0.2.9.

WebSphere Applications Server V6.0 for i5/OS PTF group level 7, available during October 2006, is at WebSphere Application Server - Express V6.0.15 level. You can run all of the October 2006 Web enablement enhancements with 6.0.2.9 except the Web Performance Monitor and Advisor. At the time this book was published, the Monitor and Advisor required 6.0.2.15 or later.

- ▶ Installs V5R4 IBM HTTP Server, 5722-DG1. This includes installing PTFs.
- ▶ Configures a WebSphere Application Server profile (corresponds to an instance in pre-V6.n versions of WebSphere Application Server) and an associated HTTP server instance, generates the WebSphere Application Server plug-in, and starts both servers.
- ▶ Configures iSeries Access for Web to run in that WebSphere Application Server profile.
- ▶ Installs and configures an icon on the Windows workstation desktop that links to a FirstSteps static HTML Web page. The FirstSteps Web page is designed to provide an easy way for you to get started using what the wizard deployed to your i5/OS.

Included in this product is a set of sample Web applications, ready for you to run. Each Web application modernizes (refaces) the flight400 application, a well-known sample RPG, Control Language (CL) application that uses an i5/OS Display file to format the 5250 screen dialogue with the 5250 workstation user.

Figure 14-2 shows one of the flight400 5250 screens that is processed by the WebSphere Application Server applications shipped with IBM Express Runtime Web environments for i5/OS.

```

Flights Reservation System - Create Order      13:49:43  3/09/04      AS07

Type choices, press F10 to Make Reservation

-----
|          FLIGHT INFORMATION          |          TICKET ORDER INFORMATION          |
|-----|-----|
| Airline:      Flight: 0000000      | Order Number.....:      PENDING      |
| Date of Flight...: 00 00 2004      | Customer...:              |
| From City.:   |          Class of Service - First.....:  |
| Depart Time.....:          Business.....:  |
|              |          Economy.....:      X      |
| To City...:   |          Number of Tickets.....: 01      |
| Arrival Time.....:          Price $.....:  |
|              |          Tax $.....:              |
|              |          Total Due w/ Tax $.....:      |
|-----|-----|
F2=Refresh  F4=FROM Cities  F5=T0 Cities  F6=Flights  F7=Customers
  
```

Figure 14-2 Example flight400 application screen using the 5250 interface

By using the IBM-supplied HATS and WebFaced versions of the flight400 application you can consider ways to modernize your 5250 applications to be accessible over the Internet. You can observe some of the advantages of JavaServer Pages (JSPs) and their follow-on constructs to make your applications available to an expanded set of users over the Internet.

For more details on the flight400 applications see *Mastering the IBM WebFacing Tool*, SG24-6331.

The included sample Web applications are as follows:

- ▶ Host Access Transformation Services (HATS)

A sample application that demonstrates the presentation of an RPG program named flight400, which was modernized using HATS.
- ▶ WebFacing

A sample application that demonstrates the presentation of flight400, which was modernized using WebFacing. The WebFacing sample application requires Microsoft Internet Explorer V5.5 or later as the Web browser. This application does not work with any other Web browser.
- ▶ Web services

This sample application queries flight400 by accessing it using a Web service. The first letter of the name being queried always begins with a capital letter. For example, a query of the name *a* does not produce a result, whereas a query of *A* does.

For a short definition off Web services refer to 12.1.5, “Web services” on page 273.

Important: The installation CD or DVD provides a readme file that must be read for successful installation and deployment of the IBM-provided applications. This readme file provides important information not included in this book.

You can run with the fix levels (or no fix level specified, such as for TCP) specified in the readme file. However, we recommend using the fix levels listed here.

14.2.1 Hardware and software requirements for 5733-SO1

The following are i5/OS partition and PC workstation requirements to take advantage of the products and applications installed using IBM Express Runtime Web Environments for i5/OS. We use the term *partition* here to also include a System i5 configuration running i5/OS with no partitions (that is, a system-wide partition).

► i5/OS partition

– i5/OS V5R4 (5722-SS1)

You should use the IBM Workload Estimator to ensure proper sizing of processor, memory, and disk configuration, depending upon how many users of the IBM-supplied applications are active at the same time.

You can run WebSphere Application Server on an i5/OS partition with a minimum of 1000 CPW, 4 GB of memory, and 4 disk drives, with one user active. However, with multiple applications or multiple active users reasonably busy, you will need more hardware resources (for example, 2000 or more CPW, 8 GB or more of memory, and 8 or more disk drives).

Use the IBM Workload Estimator and input workload parameter values you have thoughtfully estimated.

- Extended Base Directory Support (5722-SS1, option 3)
- AFP Compatibility Fonts (5722-SS1, option 8)
- Host Servers (5722-SS1, option 12)
- Qshell (5722-SS1, option 30)
- Java Developer Kit 1.4 (5722-JV1, *BASE, option 5, and option 6)
- Toolbox for Java (5722-JC1)
- TCP/IP Connectivity Utilities (5722-TC1)
- Configured for TCP/IP and DNS
- iSeries Access Family (5722-XW1)

This must be installed on the i5/OS system where iSeries Access for Web is accessing resources.

– IBM Query for iSeries (5722-QU1)

If you want to run reports using the HATS or WebFacing sample applications.

– Recommended PTFs:

- i5/OS Cumulative PTF Group SF99540 Level 6066 or later
- DB2 Universal Database for iSeries PTF group SF99504 Level 4 or later
5733-SO1 can run at level 2.
- Java PTF Group - SF99291 Level 4 or later
5733-SO1 can run at level 2.

- TCP/IP PTF Group - SF99315 Level 2 or later
5733 can run with no PTF group.
- IBM HTTP Server for i5/OS - 99114 Level 4 or later
5733-SO1 installs this level.

Note: Over time, later PTF group levels than those listed here become available. We strongly recommend, when running WebSphere Application Server-based Web applications, that you install the latest PTF group levels periodically.

- ▶ PC workstation
 - One of the following Windows operating systems
 - Windows XP Professional SP2
 - Windows 2000 Server SP4
 - Windows 2000 Advanced Server SP4
 - Windows 2000 Professional SP3
 - Windows Server 2003, Standard Edition SP1
 - Windows Server 2003, Enterprise Edition SP1
 - Configured to use TCP/IP and DNS
 - One of the following Web browsers
 - Windows Internet Explorer 6 or later
 - Firefox 1.5 or later

Note: Windows Internet Explorer 6 or later is required to run the WebFaced version of flight400. Other browsers may not work.

- ▶ Workstation hardware minimums
 - 512 MB of memory. 1 GB recommended.
 - Intel Pentium® III class processor with a minimum clock speed of 600 MHz. Intel Pentium IV class processor with a minimum clock speed of 1.2 GHz is recommended.
 - Local area network (LAN) connection.
 - SVGA monitor with a minimum 1024 x 768 video resolution configured to display a minimum color depth of 256 colors.
 - 1.5 GB free disk space.

14.2.2 Installing the Web Enablement Environment

When installed on the system, the entire HTTP Server and WebSphere Application Server is referred to as the Web Enablement Environment. We do not show the install steps and windows in this book. However, here is the summary of the install steps:

1. Place the first CD or the DVD in a Windows workstation attached to the same network as your target i5/OS partition.

Important: Before starting the wizard you must consider any firewalls on your PC workstation and any between your workstation and the target i5/OS partition. There may be IP ports used by the installation process that are blocked by the firewall. In some cases the firewall blockage may not be obvious, and the installation process appears to be continuing when it is not.

Watch the PC workstation for any firewall messages and respond promptly. Also watch progress bar percentages on the installation wizard closely for no increase in progress percentage after several minutes. The initial progress indication is for work done solely on the PC workstation. Then the process on the i5/OS partition starts. It is at this time that a firewall problem could occur. The installation process has an overall installation timer that could eventually time out if you have this problem and do not detect it.

2. Follow the wizard screens, which include accepting the license agreements for the software being installed and specifying the i5/OS target TCP/IP host name and user ID and password used to perform the install. This user ID must have *SECOFR authority.
3. Specify your own names for the automatically created HTTP server and WebSphere Application Server Express V 6.0 instance as a result of the installation.

For the examples used in this chapter we specified WEHTTPSVR and WEAPPSVR, and we let the installation processes select the ports that would be used by WEHTTPSVR and WEAPPSVR. In our example the WEHTTPSVR port became 10000. You can see the ports used using the IBM Web Administration for i5/OS browser interface for the HTTP servers tab and then the Application Servers tab.

When the installation is done, the HTTP server and the associated instance are up and running with iSeries Access for Web, the flight400 5250 application, the flight400 WebFaced application, the flight400 HATS application, and the flight400 Web Services application set up, and the WebFacing server under i5/OS are all up and ready to run.

4. During the installation process progress messages are displayed on the PC workstation's Deployment window. Anytime during the process and after it has completed you can also view more detailed log messages.
5. If you have any of the products already installed, such as the HTTP Server for i5/OS, that is noted in a message and progress continues. Depending on the versions and fix levels of any software already installed, the deployment process provides updates through the installation process.
6. Depending upon your network activity, existing workloads on the i5/OS partition, and how many of the products being installed are already on your system, the installation process can take an approximate minimum of 65 minutes up to approximately 4 or more hours.
7. Observe the 100% complete message. At that time your desktop has a shortcut to the FirstSteps Web page shown in Figure 14-3 on page 346.

14.2.3 Starting to use the Web Enablement Environment

When finished using the installation wizard, the applications are deployed on the target i5/OS partition, and the HTTP server and associated WebSphere Application Server profile are ready to run the IBM-provided applications. At the end of the initial installation they will have been started and ready for your use.

On the workstation on which you ran the wizard to install the Web Enablement Environment, after the installation has completed successfully, you will have a desktop icon that is your entry point to the Web Enablement Environments for i5/OS product and the first local Web

page called FirstSteps. The icon contains a URL that launches a browser to the FirstSteps Web page.

Later you can use the desktop icon or the URL example listed below from a browser on any workstation to access the FirstSteps Web page:

```
http://rchas55:10000/webenv/FirstSteps.html
```

Substitute your partition's host name in place of our *rchas55*.

The FirstSteps Web page has additional documentation and also links to run iSeries Access for Web, the HATS version of flight400, or the Webfaced version of flight400, or to use Web Services to query the flight400 application.

The initial installation performed by 5733-SO1 starts the WebFacing server. If you later end this server or shut down the partition, you must then ensure that the WebFacing server is started before you try the WebFaced version of flight400. You can view the status of the WebFacing server or start it using any of the following interfaces:

- ▶ Using the iSeries Navigator interface, select **systemname** → **Network** → **Servers** → **TCP/IP** → **WebFacing**.
- ▶ Using the iSeries Navigator interface, select **system name** → **Network** → **TCP/IP Configuration**. Right-click and select **Properties** → **Servers to Start**.
- ▶ Using the i5/OS command `STRTCPSVR SERVER(*WEBFACING)`.

After the initial installation of the Web Enablement Environment you may shut down and start the HTTP server and its associated WebSphere Application Server profile created as part of the installation wizard process in any of the following ways:

- ▶ Using the i5/OS partition's IBM Web Administration for i5/OS Web page. You must have the HTTP Administration server (`STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`) that comes with i5/OS up and running. By default this uses port 2001.

One way to do this is using the Start TCP Server command `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`. Once the *ADMIN server is running in subsystem QHTTPSVR perform the following:

- a. From your browser, enter `http://yoursystemname:2001`.
 - b. On the i5/OS Tasks window, select the **IBM Web Administration for i5/OS** link.
 - c. From the next Web page you can access all HTTP servers and Web Application Servers you have defined. Select the **Manage** tab and then either the **All HTTP Servers** or the **All Web Application Servers** tab.
 - d. From the next window you can start or stop the appropriate server. The default values shown are MYHTTPSVR and MYAPPSVR. In the examples shown in this book we used WEHTTPSVR as the name of our HTTP server and WEAPPSVR as the name of our Web Application Server profile and server (same name).
- ▶ Start the HTTP server from either:
 - An i5/OS command:
`STRTCPSVR SERVER(*HTTP) HTTPSVR(WEHTTPSVR)`
 - The i5/OS Web Administrator server browser interface, selecting to start WEHTTPSVR.

- ▶ Start the application server using either:
 - From an i5/OS 5250 workstation enter the Start Qshell command (strqsh or qsh). Once in Qshell mode enter:
 - i. `cd /qibm/proddata/websphere/appserver/v6/base/bin`
 - This gets the current path under the i5/OS Integrated File Systems to the one containing WebSphere Application Server V6.0 software
 - ii. `startServer -profileName WEAPPSVR`
 - You can also stop the application server later with `stopServer -profileName WEAPPSVR`.
 - From the i5/OS Web Administrator server browser interface, select the **All application servers** tab and start WEAPPSVR.
- ▶ Use the desktop icon or the URL example we previously listed (<http://rchas55:10000/webenv/FirstSteps.html>) to access the FirstSteps Web page.

Figure 14-3 shows an excerpt from the FirstSteps Web page on your system.

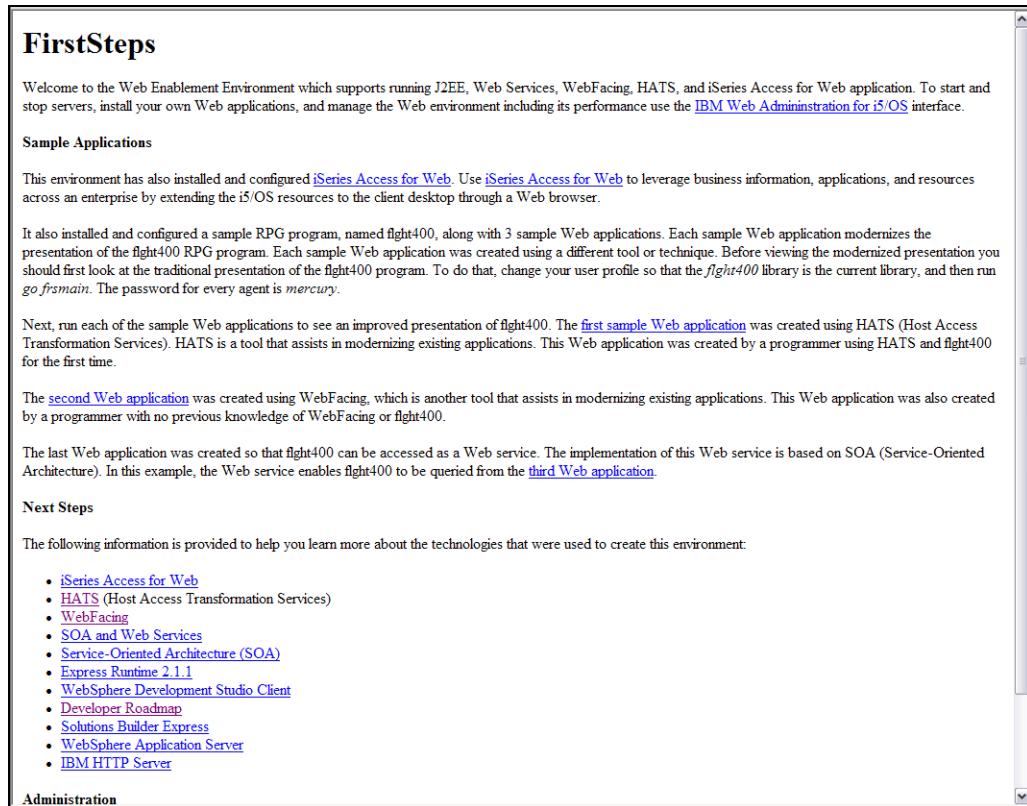


Figure 14-3 Web Enablement Environment FirstSteps window

To run flight400 you must place library flight400 in the active job's library list as the current library. You can do this in any of the following ways, depending upon what works best for your i5/OS run time environment:

- ▶ Place library flight400 in the system value QUSRLIBL (User Library List).
- ▶ Place library flight400 in the system value QSYSLIBL (System Library List). This has the disadvantage that if you want to remove any library name for QSYSLIBL, a partition start or system IPL (no partitions) is required.

- ▶ Create a job description unique to the i5/OS user profile you will sign on to run that contains flight400 in the job description's initial library list (INLLIBL) parameter.
- ▶ Based upon the i5/OS user ID you sign on as, call a CL program that includes statements similar to the following (Example 14-1). This CL program is not called when initiating the WebFaced version of flight400. You must use one of the other techniques to set the current library to flight400 to run the WebFaced flight400 application.

Example 14-1 CL excerpt to add library flight400 to the user's library list

```

:                                     :
ADLE:      :                               :
          : ADDLIBLE   LIB(FLGHT400) /* Used for WEE - +
          :             Hats/WebFace/Webservice. Ignore error if +
          :             not found */
          : MONMSG     MSGID(CPF0000) /* If FLGHT400 not installed +
          : CHGCURLIB  CURLIB(FLGHT400)
          : MONMSG     MSGID(CPF0000) /* If FLGHT400 P not installed +
          :             keep trucking */
          : CALL      PGM(QCMD) /* Back to whatever you were going +
          :             to do. */
:                                     :
:                                     :
:                                     :
:                                     :

```

When you performed the installation wizard steps you selected ports that would be used to run your WebFaced version of flight400. Use the URL:

http://<your host name>:<http server port number>/flight400/

In our example this was:

http://<rchas55:10000/flight400/

Figure 14-4 shows an example of the flight400 5250 application running in WebFaced mode, from a browser.

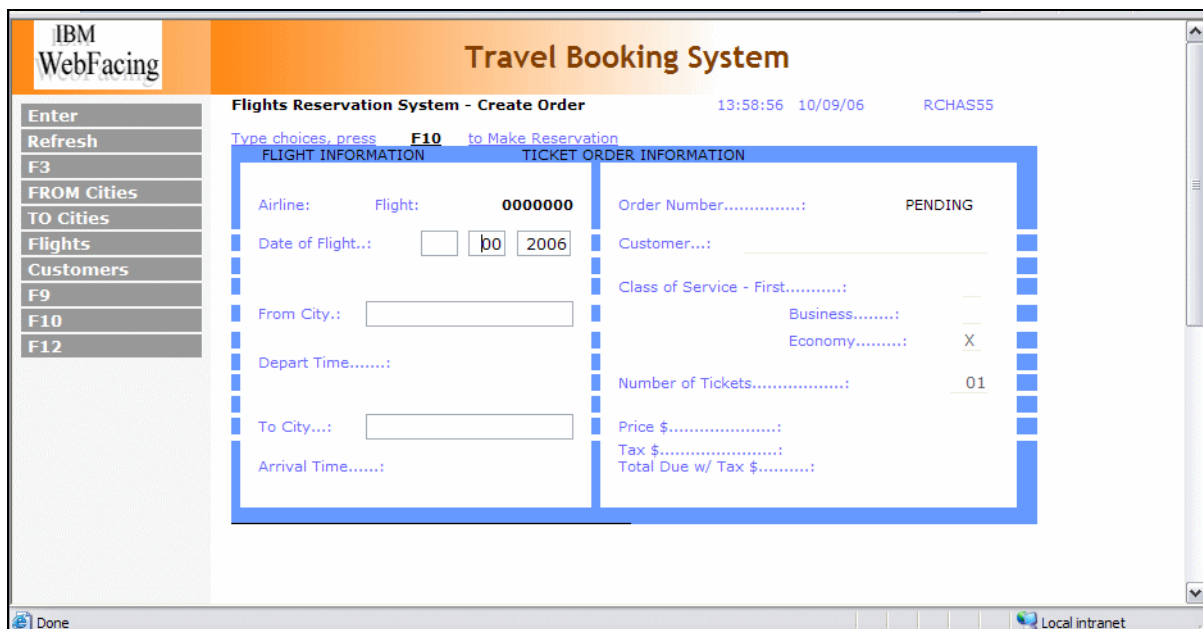


Figure 14-4 Example flight400 application screen using the WebFaced interface

Using the URL shown for the WebFaced example as a base, the URL for the HATS version of flight400 is:

http://<your host name>:<http server port number>/flight400-hats/

In our example this was:

http://rchas55:10000/flight400-hats/

Figure 14-5 shows an example of running the flight400 application using HATS. The light brown background with brown text area within the 5250 screen area makes it hard to see the application text in a black and white PDF print of this book.

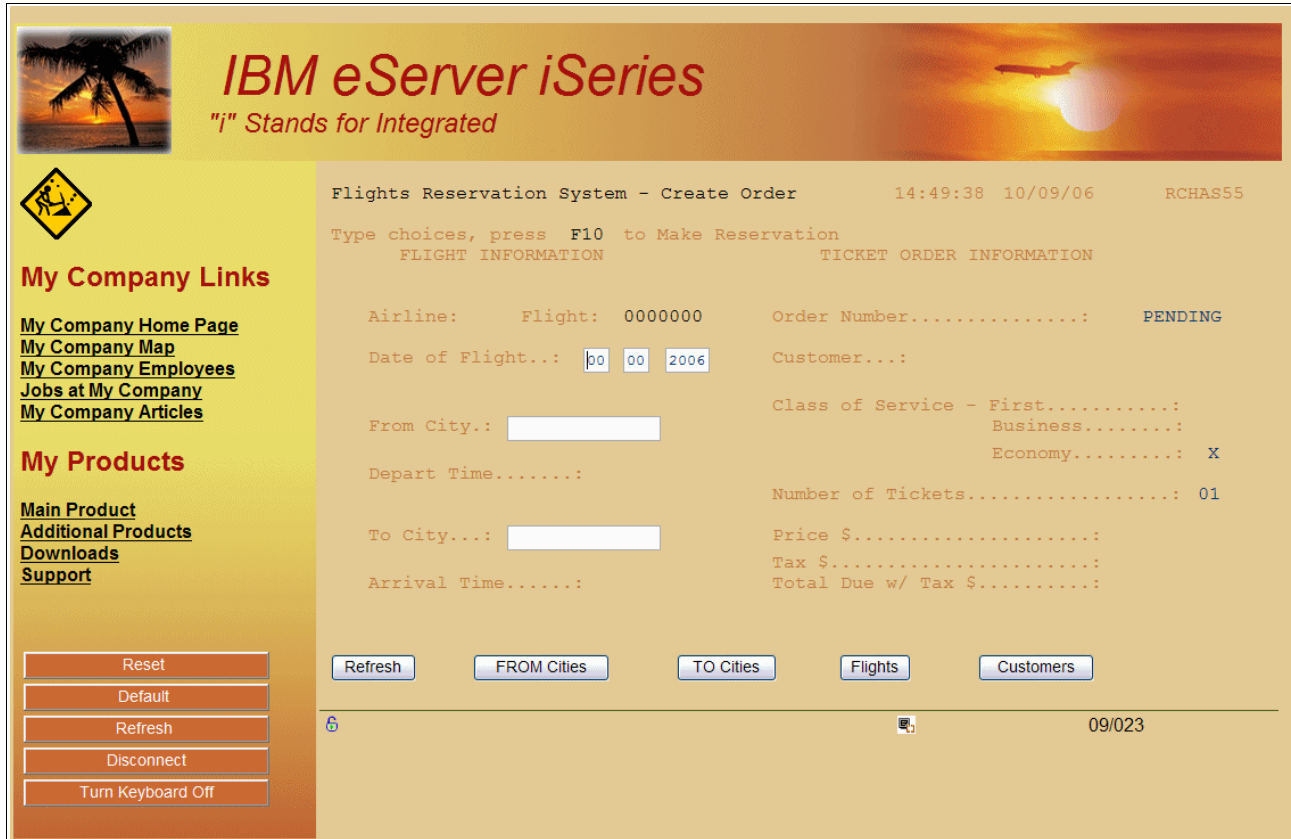


Figure 14-5 Example flight400 application screen using the HATS interface

IBM has customized the flight400 application using a subset of the HATS and WebFacing rendering capabilities available using the WebSphere Development Studio for iSeries Client application development product.

Figure 14-6 shows the first Web page for iSeries Access for Web. The URL for the example page shown for iSeries Access for Web is:

http://<your host name>:<http server port number>/webaccess/IWAMain

In our example this was:

http://rchas55:10000/webaccess/IWAMain

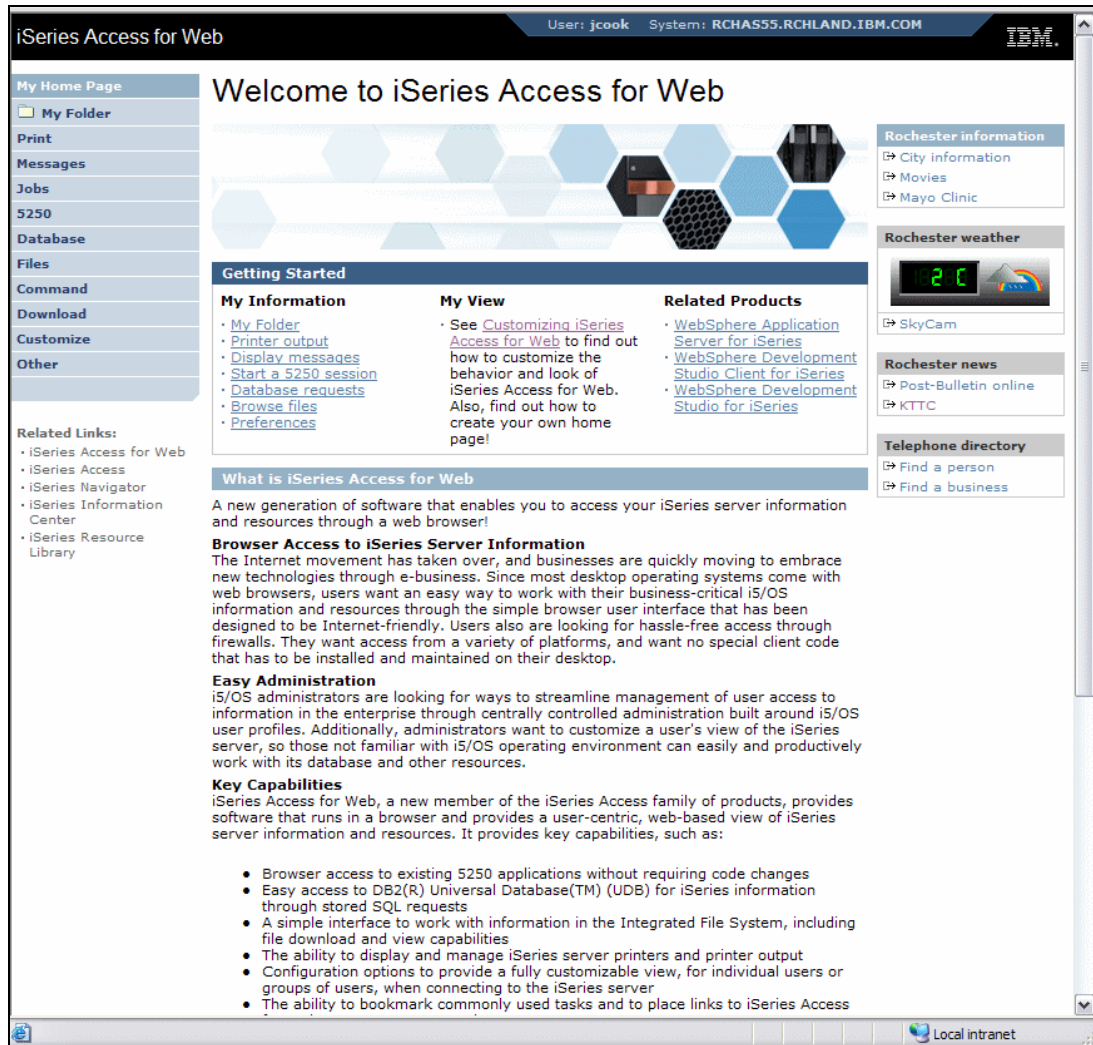


Figure 14-6 iSeries Access for Web home page

iSeries Access for Web provides a browser interface to i5/OS supporting the 5250 command interface, as well as links to functions that include managing your output queues, printers, messages, viewing active jobs, running SQL statements, and more.

Installing the IBM HTTP, WebSphere Application Server, and middleware along with applications gets you up and running with a Websphere Application Server V6.0 application serving environment with minimum of time and utilization of your company's human resources. You can quickly see ways you might:

- ▶ Further integrate links among iSeries Access for Web and WebFaced or HATS-ed applications.
- ▶ Consider purchasing the WebSphere Development Studio for iSeries Client application to do further 5250 screen customization as well as develop new Web applications.

Note that as you run through all of the flight400 application options windows, you will notice different ways the HATS or WebFaced customization provides functions compared to the 5250 workstation version. Some 5250 functions, such as use of function key and menu options as links or pull-down lists were chosen differently between the HATS and WebFaced versions of flight400. This way you evaluate the look and feel of the flight400 application using a 5250 workstation compared to the Internet HATS version (and also a WebFaced version installed with Web Enablement Environment) without having to first install the WebSphere Development Studio for iSeries Client application development product on your own PC workstation

The customized options chosen to demonstrate alternatives may not to be the best options for you. This way you can more easily consider parts of your application that you would deliver to the Web as basic refaced interfaces or take advantage of more advanced interfaces such as adding graphics to the window or even adding a new function that uniquely uses more advanced interfaces and functions available to Web users. This more advanced customization or addition of new applications could be possible using i5/OS application development tools, such as the WebSphere Development Studio for iSeries Client product.

14.2.4 Using IBM Express Runtime Web Environments for i5/OS with applications

IBM business partners or Independent Software Vendors may want to consider integrating their applications into IBM Express Runtime Web Environments for i5/OS, just as flight400 and iSeries Access for Web have been when using IBM Express Runtime Web Environments for i5/OS.

For more information about Express Runtime Web Environments for i5/OS and this specific topic refer to the IBM Redbook *WebSphere Application Server for i5/OS Handbook: Version 6.1*, SG24-7221.

Note: IBM Express Runtime Web Environments for i5/OS, 5733-SO1, should not be confused with another IBM ease-of-use application solution deployment produced from IBM - Express Runtime V2.1.1

57333-SO1 was created using this Express Runtime product. For information about Express Runtime V2.1.1 see 13.3, "IBM Express Runtime" on page 334. There are IBM Redbooks on Express Runtime V2.1.1 and Express Runtime V2.1.

14.3 Web performance management tools

Two new Web performance management tools became available during 4Q 2006:

- ▶ Web Performance Monitor (November 2006 availability via PTF Group SF29114 Level 6 and PTF SI25855)
- ▶ Web Performance Advisor (November 2006 availability via PTF Group SF29114 Level 6) and PTF SI25855)

Important: Whenever starting to use new WebSphere Application Server-based functions, we recommend ensuring that you have the latest available sets of PTFs (fixes) in the following component areas:

- ▶ The latest cumulative PTFs for your i5/OS version
 - ▶ The following PTF groups for the corresponding i5/OS release:
 - The latest Java PTF groups.
 - The latest DB2 Universal Database for iSeries PTF groups.
 - The latest HTTP PTF groups.
 - The latest WebSphere PTF groups for your WebSphere Application Server version and product these two Web performance tools support. As of December 2006, these are:
 - WebSphere Application Server 5.1.1.x
 - WebSphere Application Server V6.0.x
- WebSphere Application Server 6.1.x, any Portal Server, and any Workplace releases are not supported.

Both of these tools are accessed from the i5/OS-provided default HTTP Administration server, commonly referred to as the IBM Web Administration for i5/OS. This product has been available for many i5/OS releases and continues to get enhanced with each new i5/OS release and frequently within a release cycle.

This Web Administration for i5/OS server can be started using either of the following methods:

- ▶ By the Start TCP/IP Server command `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`.
- ▶ By the iSeries Navigator interface (**systemname** → **Network** → **Servers** → **HTTP Administration**).
- ▶ By specifying *YES for the Autostart (AUTOSTART) parameter of the Change HTTP Attributes (CHGHTTPA) command. This tells i5/OS to automatically start the HTTP server when TCP/IP is started by the Start CTP/IP (STRTCP) command.

Though you can change the default for port number, you typically access the i5/OS Administration server by entering the following from your browser:

`http://hostname:2001`

Then:

1. On the i5/OS Tasks page, select the **IBM Web Administration for i5/OS** link.
2. Select the **Manage** tab.
3. Select the **Application Servers** tab.
4. From the next window showing all defined servers, select a specific application server you want to use with the Web Performance Monitor or the Web Performance Advisor.
5. For that server, click the **Manage Details** button to get the window showing the monitor or advisor links.
6. Click either the Web Performance Monitor or the Web Performance Advisor link in the lower left navigation area.

In our example shown in Figure 14-7, we have selected **WEAPPSVR**. You see the links to the Web Performance Advisor and Web Performance Monitor.

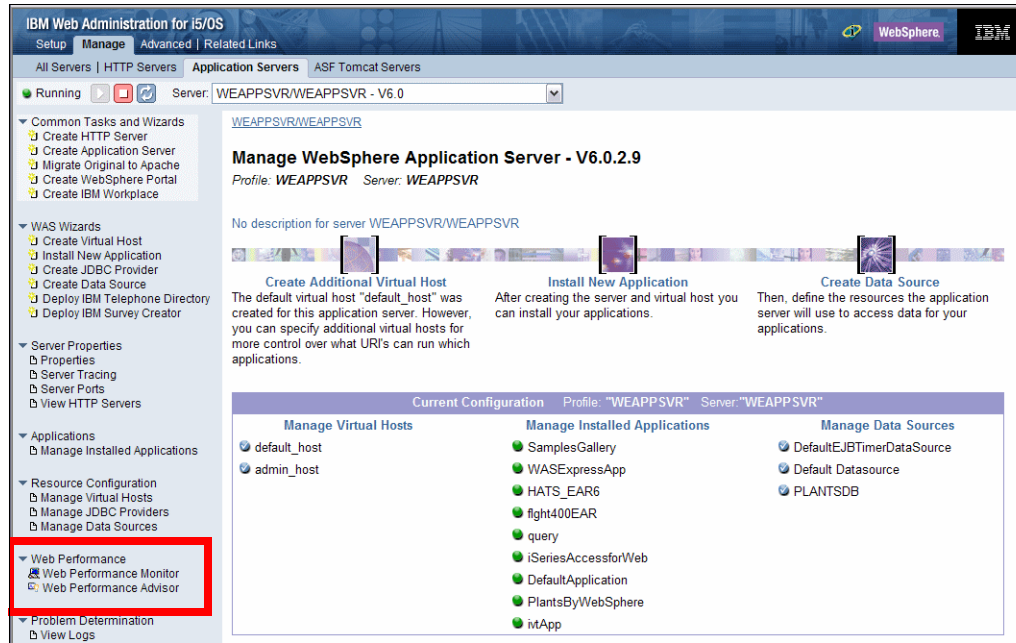


Figure 14-7 Accessing the Web Performance Monitor and Performance Advisor

Click either link shown.

Important: Through January 2007:

- ▶ The Web Performance Advisor supports WebSphere Application Server 5.0.x, 5.1.x, and 6.0.x, Portal and Workplace.
- ▶ The Web Performance Monitor supports WebSphere Application Server 5.1.1.x and later servers.

If either of these tools does not support the version of the server you chose, the link will not be displayed.

The following sections provide additional overview information about these two performance management tools.

14.3.1 Web Performance Monitor

The Web Performance Monitor collects unique Web application performance and transaction data for just the jobs that comprise the selected Web environment. This includes the HTTP server, application server, and associated database jobs used by WebSphere applications running in the selected application server.

The data collection is performed by gathering performance data generated by Application Response Measurement 4.0 (ARM) instrumentation already shipped with:

- ▶ HTTP Web servers from IBM
- ▶ WebSphere Application Server, versions starting with V5.1.1.x through V 6.0.x
- ▶ DB2 Universal Data Base for iSeries (i5/OS)

ARM is a standard of The Open Group (<http://www.opengroup.org>), composed of a set of APIs that can be used to collect response time information for enterprise applications. The ARMs were originally instrumented for use with WebSphere Enterprise Workload Manager (EWLM).

EWLM monitors and records activities across multiple system (partition) and operating system environments, whereas the Web Performance Monitor monitors a single i5/OS partition.

Overhead is introduced to your WebSphere transactions while the monitor is running, so you typically would not run the monitor during multiple user production mode. You use it to understand the performance characteristics of your WebSphere application and associated HTTP server and application server. You are looking for long response times or high transaction counts for certain transactions or functions relative to others performed by the actual application.

For test purposes, with only a few WebSphere application users active, the overhead is barely noticeable.

This new monitor information can help identify the number of Web-based transactions and average transaction response times over the time period of collection, and the relative amount of time spent within sub portions of a transaction, such as within the HTTP server, within the Web applications, and within the database functions of transactions if the applications are using WebSphere Application Server data source interfaces.

Understanding the components of transaction processing becomes especially important when applications consist of many jobs and threads running within a single system and on multiple systems, such as when the database functions and application processing functions are performed on different systems.

The monitor can:

- ▶ Monitor end-to-end performance.
- ▶ Provide a WRKACTJOB-like level of information on the initial window (that can be refreshed).
- ▶ Provide job-level and transaction-level statistics among the selected HTTP server, WebSphere Application Server applications, and, optionally, any data source access by the applications to DB2 for i5/OS database server jobs.
- ▶ Provide more detailed job management functions by linking to optionally started iSeries Navigator tasks on the Web.

Perform the following to run the Web Performance Monitor:

1. Access the Web Administration for i5/OS from your browser and select a specific applications server. For that server, select the **Manage Details** button. Click the **Web Performance Monitor** link at the left of the details page. The browser will open the Web Performance Monitor page, as shown in Figure 14-8.

In our example we use the HTTP server WEHTTSPVR and WebSphere Application Server profile/server WEAPPSVR we had the wizard of the IBM Express Runtime Web Environments for i5/OS create for us in 14.2.2, “Installing the Web Enablement Environment” on page 343.

2. Click the **Start** button and processing begins. An hourglass is shown in front of the current item and items are checked off as they are completed.

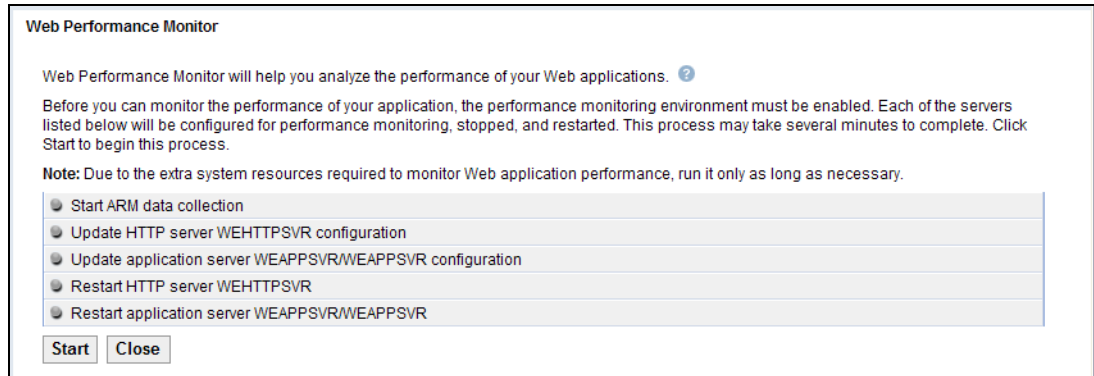


Figure 14-8 Web Performance Monitor - start example

When the start process has completed, your HTTP server and WebSphere Application Server instance or profile are enabled for performance monitoring.

The Web Performance Monitor shows the job and transaction statistics for your Web-based jobs as they run on the i5/OS partition. Notice that it shows a WRKACTJOB-like screen, as shown in Figure 14-9.

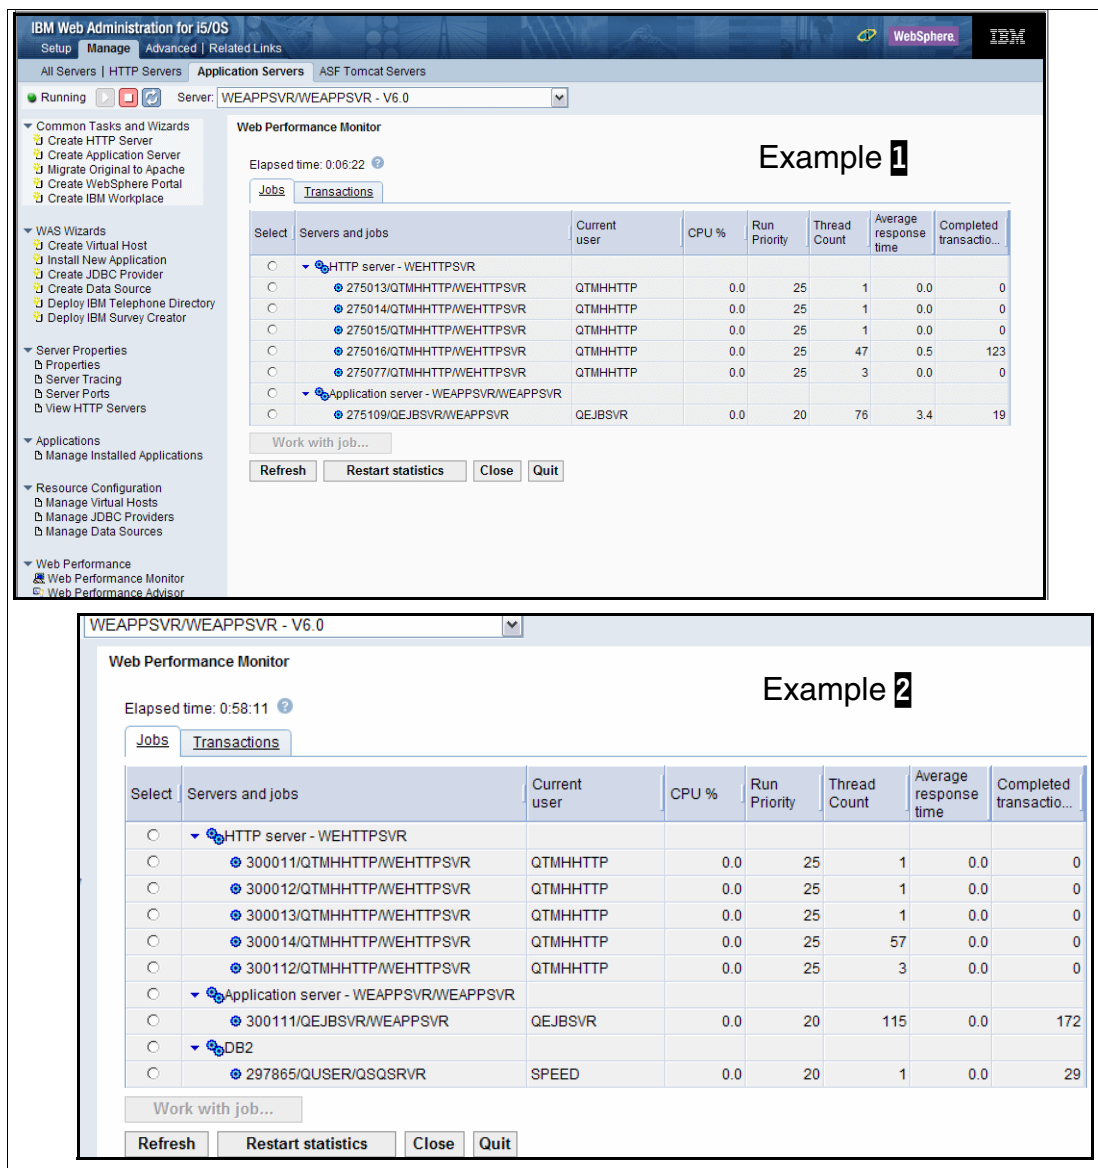


Figure 14-9 Web Performance Monitor - jobs example

If you select a job, the **Work with job** button is activated. If you then select that button, you are launched to the iSeries Navigator tasks on the Web application (included in i5/OS 5722-SS1, Option 3), provided that:

- ▶ On V5R4 you have started the WebSphere Application Server server or profile SYSINST/ADMIN.
- ▶ On V5R4 you have PTF SI24255 installed. With this PTF installed, iSeries Navigator tasks on the Web runs in the integrated Web application server.

iSeries Navigator tasks on the Web provide a large subset of the iSeries Access for Windows iSeries Navigator functions from a browser interface rather than from a Windows workstation iSeries Navigator interface. Displaying active job information is one of its functions.

If this application is not activated, you will get an error window. To learn more about iSeries Navigator tasks on the Web in this book, refer to:

- ▶ 14.1, “iSeries Access for Web, iSeries Navigator tasks on Web, and new integrated Web application server” on page 338
- ▶ 9.1, “iSeries Navigator tasks on the Web” on page 190.

You can also show a transactions view of the performance information by clicking the **Transactions** tab. We do not show an example of the Transactions window in this book, but in that view you can specify preferences and user filtering to control what is shown on the screens. For example, show only the activity for user USERxxx.

In our jobs example windows shown in Figure 14-9 on page 355, you see very little CPU utilization, but some work has been done as evidenced by:

- ▶ HTTP server job 275016 has processed 123 requests (completed transactions).
- ▶ The associated WebSphere Application Server job - 275109 has processed 19 completed transactions with an average response time of 3.4 seconds.

Further discussion of the reported metrics and how to best use the Web Performance Monitor is beyond the scope of this book. However, it is apparent that you need to be able to work with someone who understands the application implementation in order to productively use the statistics that are presented to you. You also must be aware of the following.

- ▶ One browser request (mouse click, function, or Enter key) sent to the system results in one request being sent through the HTTP server. However, often the response to the request contains links to files (such as images) to be part of the end user's view of the response.
- ▶ This results in the browser sending additional requests to the HTTP server that may be handled within the HTTP server or routed to a processing program. These additional request and response pairs are counted as a separate transaction even if they do not get processed by the application program, such as the flight400 application.
- ▶ One effect of this, as shown in our example 1 window, is that the HTTP transaction counts are higher than the WebSphere Application Server counts, and HTTP server average response times are lower than the WebSphere Application Server average response time.
- ▶ An application, such as the trade application we installed in WEAPPSVR, can be run with a port number such that the browser request goes directly to the application server. That is, the request/response is never processed by the HTTP server (WEHTTTPSVR, in our example). Therefore, there is no associated increase in transaction counts for the HTTP server, while there could be an increase in the application server statistic values. We show an example of this in our example 2 window in Figure 14-9 on page 355. This example shows an application server transaction count of 172, but zero transactions for our HTTP server.
- ▶ In our two example windows you see that example 1 does not show any associated i5/OS DB2 database server job metrics, while example 2 does show some DB2 database server statistics. You may see an associated database server job (QSQSRVR jobname prefix) or QZDASOINIT job name prefix), depending on how your application interfaces to DB2 for i5/OS.

Under i5/OS an application can access the i5/OS database with any of the following hierarchical level interfaces, starting with the application layer at the lowest level:

- Application layer: On i5/OS this can implicitly or explicitly be:
 - Native file level, non-SQL interface (open file, read, write, update, delete, and so on).

- SQL, application level interface (SELECT, and so forth), no interface through a JDBC Provider.
 - JDBC Provider - Data Source, using the i5/OS-unique native JDBC driver. This driver is unique to i5/OS and is often referred to as *DB2 UDB for iSeries (Native XA)*. This interface performs all SQL functions under one of the associated i5/OS database server jobs with the QSQSRVR job name prefix.
 - JDBC Provider - Data Source using the industry standard for database access driver. On i5/OS this is known as the Java toolbox driver interface and performs all SQL functions under one of the associated i5/OS database server jobs with the QZDASOINIT job name prefix.
- WebSphere Application Server Server layer
 - JDBC Provider - Data Source using the i5/OS-unique native JDBC driver (under the QSQSRVR job name prefix).
 - JDBC Provider - Data Source using the industry standard for database access driver (under the QSQSRVR job name prefix).
 - WebSphere Application Server Node layer
 - JDBC Provider - Data Source using the i5/OS-unique native JDBC driver (under the QSQSRVR job name prefix).
 - JDBC Provider - Data Source using the industry standard for database access driver (under the QSQSRVR job name prefix).
 - WebSphere Application Server Cell layer
 - JDBC Provider - Data Source using the i5/OS-unique native JDBC driver (under the QSQSRVR job name prefix).
 - JDBC Provider - Data Source using the industry standard for database access driver (under the QSQSRVR job name prefix).

In example 2 we included transactions performed by the Trade (stock transaction) application used as a WebSphere Application Server benchmark within IBM. This application uses the *DB2 UDB for iSeries (Native XA)* JDBC Provider - Data Source, defined at the node level.

Use the WebSphere Administrative Console interface to your WebSphere Application Server server - profile → **JDBC Providers** link to view what is defined.

- ▶ For now it seems clear that if browser users have a response time concern and you see high average response time for the WebSphere Application Server application, the primary contributor to the high response time is at the application level. The performance problem could be in the area of database access, but you need to use other i5/OS-based performance tools to analyze this area.
- ▶ Once you have a look at the monitor statistics you may next consider looking at current i5/OS system values, WebSphere Application Server values, and the efficiency of the application itself. This leads us to the Web Performance Advisor topic after we end the currently active Web Performance Monitor in our example scenario.

To end the Web Performance Monitor, click **Quit** on the Jobs (shown in Figure 14-9 on page 355) or the Transactions window. This brings up the window shown in Figure 14-10. Click **Stop**.

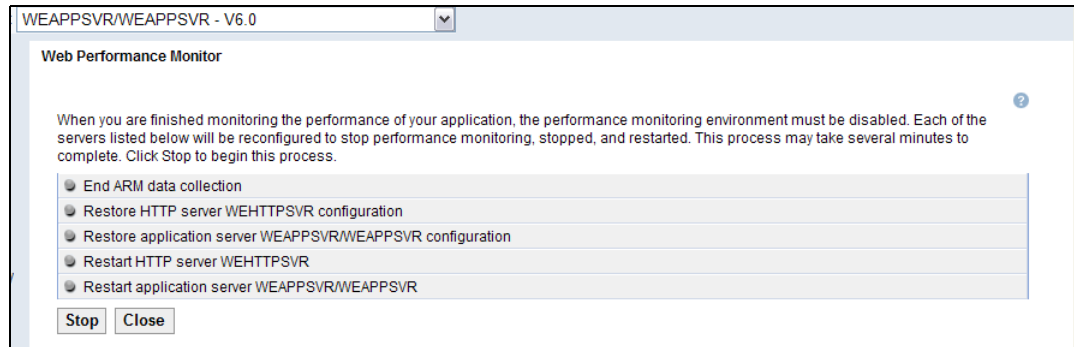


Figure 14-10 Ending the Web Performance Monitor

You can watch progress check marks as each step is completed. When complete, your HTTP server and associated application server are back up and running without the monitoring capability enabled.

14.3.2 Web Performance Advisor

The Web Performance Advisor was made available during November 2006 with PTF SF99114. This advisor is unique to i5/OS, providing a set of wizards and attribute management tools that can be used to evaluate and improve the performance of a Web environment.

One of the primary advantages of using this Advisor tool, which you can invoke in real time during normal workload operations, is that it interrogates the named WebSphere Application Server instance and extracts several parameter values that affect performance and presents them all in a single, easily accessible interface. This minimizes the typical number of mouse clicks needed to find this information using the WebSphere Application Server profile's administrative console interface.

The primary advisor capabilities include:

Does a quick analysis of your Web performance-related i5/OS system values, HTTP server configuration parameters, and WebSphere Application Server profile configuration parameters. Based on the evaluation, it indicates, within an area, if some changes are suggested to improve performance. If there are recommended values to change to, they are presented to you. The analysis and recommended values are based upon experience by Rochester IBM System i performance experts.

- ▶ Managing the performance attributes analyzed: This allows you to view and modify all the attributes that can have a significant impact on the performance of a Web environment, like:
 - View and modify the value.
 - View the current value's rating. (Is this value acceptable?)
 - View and set the recommended value.
- ▶ Import/export the performance attributes: Saves the performance attributes to an xml file that can be sent to IBM or an ISV for additional performance analysis. The IBM or the ISV representative can update the XML file and return it to you. Assuming that you agree on

the updated attribute values, these updates can then be imported into your Web environment.

To start the Web Performance Advisor you have to use the IBM Web Administration for i5/OS management of application servers interface to select the specific WebSphere Application Server profile - server. Select **Manage Details** and then click the **Web Performance Advisor** link (shown in Figure 14-7 on page 352).

You then get the initial advisor window shown in Figure 14-11.

You can see some system attributes and buttons to link to further information and guided assistance:

- ▶ Manage system attributes.
- ▶ Manage application server and HTTP server attributes.
- ▶ Export or import the performance profile (xml file).

You can also see an indication of the initial evaluation performed, as indicated by the text and exclamation point icons in our example, as shown in Figure 14-11. Our example shows WebSphere Application Server WEAPPSVR and HTTP server WEHTTSPVR.

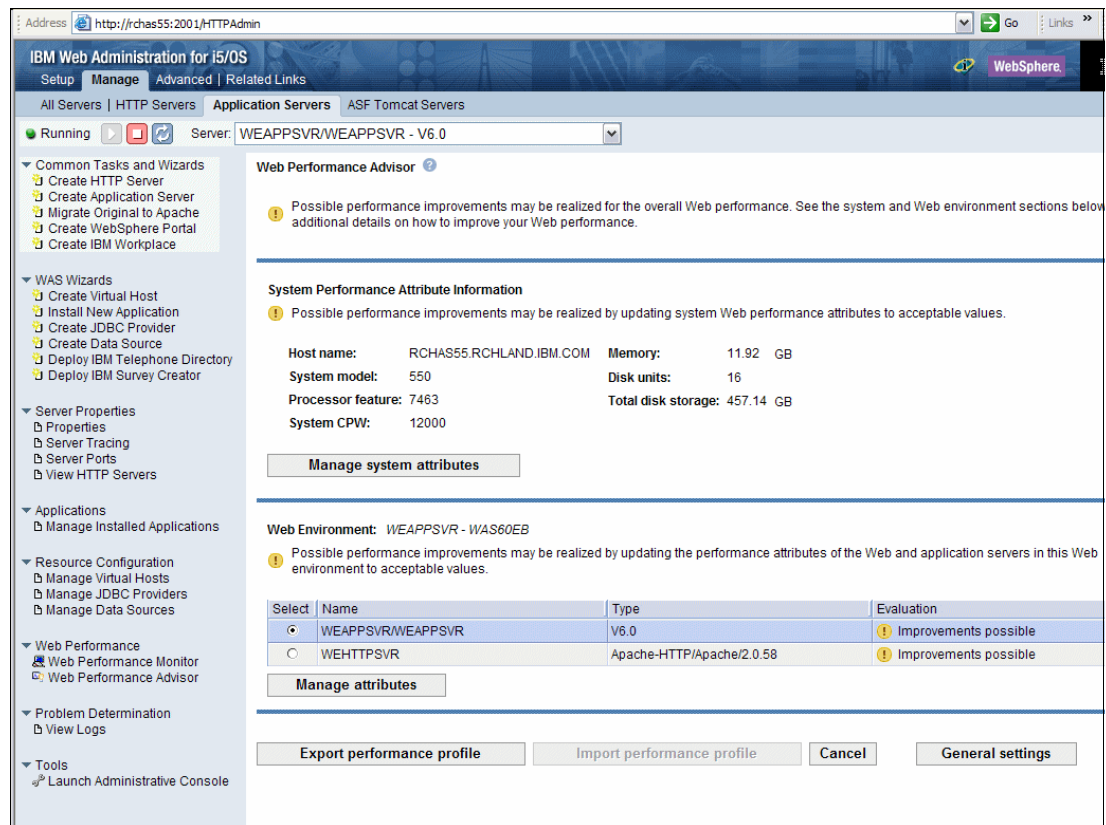


Figure 14-11 Web Performance Advisor - initial window

The exclamation point icons indicate that the Advisor considers that there are some parameters that, if changed, could improve performance.

You can view all of the servers that comprise the current environment, review their performance parameter settings, and select and manage individual attributes that affect server performance.

This first page gives a good overview over the entire machine, as it shows the host name, system model, processor feature, system CPW, memory, number of disk units, and total amount of disk.

Rchas55 is a partition. The host name, memory size, disk units, and total disk storage are for the Rchas55 partition. Other values are system wide. Processor feature 7463 indicates that this is an Enterprise Edition (full 5250 CPW capacity).

We looked up feature 7463 in any of the following documents to determine that 7463 is an Enterprise edition:

- ▶ *System i Performance Capabilities Reference i5/OS Version 5, Release 4*
- ▶ *IBM System i5 Handbook: IBM i5/OS Version 5 Release 4, SG24-7486-00*
- ▶ *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006, SG24-2155*

Note that using the Display System Value (DSPSYSVAL) for system value QPRCFEAT would also show 7463 on our example system (partition).

Figure 14-12 shows an example of System Resources tab information with *advise links* you can review and consider making changes to its associated parameter. You also see a tab for performance settings, and also the PTF Groups tab and the Web PTFs tab.

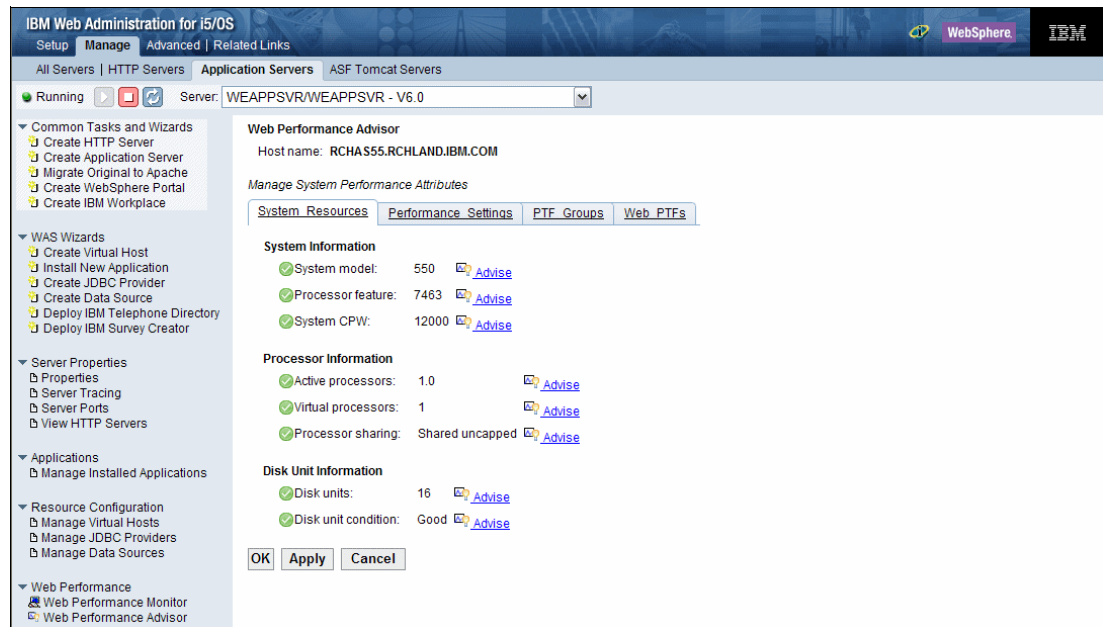


Figure 14-12 Web Performance Advisor - manage system attributes example

The PTF tabs provide a summary level list of PTFs, as illustrated in Figure 14-13.

The screenshot displays the Web Performance Advisor interface for a system named WEAPPSVR/WEAPPSVR - V6.0. The interface is divided into two main sections: 'PTF Groups' and 'Web Product PTFs'.

PTF Groups Section:

Host name: RCHAS55.RCHLAND.IBM.COM

Manage System Performance Attributes

System Resources | Performance Settings | **PTF Groups** | Web PTFs

The level and status of PTF Groups on this system

PTF Group	Level	Status	Description
SF99540	6192	Installed	CUMULATIVE PTF PACKAGE C6192540
SF99539	27	Installed	GROUP HIPER
SF99504	4	Installed	DB2 UDB FOR ISERIES
SF99322	3	Unknown	WEBSPPHERE APP SERVER V6.1
SF99315	3	Installed	TCP/IP GROUP PTF
SF99312	7	Installed	WEBSPPHERE APP SERVER V6.0
SF99303	3	Not applicable	WEBSPPHERE MQ FOR ISERIES - V6.0
SF99301	17	Unknown	WEBSPPHERE APP SERVER V6.0
SF99300			
SF99298			
SF99295			
SF99293			
SF99291			
SF99290			
SF99289			
SF99288			
SF99287			
SF99286			
SF99285			
SF99282			

Web Product PTFs Section:

Host name: RCHAS55.RCHLAND.IBM.COM

Manage System Performance Attributes

System Resources | Performance Settings | PTF Groups | **Web PTFs**

PTFs for Web related products on this system

PTF Information	Status
▶ 5722DG1 -- IBM HTTP Server for i5/OS	
▶ 5722DG1 -- Triggered Cache Manager	
▶ 5722JV1 -- Java Developer Kit 1.3	
▶ 5722JV1 -- Java Developer Kit 1.4	
▶ 5722JV1 -- Java Developer Kit 5.0	
▶ 5722JV1 -- J2SE 5.0 32 bit	
▶ 5733W60 -- WebSphere Application Server for OS/400 V6	
▶ 5733W60 -- WebSphere Application Server V6 ("Base")	
▶ 5733W60 -- WebSphere Application Server V6 Network Deploy	
▶ 5722999 -- Licensed Internal Code	
▶ 5722SS1 -- i5/OS	
▶ 5722SS1 -- Extended Base Support	
▶ 5722SS1 -- Online Information	
▶ 5722SS1 -- Extended Base Directory Support	
▶ 5722SS1 -- Example Tools Library	
▶ 5722SS1 -- Host Servers	
▶ 5722SS1 -- System Openness Includes	
▶ 5722SS1 -- ObjectConnect	
▶ 5722SS1 -- Integrated Server Support	
▶ 5722SS1 -- Qshell	
▶ 5722SS1 -- Domain Name System	
▶ 5722SS1 -- Portable App Solutions Environment	
▶ 5722SS1 -- Digital Certificate Manager	

Figure 14-13 Web Performance Advisor: PTF groups and Web product PTFs

You can click the arrow to the left of each Web product and see the list and status of that product's PTFs installed on the partition.

Click the **System Resources** tab. On that Web page click the **Advise** link next to a parameter and the Advisor presents an Advise window. Figure 14-14 shows examples of the Advise windows for several of the partition's system resource parameters.



Figure 14-14 Example of system resources advice

You see a recommended value in our examples but no capability to change a setting. The settings shown in our example need to be changed outside of the Web Performance Advisor. For example, the number of virtual processors on our 550 model could be changed using the Hardware Management Console (HMC) interface.

In the figures that follow we show examples of the Performance Settings tab Advise windows and the capability to make changes. Depending upon the actual parameter and the software affected by that value of that parameter, you may be able to make a change in real time or not be able to change the value immediately. The Advise window explains this.

Figure 14-15 shows an example of the output you get after selecting the Performance Settings tab. Note that the Performance adjustment setting and the TCP/IP Buffer Sizes are highlighted in this case, as they are out of line with recommendations. You can click the **Advise** link to get detailed information about the recommended settings.

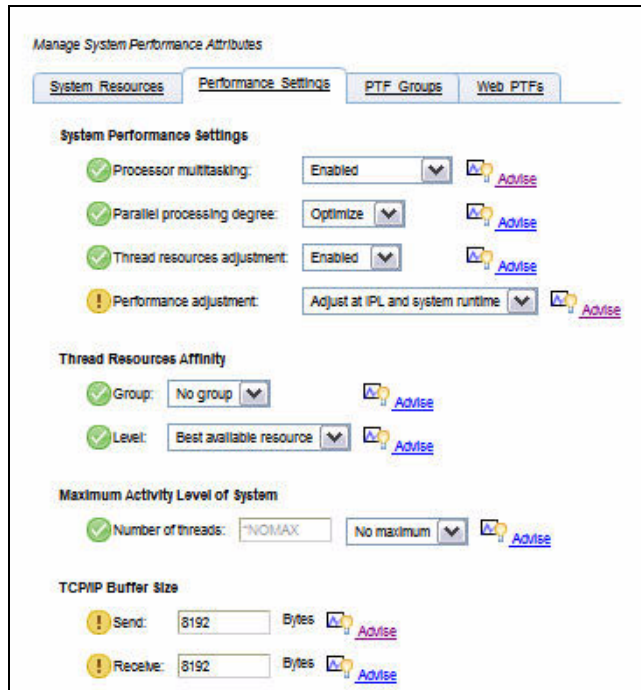


Figure 14-15 Web advisor - Performance Settings list

We selected to get the advise information about the performance adjustment parameter, the TCP/IP send buffer size, and the parallel processing degree (for queries) values, as shown in Figure 14-16.

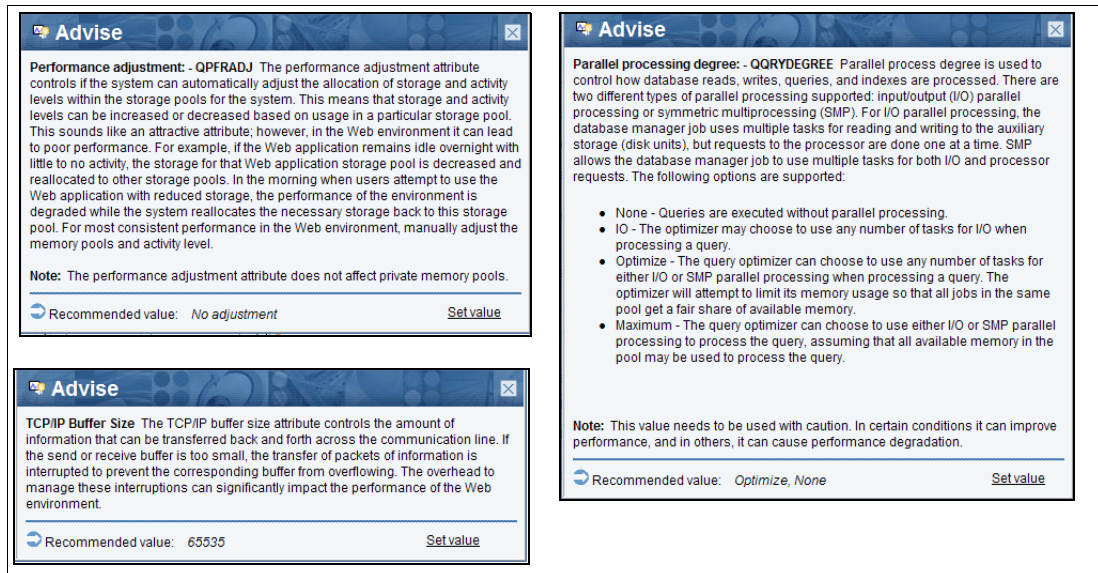


Figure 14-16 Performance Settings Advise examples: TCP/IP buffer, query parallel processing

Review the Advise window text. Note that each window has a set value link. Each time you click the **Set value** link you are returned to the main Performance Setting window with the recommended value already entered.

You would finish all your Advise window processing and then select **Apply** to have the recommended values actually set.

Note the advise text that recommends manual adjustment of pool sizes and activity levels for Web application environments.

We selected **Set value** for both the Performance Adjustment (i5/OS system value QPFRADJ) recommended value and the recommended 65535 value for the TCP/IP send buffer — an i5/OS Configure TCP/IP Attributes parameter.

Figure 14-17 shows our two changed values.

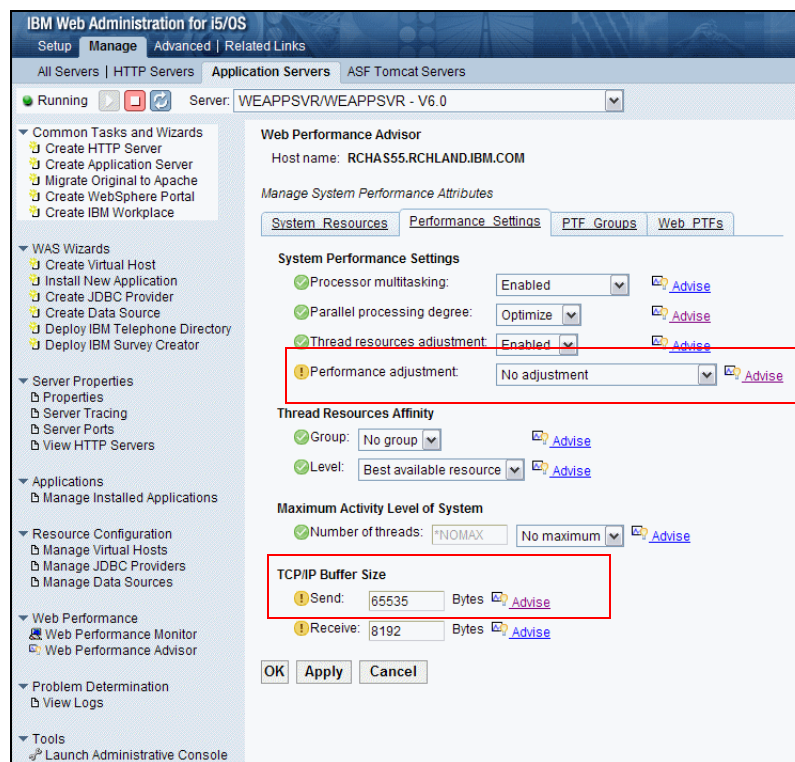


Figure 14-17 Changed Performance Settings example

You must click **Apply** to have the changes put into effect.

The Web Environment section of the Web Performance Advisor (Figure 14-11 on page 359) allows you to select either your Web application server or HTTP server.

This option is another place where Web Performance Advisor demonstrates significant advantages over the administrative console interfaces of the Web application servers you are using — for finding individual performance-related parameters and assisting you in evaluating and optionally changing a value for each of the parameters.

Instead of you wandering around several levels of windows to find the important performance parameters, the advisor retrieves them and places them on advisor windows in one place where you can quickly view, review the advise information, and, depending on the parameter, change them. Advise information is provided for each parameter presented to you.

Using the window shown in Figure 14-11 on page 359, we first select **WEHTTPSVR** and then clicked **Manage attributes** to get the window shown in Figure 14-18.

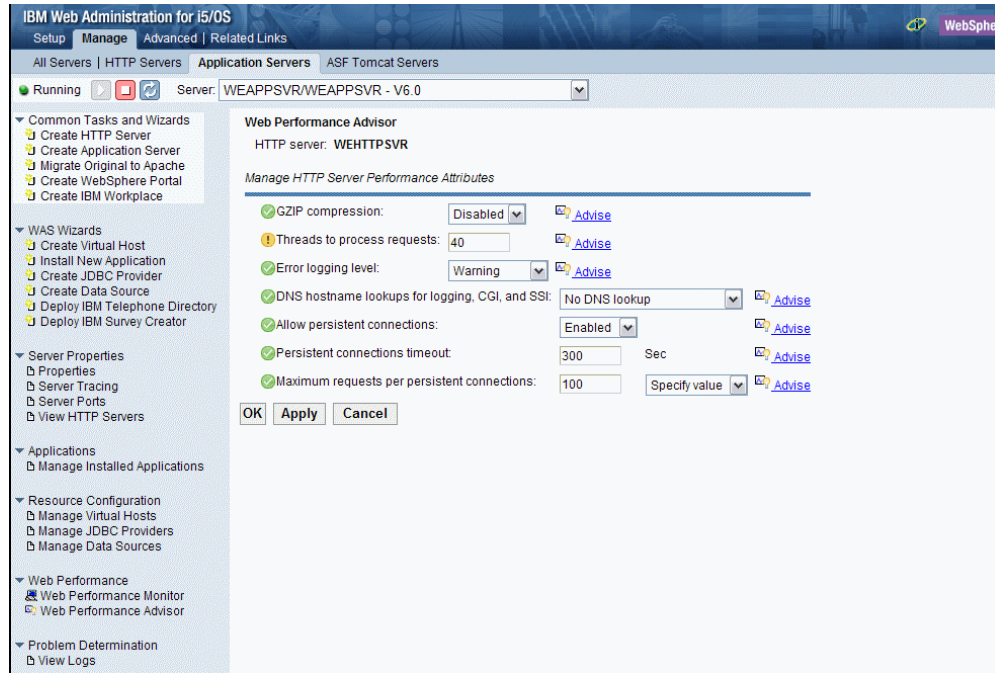


Figure 14-18 Web advisor - HTTP Server Web environment settings

Reviewing the Advise link for each parameter can provide useful information for those new to HTTP serving performance considerations and also provide helpful reminders to those who are familiar with, but not experts on HTTP performance considerations.

In our example we only show the Advise window (in Figure 14-19) for the Threads to process requests parameter highlighted by the exclamation point icon (!).

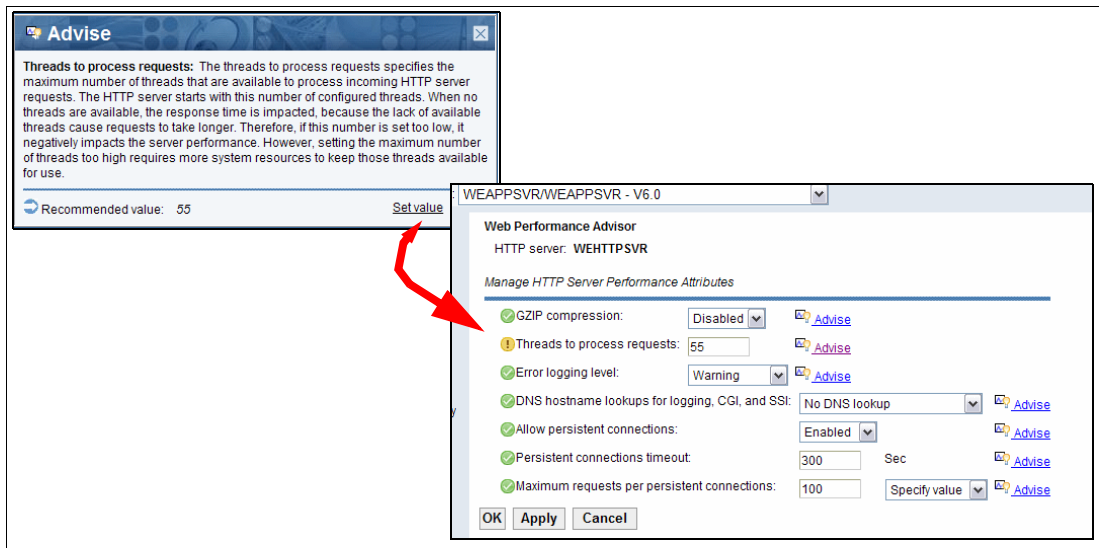


Figure 14-19 Advise window for threads to process

We click **Set value** in the lower right of the Advise window to change to the recommended 55 value, as shown of Figure 14-19. Then we click **Apply** shown in the lower right window.

Now, using the window shown in Figure 14-11 on page 359, we select **WEAPPSVR** and then click **Manage attributes**.

The four categories of Web Environment parameters for the Web application server are:

- ▶ Java Virtual Machine (JVM) parameters
- ▶ Application server resources
- ▶ Java Database Connectivity (JDBC) parameters
- ▶ Additional Web application server settings

Refer to the Web Environment area of the window shown in Figure 14-11 on page 359. Click **Manage attributes** to initially open the JVM Settings tab window, as shown in Figure 14-20.

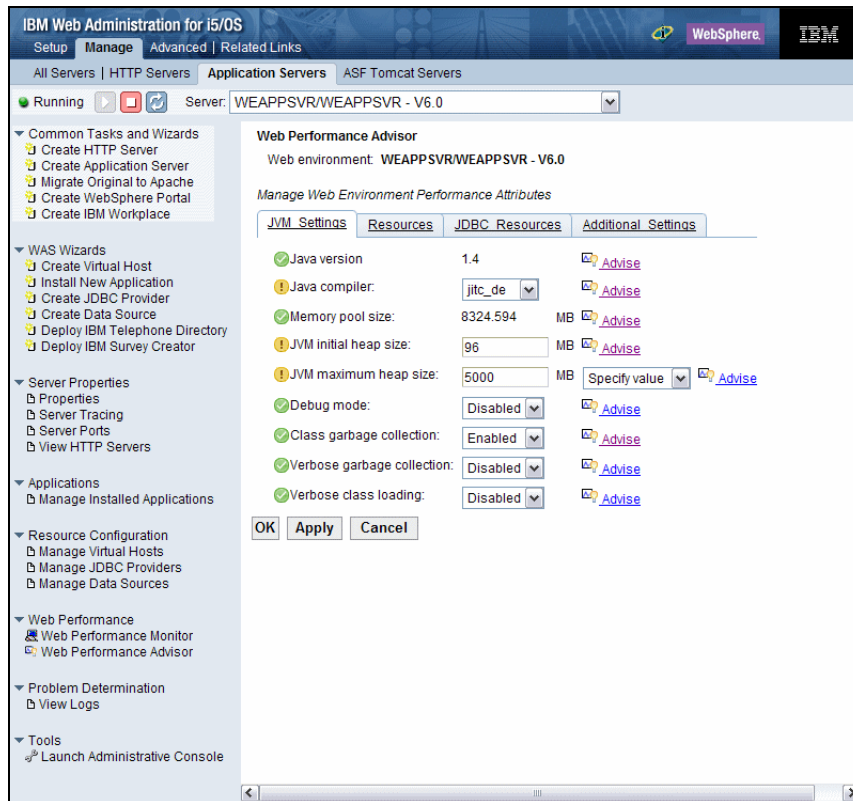


Figure 14-20 Web environment JVM settings

Note the important performance parameters in this JVM settings category. We select Advise for the three parameters with the exclamation point:

- ▶ Java version (initial setting jitc_de - direct execution)
- ▶ JVM initial heap size (initial size 96)
- ▶ JVM maximum heap size (initial size 5000)

Figure 14-21 shows the Advise windows for each of these settings.

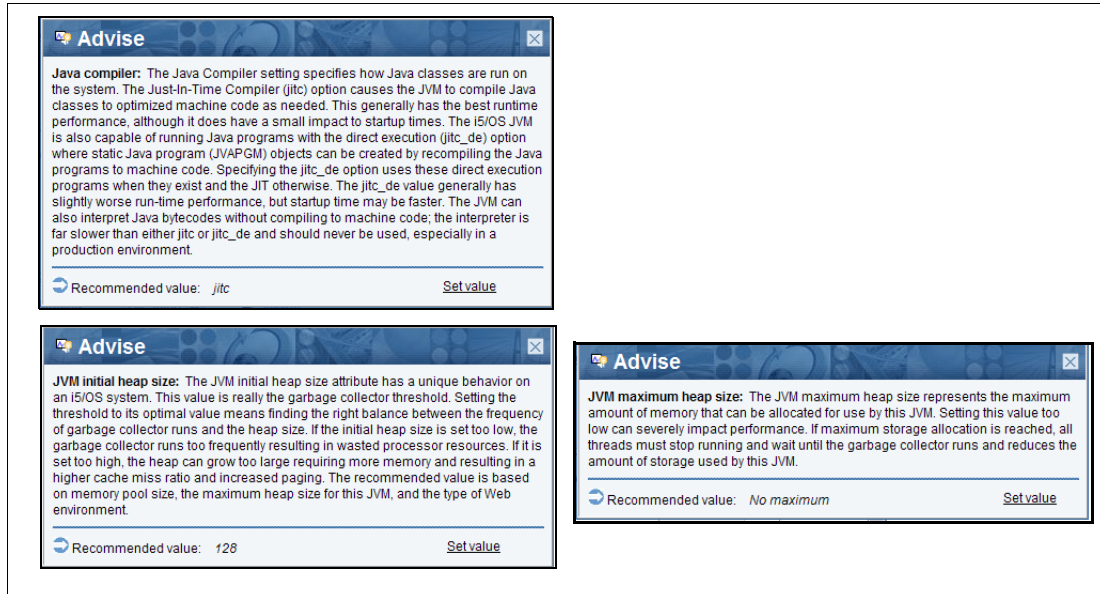


Figure 14-21 JVM settings advise examples

You can review the text in all of the Advise windows shown.

One of the important WebSphere Application Server for i5/OS parameters is the Java compiler one. The Direct Execution setting, `jltc_de`, was originally the fastest performing compiler option under i5/OS. However, in the last few years the standardization of Java compilers across IBM has resulted in the Just In Time (JIT) setting (`jltc`) being the best performing choice under i5/OS.

We decided to make some of the advised changes. Figure 14-22 shows all of the changes we selected to make, before applying them.

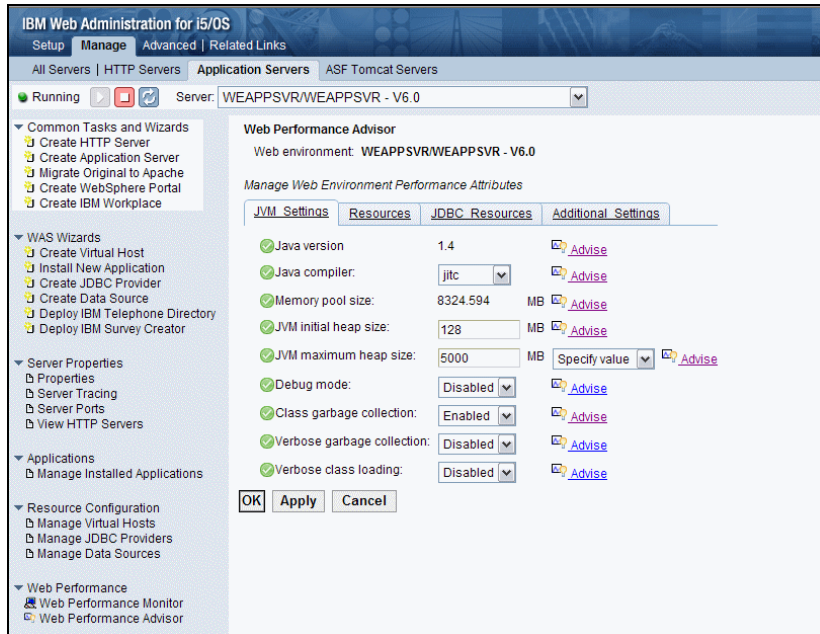


Figure 14-22 JVM settings after completing all changes to be applied

We clicked **Apply**. After a few seconds the apply finished.

We also selected the **Resources** tab, **JDBC Resources** tab, and then the **Additional Settings** tab. In this book, however, we show only the windows representing these settings. We do not show additional advise function examples. The information in these tabs gets into WebSphere Application Server details that are beyond the scope of this book.

The content in this chapter should be sufficient to encourage you to take advantage of the Web Performance Advisor functions initially delivered during November 2006. Combined with the other *Web enablement* enhancements described in this chapter, i5/OS offers a much easier way to get started using WebSphere Application Server-based environments than you having to create your own applications.

Click the **Resources** tab to get the window shown in Figure 14-23.




Figure 14-23 Advisor - Web Environment Resource

The resource settings cover i5/OS-specific settings for the default subsystem in which WebSphere Application Server for i5/OS V6.0 runs, Java heap size parameters (also covered under JVM settings), and several WebSphere specific settings.

Click the **JDBC Resource** tab to get the window shown in Figure 14-24.

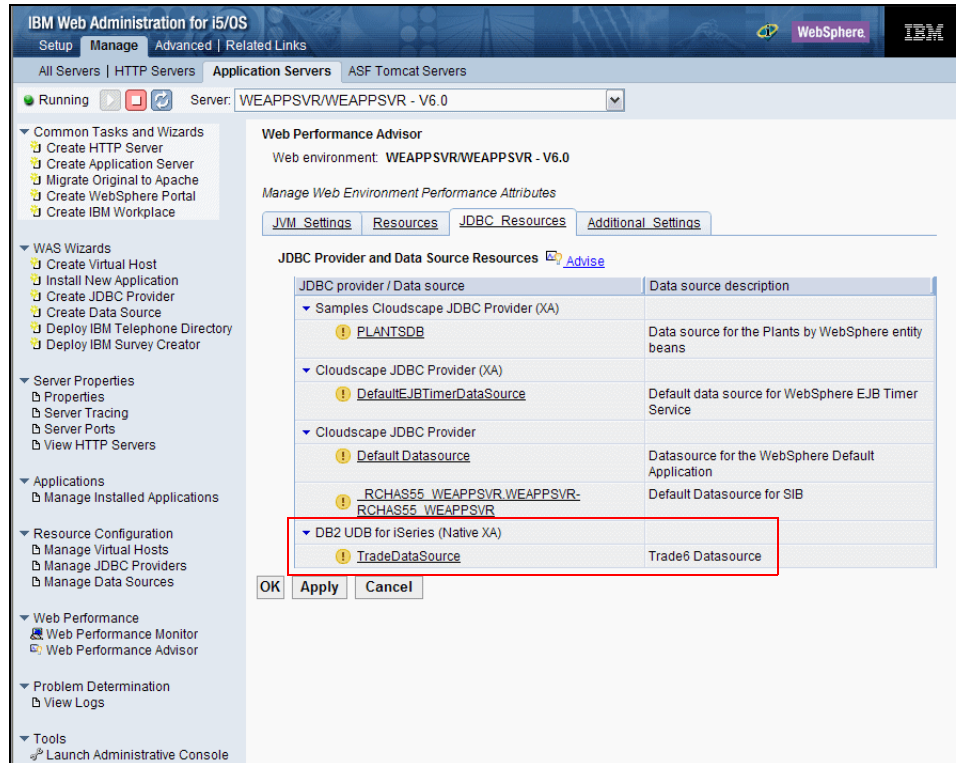


Figure 14-24 Advisor - Web environment - JDBC Resources

A JDBC provider tells the application server what type of access method should be used when it needs to access information from a database. A data source contained within a JDBC provider links to the actual database. Under i5/OS there are two commonly used i5/OS jobs that can be used under these interfaces, regardless of the names you see for JDBC provider or data source. In the actual definition of the JDBC provider-data source you specify one of two:

- ▶ Local - IBM Developer Kit for Java JDBC driver (the Native JDBC driver)
- ▶ Remote - IBM Toolbox Java JDBC (™) driver

Local functions will run in one of the database server jobs included under i5/OS with the QSQRVR job name prefix. This interface is unique to i5/OS.

Remote functions will run in one of the database server jobs included under i5/OS with the QZDASOINIT job name prefix. This interface is common across Java implementations on all operating systems.

The JDBC resources are defined under your WebSphere Application Server profile, in our example WEAPPSVR. Remember that the name itself does not identify which local or remote SQL interface under i5/OS is being used.

For WEAPPSVR, we have several JDBC providers or data sources listed, some of which you may not expect if you are new to WebSphere. It is up to the Web applications to use these data sources or go to the appropriate database directly from the application code. For example, in this chapter neither iSeries Access for Web or the flight400 applications use the WebSphere Application Server JDBC data sources.

However, without getting into details, there are some advantages for the applications specifically designed to run under WebSphere Application Server to interface to the appropriate database through the WebSphere Application Server JDBC resources.

In our example application server, WEAPPSVR, you can see JDBC provider/data source resources that include those shipped with IBM-provided, but optionally installed, sample applications that come with WebSphere Application Server on most supported operating systems.

Note the TRADE6 data source for the TRADE6 internal IBM application, which we optionally installed into WEAPPSVR.

With this JDBC Resources set of parameters, you can view the Advise window information. A more in-depth discussion of the performance aspects of the Advise window information in this area is beyond the scope of this book.

Last, but not the least, we show a screen example of the Additional Settings tab window in Figure 14-25.

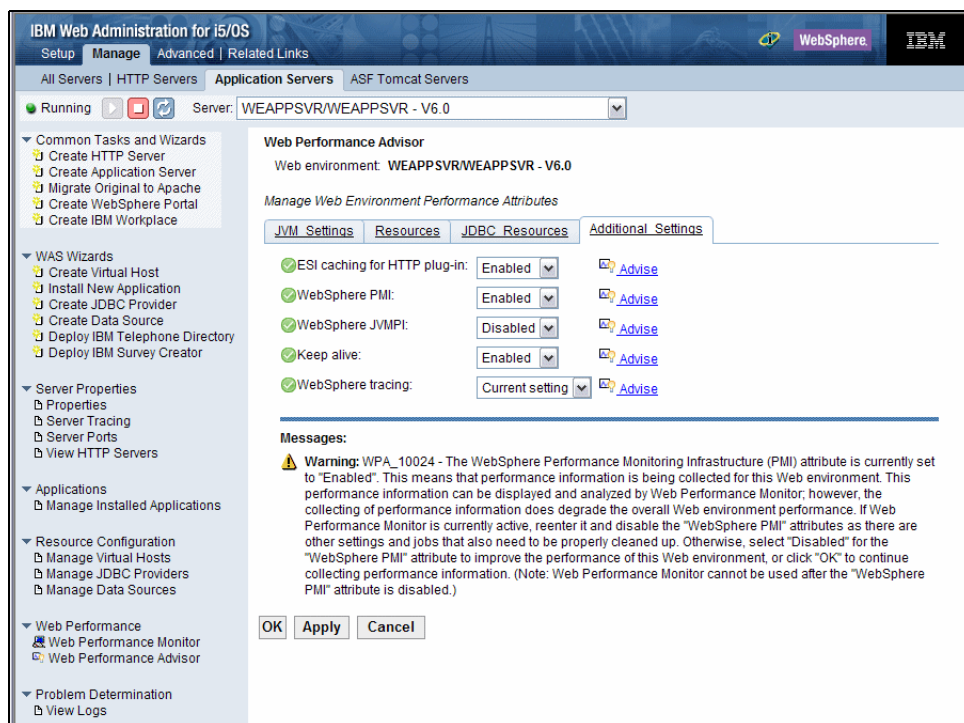


Figure 14-25 Advisor - Additional Settings

Additional settings contains miscellaneous performance-related parameters. Note the parameters set to enabled and disabled with no exclamation points shown. However, as indicated under Messages, note that the WebSphere Application Server profile WEAPPSVR does have the WebSphere Performance Monitoring Infrastructure (PMI) enabled to collect performance information.

14.4 WebSphere Application Server performance data

With a set of PTFs installed and with WebSphere Application Server V6.0 or V6.1 active with applications, Collection Services will automatically include a subset of the WebSphere

Application Server performance data in its collected data. When the “create performance database files” option is selected, new QAPMWASxxx files are produced:

- ▶ QAPMWASCFG - WebSphere server configuration
- ▶ QAPMWASSVR - WebSphere servers
- ▶ QAPMWASAPP - WebSphere applications
- ▶ QAPMWASEJB - WebSphere EJB data
- ▶ QAPMWASRSC - WebSphere pooled resource data

Refer to V5R4 Information Center for the field definitions for these files. Figure 14-26 shows an excerpt for file QAPMWASSVR.

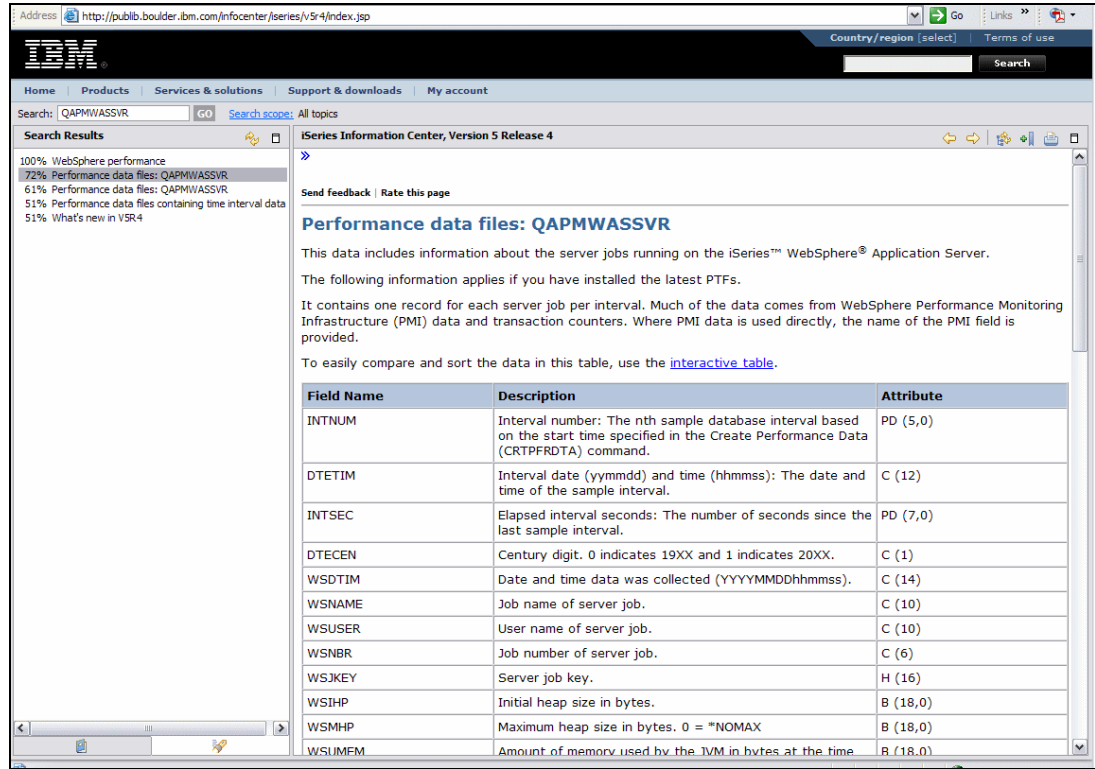


Figure 14-26 Example of QAPMWASSVR performance data files

Through October, 2006 there were no known applications that processed the data in these new WebSphere Application Server performance database files. We anticipate that some Independent Software Vendors (ISVs) will explore graphically displaying this data.

This data will eventually be summarized and processed by Performance Manager for System i (PM for System i) as part of its trend analysis functions.

See the following Web site to determine the PTFs required to install the capability to collect this performance data.

<http://www.ibm.com/servers/eserver/series/software/websphere/wsappserver/news/sitenews.html>

14.5 Simplified use of SSL with i5/OS HTTP Administration server

You can use the HTTP Administration server (*ADMIN) graphical interface to configure itself to use SSL-based encryption with a short set of wizard-like windows. When you use this new wizard support you get very clear instructions on each Web page as you proceed through the set up windows.

Setting up encryption of either or both TCP/IP connection handshaking and data transmission for the first time can be intimidating, as you need to learn and use new terminology and you need to set it up for each application you want to secure.

This whole process is commonly referred to as configuring and using the industry-wide Secure Sockets Layer (SSL) support, which introduces the terms Certificate Authority (CA), digital certificate, signed certificate, and certificate store.

With this new Web enablement support you can quickly set up and use basic SSL support for the Administration instance provided with HTTP Server for i5/OS, 5722-DG1.

Usually, the HTTP Administration server (*ADMIN) shipped with 5722-DG1 uses port 2001. You connect your browser to it entering a URL like:

```
http:rchas55:2001
```

During the SSL setup dialogue you are asked whether you want to keep using port 2001 (not secure), as well as the new secure port (2010) that will be used. This is your choice.

Assuming that you complete the SSL set up for *ADMIN, you would connect using the secure URL format:

```
https:rchas55:2010
```

The following figures show some of the windows you see setting up the administration server to use SSL.

By observing the text on the wizard windows you can become familiar with the SSL setup enough to consider configuring other applications within i5/OS to use SSL encryption. Setting up a digital certificate and assigning it to, for example, Telnet, is a little more complex than the windows you see here. However, by recalling the windows you see here and using Information Center articles on setting up SSL and digital certificate processing, you should feel comfortable about securing other applications, such as Telnet or Management Central.

Figure 14-27 shows the first two windows you need to start to set up the Administration server to use SSL. To get to the window on the upper left, we entered `http:rchas55:2001`.

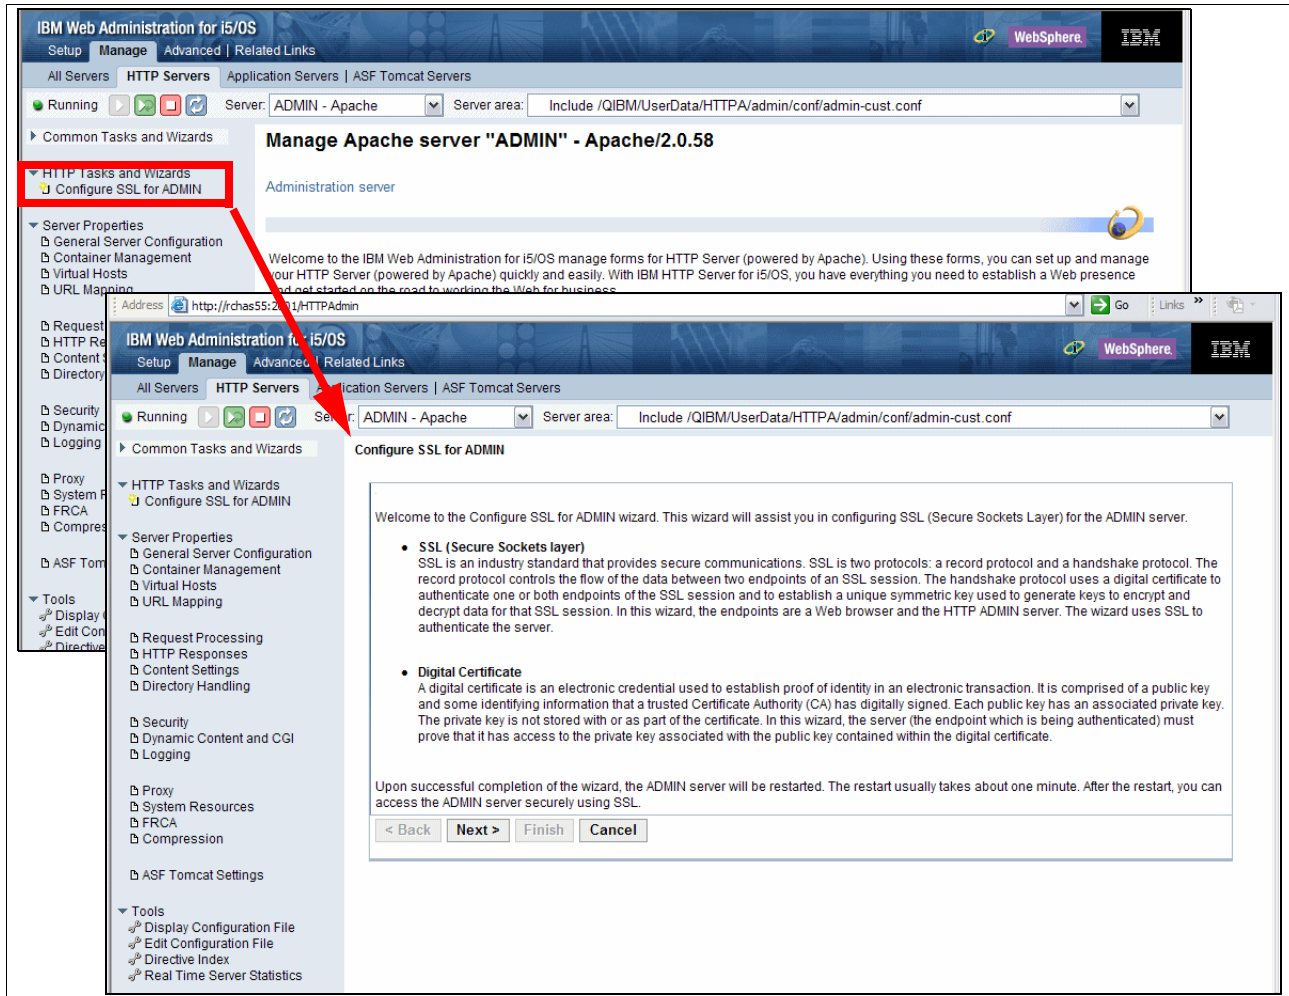


Figure 14-27 Setting up the HTTP Administration server to use SSL (1 of 2)

Review the SSL text and digital certificate descriptions in the lower right window. Click **Next**. On the next window (not shown here), you determine whether you want to disable port 2001 at the end of configuring the HTTP Administration server to use secure port 2010. This is your choice. You can disable use of port 2001 later if you wish.

Click **Next** on that window and you get the upper left window shown in Figure 14-28.

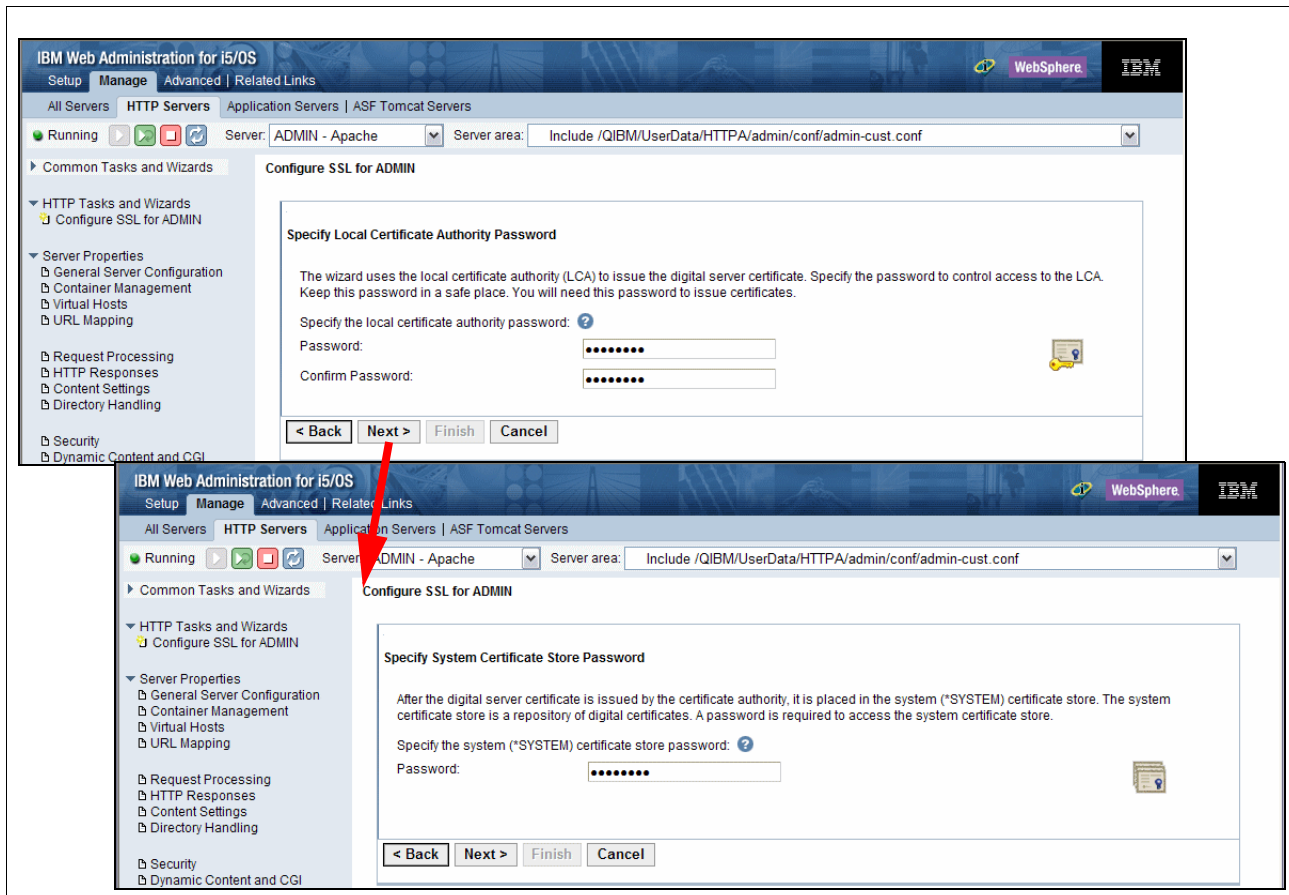


Figure 14-28 Setting up the HTTP Administration server to use SSL (2 of 2)

In this window you specify (twice for confirmation) the password for the Local Certificate Authority (LCA) if the LCA does not yet exist. i5/OS comes with a capability to have an LCA.

If this is the first time that you are using the LCA you can make up any password that you can remember.

Entering a successful password enables the wizard to create an LCA. If an LCA had already existed you would have a single password prompt.

Make sure that you save the LCA password, as you would need it later to make changes, and enables you to do more SSL and digital certificate work.

After successfully entering a password, click **Next** to continue to the window shown in the lower right of Figure 14-28.

In the lower right window you must enter the password for the i5/OS-provided *SYSTEM certificate store. This is where the digital certificate created by this process will be stored.

Since this is the first time we are using *SYSTEM, we enter the password twice and save it in a secure location. This will create the *SYSTEM store as part of this process. If the *SYSTEM store already exists, you must enter the correct password.

After entering the correct password click **Next**. This brings up the Summary window shown in Figure 14-29.

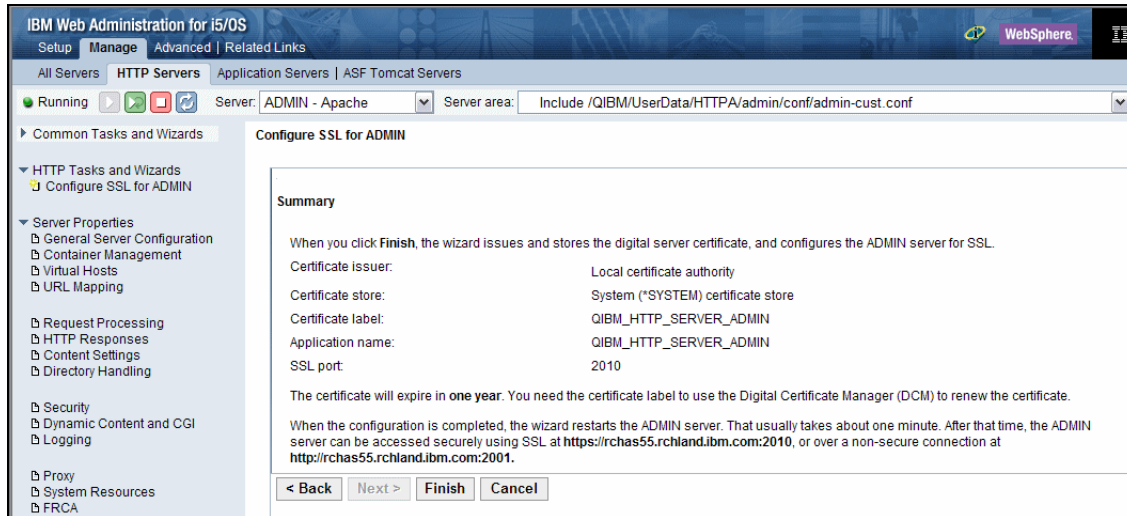


Figure 14-29 Summary: setting up the HTTP Administration server to use SSL

Note the text at the bottom of the Summary window. The text reminds you of the status of the original HTTP Admin server port 2001 and the new secure port of 2010. In our scenario we selected to keep both 2001 (non secure data) and 2010 (secure data) available.

We kept both ports active based upon our workload environment.

Click **Finish** to have the new certificate created and saved, and the HTTP Administration server back up and running and ready to process unsecure requests over port 2001 and secure requests of port 2010.



i5/OS printing and printer output enhancements

This chapter covers the enhancements to spool management, including backup and restore, printing control and forms design, and hardware announcements, including radio frequency identification (RFID) support.

It also summarizes new printers supported during 2006 and provides a Web site you can access to keep up to date on the entire list of i5/OS supported printers.

15.1 Introducing printing and printer output enhancements

Figure 15-1 depicts the various printing hardware devices and new printer software support introduced during 2006.



Figure 15-1 Printing systems division: printers and printer software

15.2 i5/OS printing and printer output management

Print and output management continues to be a strong opportunity for businesses to gain a strong ROI with new solutions. Many businesses continue to gain cost savings from printer, fax, and other output device consolidation. Transformation to electronic delivery and storage of output continues to be a key element of business re-engineering, content management, and regulatory compliance.

The Infoprint® line of output middleware products continues to evolve. Infoprint Designer provides for integrated graphical re-design of existing documents and reports. Infoprint Server addresses the exploding demand for PDF and electronic document delivery. Virtually every System i organization is seeking to adjust the mix of print and electronic distribution of business communications. New to this middleware set is Mapping Suite. Mapping Suite combines a high function query tool and advanced document design with integrated document production.

Important: IBM remarkets Mapping Suite. The Mapping Suite product set is owned, warranted, and supported by Cogninord.

New support in i5/OS V5R4 to save and restore spooled file data and attributes helps companies with record retention and disaster recovery. Now system save/restore commands provide a new spool file data parameter that can be used to save all spooled files or just new files. APIs are available to select and subset files to be saved or restored based on an extensive list of attributes including output queue, job name, job number, user name, spooled file name, and create date/time ranges. An expiration date can also be set on spool files.

For companies that use traditional system menus, the save and restore menu options have been enhanced to offer the option to save and restore all spool files. Because spooled files are saved directly to tape, the new support is much faster than most existing methods of save/restore of spooled files, many of which used to convert the file to a document, then save the document. The new support is also integrated with Backup and Recovery Media Services (BRMS). There are new QSPGETF and QSPPUTF APIs to make the spooled file library list become the library list of the job running the QSPPUTF API.

In industrial environments, reliable output management and radio frequency identification (RFID) capability are becoming critical to supply chain efficiency. To keep your information flowing, IBM offers IBM Infoprint 6700 (shown in Figure 15-1 on page 378, middle right) thermal printers. These printers provide flexible media handling and support for a wide range of bar code formats, including 2D barcodes. By integrating Infoprint 6700 printers with IBM middleware architecture, you can monitor printer activities and output processes. With additional support for RFID tags, as well as an upgrade path to the next generation of RFID technology, IBM can provide you with an end-to-end output management solution.

A relatively new printer category is the MFP, or multi-function printer. Also termed *all-in-one devices*, they combine scan, copy, fax, and print in one integrated machine footprint. The IBM Infoprint line of MFPs now totals seven devices ranging from 25 pages per minute (ppm) to 55 ppm, and includes both color and monochrome options.

The IBM Infoprint product line for monochrome and color laser printers is all new, featuring the Infoprint 15xx and 16xx series. These printers cover the range from 24 ppm to 50 ppm and monthly volumes from 15,000 to 300,000 impressions. All support PCL printing (via Host Print Transform) and most offer IPDS™ print (via Print Services Facility™ (PSF)) or PDF printing (via Infoprint Server) as options.

Note that Printer Command Language, more commonly referred to as PCL, is a Page description language (PDL) developed by HP as a printer protocol and has become a de facto industry standard. Originally developed for early Inkjet printers in 1984, PCL has been released in varying levels for thermal, matrix printer, and page printers. HP-GL and PJJL are supported by later versions of PCL.

The following sections provide additional details. For more information than what is covered in this section see:

<http://www.printers.ibm.com/>

15.3 Native save/restore spooled files

Backing up spooled files onto backup media reduces the number of active spooled files, improves spooled file operations performance, and protects spooled file data from loss.

15.3.1 History and changes

In the first releases of OS/400, spooled files did not survive IPL. Subsequent releases of OS/400 added the ability to save spooled file data in different ways. APIS QSPGETF and QSPPUTF allow users to save spooled files to physical file members with most of their attributes intact.

The reason spooled files are difficult to save is that they are not objects. Prior to V5R4, any attempts to back up spooled files encountered problems due to the nature of spooled files and the way that they are identified. The spooled file was identified by the spooled file name,

but also by the qualified job name of the job that created the spooled file and the spooled file number that was also allocated based on that job.

Formerly, when an output queue is saved, its description is saved but not the contents (spooled files). Therefore, a number of work-around solutions evolved. With a combination of spooled file APIs, user space APIs, and list APIs, programs could be written to back up spooled files. Another way is to copy spooled file information to user spaces and space, from where it can be backed up.

The Copy Spooled File (CPYSPLF) command copies textual data from spooled file to database file. However, print attributes such as spacing and advanced function attributes are not saved.

BRMS provided spooled file backup by saving spooled files as documents into a folder. How this works is that the original spooled file attributes are retained during save operation (including names, user names, user data fields, and job names). During restore, new spooled files are created (new job numbers, system dates, and times), so the true identity of spooled files is lost. In addition, the library list that was saved with the original file was replaced with the library list of the job doing the restore. This caused problems printing the spooled file if it used any resources that were specified using *LIBL for the library name.

When complying with regulatory issues, this can become an issue.

V5R2 changes to spooling support allowed a spooled file to exist even though the job that created it had gone away. This was added with the QSPLFACN system value and SPLFACN JOBID attribute.

Additional fields were created to uniquely identify spooled files:

- ▶ Job system name (the system on which the job that created the spooled file ran).
- ▶ Create date and time.

These fields were needed in case spooled files were created by two different jobs with the same qualified job name either on the same system or two different systems. V5R3 allowed spooled files to exist in independent disk pools, which might be switched between systems.

Starting with V5R4 you have the capability to specify, for output queues, whether you want to save the actual spool file data (and its attributes) with the i5/OS commands, as well as with the Backup, Recovery and Media Services (BRMS).

On restore functions you can optionally specify to restore the saved spool file data or just the output queue. Support was added to display the save/restore media and see the list of spooled files with the commands DSPTAP, DSPOPT, and DSPSAVF. This is similar to existing support for displaying members within a database file.

The new method is faster than most existing methods of save/restore of spooled files. It does not require copies of the spooled file into database files or document library objects (DLOs). The actual spooled files with their spooled file attributes are saved and restored directly. V5R4 allows users to save and restore spooled files such that the restored spooled file is identical to the original including:

- ▶ Total identity
- ▶ Spooled file name
- ▶ Spooled file number
- ▶ Qualified job name
- ▶ Job system name
- ▶ Create date and time
- ▶ Most spooled file attributes
- ▶ Limitations covered later in presentation
- ▶ Spooled file job library list

This was one of the major limitations of the QSPGETF and QSPPUTF APIS — the spooled file library list would become the library list of the job running the QSPPUTF API.

Save/restore spooled files integrated into existing save/restore commands and APIs. These commands include save and restore objects (SAVOBJ/RSTOBJ) and libraries (SAVLIB/RSTLIB).

When using save commands, the user can specify either that spooled files in the saved output queue be saved or just output queue (no spooled files). There is a new parameter for save commands (SPLFDTA).

For restore commands, the user can specify whether to restore any spooled files saved with the output queue (SPLFDTA(*NEW)) or not to restore the saved spooled files (SPLFDTA(*NONE)).

15.3.2 Saving and restoring spool file examples

In this section we give examples of saving and restoring spool files.

First method: through options added to save and restore commands

A Spool File Data (SPLFDTA) option added to SAV/RST-related CL commands:

- ▶ SAVOBJ and RSTOBJ
- ▶ SAVLIB and RSTLIB
- ▶ SAVRSTOBJ and SAVRSTLIB

On SAVE commands it looks like:

```
Spooled file data . . . . . *NONE *NONE, *ALL
```

The default value for SPLFDTA, *NONE, makes the command work just like it did prior to V5R4, where only the output queue attributes are saved, but not any spooled files. By specifying *ALL for this parameter, all of the spooled files on any output queues being saved will also be saved. These spooled files will then be restored by default when the output queues are restored.

On RESTORE commands it looks like:

```
Spooled file data . . . . . *NEW *NEW, *NONE
```

In this case, the default value of *NEW indicates that any spooled files saved with the output queue should be restored if they do not already exist on the system. The user must specify

*NONE if spooled files that were saved are not to be restored. It is important to note that spooled files will not be restored if a spooled file already exists on the system with the same identity. If the user wants to replace a spooled file that already exists, the spooled file must be deleted before the restore.

Second method: QSRSAVO and QSRRSTO APIs

Spooled files can be saved and restored via the Save Objects List (QSRSAVO) and the Restore Objects List (QSRRSTO) application programming interface (APIs). APIs enable selection of spooled files to be saved (spool file identify, output queue, spooled file name, job name, owning user, user data, form type, and create date and time).

The QSRSAVO API enables calling a program to set the expiration date of the saved spooled file residing on the system that the spooled file was created on. Later a program can use the DLTEXPSPLF command to physically delete the saved spooled files off the original system because it has expired.

The QSRRSTO API provides additional options. Restore a copy of a spooled file if the spooled file already exists on the system. The restored spooled file will be identical to the original, but with a new create date and time that corresponds to the date and time of the restore operation. Restore spooled files in a different output queue than the one from which they were originally saved. The enables the calling program to set the expiration date for the restored spooled file if a spooled file already exists on the system with the same identity.

One additional option on the QSRSAVO API is the ability to set a new expiration date for the spooled files left on the system. This ability could be used to save a set of spooled files and set the expiration date for the copy of the spooled files that stay on the system such that at some future time, the ones on the system can be automatically deleted.

BRMS uses the new support, and while there are no user interface format changes, there are additional options and performance improvements.

Work with Output Queue								
Queue:	QPRINT	Library:	QGPL	Status:	RLS			
F11=View 1:								
Opt	File	File Nbr	Job	User	Number	Date	Time	
	QPDSPJRN	2	QPADEV0003	JOONWOOK	012143	12/13/05	10:44:33	
	QPDSPJRN	3	QPADEV0003	JOONWOOK	012143	12/13/05	10:45:25	
	QPDSPJRN	4	QPADEV0003	JOONWOOK	012143	12/13/05	10:46:38	
	QPDSPJOB	5	CRTPFRTDA	QSYS	012222	12/13/05	15:15:26	
	QPRINT	1	COMPACT	QNOTES	012526	12/13/05	15:47:31	
F11=View 2:								
Opt	File	User	User Data	Sts	Pages	Copies	Form Type	Pty
	QPDSPJRN	JOONWOOK		RDY	2	1	*STD	5
	QPDSPJRN	JOONWOOK		RDY	1	1	*STD	5
	QPDSPJRN	JOONWOOK		RDY	1	1	*STD	5
	QPDSPJOB	QSYS	FFDC	RDY	15	1	*STD	5
	QPRINT	QNOTES		RDY	1	1	*STD	5

Figure 15-2 Screen capture of view 1 and 2 of output queue QPRINT

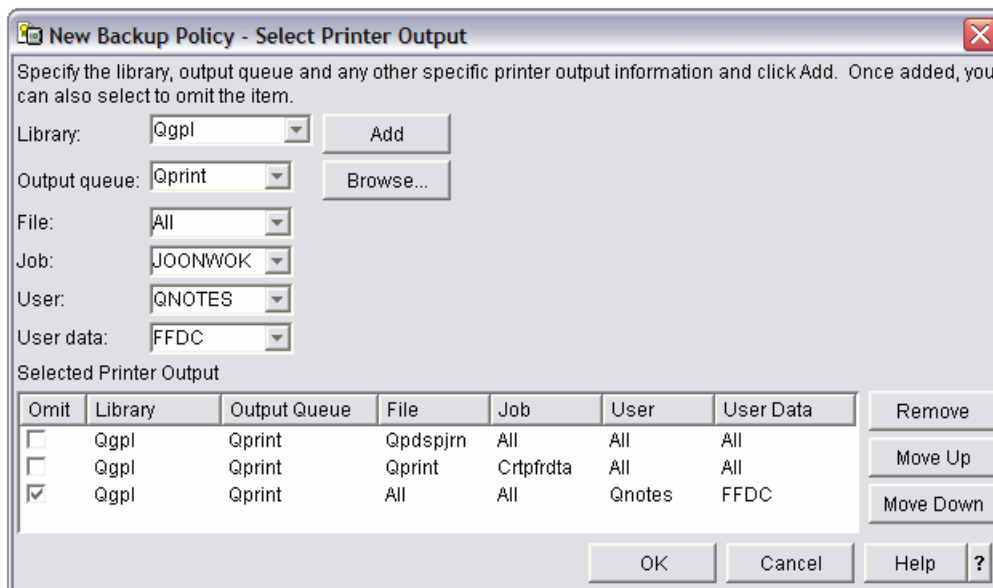


Figure 15-3 BRMS policy: printer output selection page

15.3.3 OUTQ and spooled files in iASP considerations

Placing spooled files in independent ASPs can be used to isolate low-used spooled files to off-line Direct Access Storage Devices (DASD), and also in a high availability scenario.

You place spooled files in an iASP by creating an output queue library in an independent ASP. Specify the value of *OUTQASP for the SPLFASP parameter when creating output queue (CRTOUTQ).

- ▶ Spooled files will be physically placed in the same iASP that the output queue resides in.
- ▶ The value of *SYSTEM for SPLFASP places spooled files in system ASP.

In V5R3, a library in an IASP can contain an output queue by using the parameter SPLFASP(*OUTQASP) on CRTOUTQ. The IASP Specification is supported in commands:

CHGSPLFA, DLTSPLF, HLDSPFL, RLSSPLF, WRKSPLF

With the support of *OUTQ in *SPLF object types in iASPs, at V5R3 of the OS/400 operating system the user is able to store up to 10 million spool files in a single iASP. Storing spool files in an iASP for archiving purposes is therefore an option. When storing spool files in an iASP, the user is enabled for making these spool files inaccessible by a vary off of the iASP. Limiting the access to spool files in an iASP to users with *JOBCTL or *SPLCTL authority can easily be accomplished by denying them authority to the iASP device description. The user can now also specify a writer to start on an output queue when his job has the iASP of that output queue in its namespace.

Since the spool files in an iASP will always be decoupled from a job when that job ends, and after a switch or failover, the way to find spool files in output queues with a large number of entries is by JOBSYSNAME and CRTDATE attributes.

The new functions in OS/400 V5R3 will make the spooled file recovery extremely fast during the vary-on phase of an iASP (compared to the same recovery time required during a server IPL). In addition to this enhancement, cross-reference synchronization is also faster with OS/400 V5R3.

The ability to store spooled files in an iASP also causes changes to the output files QASYSFJ5 and QASYPOJ5 containing the audit entries for printed or spooled output. When SystemValue QAUDLVL or QAUDLVL2 contains the entry *PRTDTA, the journal entries in the QASYPOJ5 file are used for security auditing of generating printed output (for example, printing with parameter Spool *NO will cause an entry of the type PO). A *SPLFDTA entry in SystemValue QAUDLVL or QAUDLVL2 will cause entries with journal type SF in output file QASYSFJ5 for actions such as getting data from a spooled file (QSPGETSP).

Table 15-1 V5R3 iASP object support

Supported			Not supported	
*ALRTBL	*IGCDCT	*QMFORM	*AUTL	*IPXD
*BLKSF	*JOB	*QMQR	*CFGL	*JOBQ
*BNDDIR	*JRN	*QRYDFN	*CNL	*JOBSCD
*CHTFMT	*JRNRCV	*SBS	*COSD	*LIND
*CHRSF	*LIB	*SCHID	*CRG	*MODD
*CLD	*LOCALE	*SPADCT	*CSPMAP	*M36
*CLS	*MEDDFN	*SPLF New	*CSPTBL	*M36CFG
*CMD	*MENU	*SQLPKG	*CTLD	*NTBD
*CRQD	*MGTCOL	*SQLUDT	*DDIR	*NWID
*CSI	*MODULE	*SRVPGM	*DEVD	*NWS
*DIR	*MSGF	*STMF	*DOC	*PRDAVL
*DTAARA	*MSGQ	*SVRSTG	*DSTMF Alert	*PRDDFN
*DTADCT	*NODGRP	*SYMLNK	*EDTD	*PRDL
*DTAQ	*NODL	*TBL	*EXITRG	*RCT
*FCT	*OUTQ New	*USRID	*FLR	*SOCKET
*FIFO	*OVL	*USRQ	*IGCSRT	*SSND
*FILE	*PAGDFN	*USRSPC	*IGCTBL	*S36
*FNTRSC	*PAGSEG	*VLDL	*IMGCLG	*USRPRF
*FNTTBL	*PDG	*WSCST		
*FORMDF	*PGM			
*FTR	*PNLGRP			
*GSS	*PSFCFG			

Alert: SAV of QNTC previously mentioned as supported.

15.3.4 Restrictions and limitations

In this section we discuss restrictions and limitations.

Spooled files cannot be saved

Spooled files cannot be saved:

- ▶ When the spooled file status is RDY and spooled file is SCHEDULE(*IMMED) and still open.
- ▶ When the spooled file status is OPN.
- ▶ When the spooled file status is FIN.
- ▶ When the save operation cannot lock the spooled file.
- ▶ When the user is not authorized to save the spooled file.
- ▶ When the spooled file is damaged.
- ▶ No support for restoring over an existing spooled file.
- ▶ CL command SAVCHGOBJ will not support spooled files.

Print fidelity

In order for a user/customer to have 100% print fidelity the customer must do the following:

- ▶ Save/restore all external resources within the libraries that they existed in when the spooled file was created.
- ▶ Make sure the user profile (spooled file owner) exists and has the appropriate authorities to all the external resources required for the spooled file to print (IFS directories, TrueType font IFS files, font resource objects, overlays, form definitions, page definitions, page segments, and imbedded object IFS files).
- ▶ Make sure the Resource Allocation Table (RAT) has the same fonts and linked fonts as when the spooled file was created. This is only necessary for those spooled files that use TrueType linked fonts.
- ▶ Make sure the TrueType fonts (if used by spooled files) are on the system and in the appropriate directory.
- ▶ Make sure the environment variable QIBM_AFP_RESOURCES_PATH is set to the appropriate path for those spooled files that use imbedded objects but were not directory qualified.
- ▶ Make sure the imbedded objects that are not in the directories are restored into the same directories they existed in when the spooled files that used them were created.

Restrictions on iASPs

Because a spooled file must be unique across a system, an exact duplicate spooled file presents numerous problems.

The fact that one or more iASPs can be on or offline at any given time for a particular system also lends itself to some issues. Thus, upon restore the system will check to see if the spooled file exists on an iASP before allowing a restore into *SYSBAS.

Also, the system will fail the vary on of an iASP if a duplicate spooled file is found. A restore of a spooled file into an iASP group that already exists in *SYSBAS will fail. A move of a spooled file from an iASP group A to *SYSBAS will fail if the same spooled file also exists in another iASP group B. A vary on of an iASP group will fail to active state if a spooled file exists in both *SYSBAS and the iASP group.

Each spooled file that has a duplicate in *SYSBAS will be identified by a diagnostic message in the joblog of the job that attempted the vary on of the iASP group. To successfully vary on the iASP group, the duplicate spooled files in *SYSBAS must be deleted using the DLTSPFL CL command. The iASP group must then be varied off and then back on again.

Spooled file attributes that are not preserved

These are:

- ▶ System 36 identifier
- ▶ IPP job identifier

Spooled file attributes that may not be preserved

These are:

- ▶ Output queue name and library (if not restored to output queue it was saved from)
- ▶ Expiration date (if restore time is beyond the expiration date or overridden via API)
- ▶ Creating date and time (only if the option is chosen to duplicate file)
- ▶ Auxiliary storage pool number (only if target output queue on different ASP or the iASP has been recreated and was assigned a number different than it had previously)

Considerations for the APIs

When using QSRSAGO API to save a group of spooled files through the use of the selection criteria, if a spooled file would be selected more than once, it will only be saved on the media one time. Therefore, when BRMS calls the API to save spooled files in place of its current design, the customer will no longer have the spooled file saved more than once.

When using QSRRSTO API to restore a group of spooled files through the use of the selection criteria, if a spooled file would be selected more than once, only the first matching selection criteria will apply. The spooled file will only be restored once. All other matching criteria will be ignored.

15.4 Spool file, policy-based output queue management

In this section we discuss the spool file, policy-based output queue management.

15.4.1 V5R4 introduces spooled file expiration

This is a simple way to automatically expire spooled files. The function is provided via new EXPDATE and DAYS spooled file attributes. You can access these attributes via the following commands:

- ▶ Change Printer File (CHGPRTF)
- ▶ Create Printer File (CRTPRTF)
- ▶ Change Spooled File Attributes (CHGSPLFA)
- ▶ Override with Printer File (OVRPRTF)

Spooled files are not automatically deleted when expired. They must be deleted via the new Delete Expired Spooled files (DLTEXPSPLF) command. However, this is easy to automate via Job Schedule Entry.

System benefits include reduced disk storage requirements. You should also see better throughput for applications that deal with spooled files. This allows for savings through easier operator management of the system. There are indications that you may see faster IPL times. Lastly, unnecessary spooled files are not left around cluttering up the system.

The Expiration date for file (EXPDATE) attribute specifies the expiration date for the spooled file. The spooled file will expire at 23:59:59 system local time on the date specified. In your command, you may specify:

- ▶ *SAME

The expiration date does not change.

- ▶ *NONE

No expiration date is specified.

- ▶ *DAYS

The expiration date is to be calculated using the value specified for the Days until file expires (DAYS) parameter.

- ▶ date

Specify the date after which the spooled file will be eligible for removal from the system by the Delete Expired Spooled Files (DLTEXPSPLF) command. The date must be enclosed in apostrophes if date separator characters are used in the value.

► 1-366

Specify an interval in days after which the spooled file will be eligible for removal from the system by the Delete Expired Spooled Files (DLTEXPSPLF) command. The actual expiration date applied to the spooled file is calculated by adding the number of days specified to the date the printer file is opened.

Restriction: The Days until file expires (DAYS) attribute specifies the number of days to keep the spooled file. A value must be specified for this parameter if the Expiration date for file (EXPDATE) parameter has a value of *DAYS. If the EXPDATE parameter has a value other than *DAYS, no value is allowed for this parameter.

15.4.2 Action plan: gaining control over your output queue

A typical path to gain full control over your output queue involves two separate but equally important actions.

Cleaning up your existing files

There are a number of steps you will need to take to do this:

- Use Operational Assistant or the CHGCLNUP command to clean up QPEZJOBLOG and QEZDEBUG output queue entries.
- Use the WRKOUTQ command to access spooled files by queue and manually delete each unwanted file.
- Use WRKSPLF to access spooled files by specific selection criteria and manually delete each unwanted file.
- Use the V5R4 function to set the *lifetime* of each spooled file.

You can run the V5R4 Delete Expired Spool Files (DLTEXPSPLF) command manually to automatically delete all expired spooled files, or you can schedule a batch job to run this command at regular intervals. To set up the batch command, you can use either the Add Job Scheduled Entry (ADDJOBSCDE) command or the Advanced Job Scheduler product. More information about these two approaches can be found at the Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r4/index.jsp>

Remember that expired spooled files are not deleted until this command is executed or the file is deleted manually.

The Operational Assistant provides automatic cleanup of spooled files (job logs) in output queue QEZJOBLOG. The DLTEXPSPLF command deletes spooled files (if expired) located in any output queue on the system or in iASPs.

Enhanced job log management

Reducing number of spooled job logs created during system operations can increase response time of system and end-user operations.

As an example, when many jobs end at the same time, creating job logs to QEZJOBLOG output queue, the completion of job end operations will be delayed. WRKOUTQ with the QEZJOBLOG queue will create additional delays.

We advise that you change the logging levels of job descriptions to only create spooled job logs when needed. You should also spread the job log spooled files to several output queues.

You can use pre-V5R4 functions such as the job description LOG parameter and ENDJOB maximum job log entries, ENDSBS ENDSBSOPT options like *NOJOBLOG.

V5R4 provides additional enhancements to enable users to only create spooled job logs when needed. There is a new job log output system value (QLOGOUTPUT) and job attribute (LOGOUTPUT). If you specify a value of *PND, you will create a job log *on demand* (see the Work with Job Log (WRKJOBLOG) command below). If you specify a value of *JOBLOGSVR, you enable the creation of a spooled job log via a job log server after the job has ended.

Use the WRKJOBLOG command to manage job logs:

- ▶ Display job logs in *PENDING status or jobs in *Spooled status (job logs already printed to spooled files and job ended).
- ▶ Specify selection criteria for getting lists of job logs (start/end time/date of job logs, job name, user, number, generic value).
- ▶ Display and delete job logs.
- ▶ Work with the job (if the job log is in *PND status).

The WRKJOBLOG command works well if spooled job logs are created across multiple output queues to reduce contention. It allows you to see all job logs from a single screen regardless of what queue they reside in. The command also does not conflict with creation of spooled job logs.

15.5 Changes Work With Spooled Files (WRKSPLF)

New enhancements to the WRKSPLF command in V5R4 include additional filtering options, which can reduce the size of the list returned back to the user. Select spooled files by:

- ▶ Spooled file name
- ▶ Qualified job name
- ▶ Starting/ending create dates/times
- ▶ Generic values

Spool filtering can help you focus on only the spool items you want to. There are some cautions, however, if you use WRKSPLF with thousands of spool file entries.

These additional filtering options can cause WRKSPLF to interfere with intense spooling applications much like WRKSPLF USER(*ALL) can. Therefore, test using these filtering options in your environment with the other filtering criteria (user, user data, device name, ASP). Some steps to consider in large spool environments are:

- ▶ Placing spool files in a user auxiliary storage pool (ASP) can help limit the number of files being processed.
- ▶ Use the V5R4 filtering parameters of the WRKSPLF command to reduce the size of the spooled file list. It is easier to work with spooled files from a smaller list.
- ▶ Avoid the use of WRKSPLF USER(*ALL) if possible. Filter on user, user data, device name, or ASP.
- ▶ Filtering on this criteria can decrease time to build the list of desired spooled files.
- ▶ Filtering on form type *STD may be inefficient if most spooled files on the system use the STD form type.
- ▶ Consider using WRKJOB OPTION(*SPLF) or WRKOUTQ as an alternative to obtaining a list of spooled files.

- ▶ WRKSPLF has less performance impact if working with spooled files stored in an iASP.

Figure 15-4 shows the additional options in the WRKSPLF command. Note especially the options for spool file, job name, and time period.

```

Work with Spooled Files (WRKSPLF)

Type choices, press Enter.

Select files for:
  User . . . . . *CURRENT      Name, generic*, *CURRENT...
  Print device . . . . . *ALL        Name, *ALL, *OUTQ
  Form type . . . . . *ALL        Form type, *ALL, *STD
  User data . . . . . *ALL        User data, generic*, *ALL
  ASP . . . . . *ALL          1-32, *ALL, *ASPDEV
  Spooled file . . . . . *ALL      Name, generic*, *ALL
  Job name . . . . . *ALL      Name, generic*, *ALL
  User . . . . . *ALL        Name, generic*, *ALL
  Number . . . . . *ALL      000000-999999, *ALL

Time period:
  Start time and date:
  Beginning time . . . . . *AVAIL    Time, *AVAIL
  Beginning date . . . . . *BEGIN    Date, *BEGIN
  End time and date:
  Ending time . . . . . *AVAIL    Time, *AVAIL
  Ending date . . . . . *END      Date, *END
  ASP device . . . . . *          Name, *, *SYSBAS, *CURASPGRP
  Output . . . . . *            *, *PRINT

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys

```

Figure 15-4 Work with spooled files: note especially job name and time period options

15.6 Print output formatting and routing

System i printing software products are covered in depth at:

http://www.printers.ibm.com/Internet/wwwsites.nsf/vwebpublished/iseriessoftware_ww

In this section we cover Infoprint Server (5733-IP1), Infoprint Designer (5722-ID1), and Host Print Transform (HPT).

15.6.1 Infoprint Server for iSeries

Infoprint Server for iSeries (5722-IP1) is really focused on extending an enterprise's communications beyond print. It supplies a range of functions in this regard, enabling a company to start small and grow in function and sophistication over time. At a basic level, Infoprint Server provides PDF services for System i output. Every System i customer needs this core capability. A print file can be transformed to PDF and routed to print, e-mail, or the Web.

Once you start, however, you cannot stop there. With the addition of Intelligent Routing and new interfaces for application definition, you can use Infoprint Server to implement complex electronic delivery of documents and reports. This can include profiling and customization by recipients, electronic segmentation of multi-document spooled files, and complete control over e-mail content.

A typical System i customer might start by routing some month-end reports in PDF format to the Integrated File System in lieu of printing them. They might end by taking critical customer communications (statements, invoices, and so on) and enabling recipients to determine in what formats they want to be sent in.

Also note that there is an optional feature of Infoprint Server that provides for the transformation of other print data streams (that is, Postscript, PDF) into System i format for printing on System i printers.

Infoprint Server had a broad series of enhancements, including Intelligent Routing and new user interfaces, in V5R3. Recent enhancements have been delivered via PTF. An example is the set of Tiffs to support enhanced troubleshooting of PDF mapping programs.

15.6.2 Infoprint Designer for iSeries

Infoprint Designer for iSeries (5733-ID1) is a graphical document composition system for System i. More specifically, Infoprint Designer takes in the simple print line-based output of existing applications and enables you to re-engineer that output to fully electronic form. You can include a complete range of fonts, images, overlays, bar codes, and document logic.

Infoprint Designer is Windows-based with tight integration (via iSeries Access) to the System i environment. It comprises overlay, image, and layout editors that work together to produce new documents. An essential characteristic of Designer is that no application change is required. The design changes are embodied in standard System i objects.

Enhancements for Infoprint Designer are now being delivered via PTF. The nominal OS release continues to be V4R5. The current Designer version is 1.19. This is displayed using *Help – About Infoprint Designer* on the Windows interface.

Infoprint Designer is a good, integrated design system. But in this application space, one size does not necessarily fit all. Companies with complex documents (for example, insurance policies) may need other options. Feel free to contact Bill Shaffer (wcshaff@us.ibm.com) to explore the fit of Designer and additional options.

15.6.3 Mapping Suite

Mapping Suite (5639-AAC, AAD, AAE) provides comprehensive output management of business communications. Mapping middleware functionality covers the gamut from document design to print management to multi-channel delivery. The software system connects to a wide variety of input data sources that in turn feed into professional business document design. Mapping produces most of the standard print and electronic output formats required in business communications.

The base component, Mapping, provides for data connection and document composition. With a professional graphical design interface, it is geared to the redesign and repurposing of business output in today's very dynamic environment. Mapping includes the significant set of design functions required by complex documents and personalized communications. An optional product, MapReport, enables the design of new reports as well as the synthesis of document data from multiple input sources. An optional feature, MapOffice, provides for the automation of business correspondence.

Mapping output can be in standard print data streams, such as AFP, PCL, PDF, and thermal, as well as non-print channels such as PDF, Web, e-mail, and fax. More importantly, all of these distributions can be recipient-driven, ensuring that each person gets the correct communications in the most effective format. Mapping is tightly integrated with standard i5/OS print support, Print Services Facility (PSF), and Infoprint Server, as well as extending function in key areas.

For more information start at the iSeries print information at:

<http://www.printers.ibm.com/iseriess>

15.6.4 Host Print Transform

In this section we discuss Host Print Transform.

HPT now supports 2D barcodes

Two-dimensional bar codes offer customers many advantages. This is because these symbologies can encode substantially higher volumes of information in a smaller space. V5R4 HPT will support the printing of the following 2D bar code types:

- ▶ MaxiCode
 - Used by UPS
- ▶ PDF417
 - Used by Federal Express
- ▶ Data Matrix
 - High degree of redundancy for error correction
 - Can store up to 2000 characters
- ▶ QR Code
 - Developed in Japan. Can encode up to 7,089 characters in one symbol, including Kanji, Kana, Hiragana, symbols, binary, and control codes.
 - Possible to encode using DDS by specifying the Bar Code Object Content Architecture™ (BCOCA™) bar code type for QR Code (32) on the DDS Barcode Keyword and inserting special data fields into the record in front of the bar code data.

UCC/EAN 128 support enhanced for CODE128 bar codes

Changes to modifier values of 3 or 4 now ensure that the bar code data will be checked for compliance to the UCC/EAN 128 Application Identifier (AI) Standard. Message CPD6DF0 will be issued with reason code EC_1200 for incorrect data.

In the font code pages, certain positions or hex characters encode specific Function Controls, commonly called FN or FNC. There are four of these controls, which cause different actions to take place in the Human Readable Interpretation (HRI). The HRI is the encoded data in plain text, positioned typically above or below the bar code symbol. Reference the BCOCA Reference S544-3766, available online at:

<http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

See the Code 128 Code Page section.

These changes cause the FNC controls to be removed from the HRI. That is, they will not print, but will do what they are intended to do. Other modifiers, and in all other bar code types, the FNC and other unprintable characters are replaced with a hyphen.

For a modifier value of 4, the AIs in the HRI will be contained in parenthesis. The modifier value of 2 was supported prior to V5R4. This should generate the correct symbol for UCC/EAN 128. A modifier value of 1 causes checksum digits to be omitted from the symbol.



Figure 15-5 Code 128 modifier examples

Changes to bar code exception reporting

In prior releases, the invalid bar code was simply omitted. Now, bar code exceptions are logged with diagnostic message CPD6DF0. BCOCA architected exception codes used for reason codes. Secondary text lists the most common user-generated exceptions, while an online BCOCA reference is issued to interpret other exception codes.

Support for new printers in HPT

In response to your customer feedback, i5/OS Host Print Transform (HPT) support has added a number of new printers, with associated new printer device description object Manufacturer Type and Model (MFRTYPMDL) and associated WorkStation Customization Object (WSCST) object parameter values. Included in V5R4 are new IBM Infoprint Printers, as well as printers from Lexmark, Ricoh, and Panasonic.

Figure 15-6 on page 394 is an excerpt of the IBM System i support Web site table of supported Host Print Transform printers, showing their MFRTYPMDL and WSCST object parameter values.

Keep up to date with the HPT support by applying the associated PTFs made available throughout an i5/OS release cycle. Refer to Support Line Knowledge Base document 24474221 V5Rx PTF Listing for Host Print Transform (HPT) Printing. For the most complete list of supported printers and any associated PTFs required, you can also go directly to document 11944305 with this link:

http://www-912.ibm.com/s_dir/slkbases.NSF/7c4c0f5a4d52e3c4862569b4005d198a/a944982b85205c738625669700615ca3?OpenDocument

This is the link we used to get the excerpted list of supported printers shown in Figure 15-6 on page 394. This table shows most of the IBM Infoprint and Ricoh printers. The wavy red lines indicate that we are showing only excerpts from the Web, as the list of devices is quite long and will not fit in a single page. Printer manufacturer and supported devices listed at the Web site but not included in our figure include printers from:

- ▶ Canon
- ▶ Compaq
- ▶ Epson
- ▶ Hewlett Packard (HP)
- ▶ IBM
- ▶ Lexmark
- ▶ NEC
- ▶ Okidata
- ▶ Panasonic
- ▶ Ricoh
- ▶ Xerox

Manufacturer Type and Model (MFR TYPMDL)	Supported Printers	Workstation Customizing Object (WSCST)	Default Paper Source	V5R1 R510	V5R2 R520	V5R3 R530	V5R4 R540
*INFOPRINT12	IBM 4320 InfoPrint 12	QWPIBM4912	*LETTER	X	X	X	X
*INFOPRINT20	IBM 4320 InfoPrint 20	QWPIBM4320	*LETTER	X	X	X	X
*INFOPRINT21	IBM 4322 InfoPrint 21	QWPIBM4320	*LETTER	X	X	X	X
*INFOPRINT32	IBM 4332 InfoPrint 32	QWPIBM4332	*LETTER	X	X	X	X
*INFOPRINT40	IBM 4332 InfoPrint 40	QWPIBM4332	*LETTER	X	X	X	X
*INFOPRINT70	IBM Infoprint 70	QWPIBM2770	*LETTER	X	X	X	X
*INFOPRINT85	IBM Infoprint 85	QWPIBM2705	*LETTER	X	X	X	X
*INFOPRINT105	IBM Infoprint 105	QWPIBM2705	*LETTER	X	X	X	X
*INFOPRINT1116	IBM Infoprint 1116	QWPPIP1116	*LETTER	X	X	X	X
*INFOPRINT1120	IBM Infoprint 1120	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1125	IBM Infoprint 1125	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1130	IBM Infoprint 1130	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1140	IBM Infoprint 1140	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1145	IBM Infoprint 1145	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1220C	IBM Infoprint Color 1220	QWPPIPCOLOR	*LETTER	X	X	X	X
*INFOPRINT1222	IBM Infoprint 1222	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1226	IBM Infoprint 1226	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1228C	IBM Infoprint Color 1228	QWPPIPCOLOR	*LETTER	X	X	X	X
*INFOPRINT1312	IBM Infoprint 1312	QWPPIP1116	*LETTER	X	X	X	X
*INFOPRINT1332	IBM Infoprint 1332	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1334C	IBM Infoprint Color 1334	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1352	IBM Infoprint 1352	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1354C	IBM Infoprint Color 1354	QWPPIPCOLOR	*LETTER	X	X	X	X
*INFOPRINT1357C	IBM Infoprint Color 1357	QWPPIPCOLOR	*LETTER	X	X	X	X
*INFOPRINT1372	IBM Infoprint 1372	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1400C	IBM Infoprint Color 14xx Series	QWPPIPCOLOR	*LETTER	X	X	X	X
*INFOPRINT1410	IBM Infoprint 1410 MFP	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1412	IBM Infoprint 1412	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1422	IBM Infoprint 1422	QWPPIP1200	*LETTER	X	X	X	X
*INFOPRINT1500	IBM Infoprint 15xx Series	QWPLEXMRKC	*LETTER	X	X	X	X
*INFOPRINT1500C	IBM Infoprint Color 15xx Series Printer	QWPLEXMARKC	*LETTER		X	X	X
*INFOPRINT1600	IBM Infoprint 16xx Series Printer	QWPLEXMRKC	*LETTER		X	X	X
*INFOPRINT1600C	IBM Infoprint Color 16xx Series Printer	QWPLEXMARKC	*LETTER		X	X	X
*INFOPRINT2000	IBM InfoPrint 2000	QWPIBMIP2K	*LETTER	X	X	X	X
*INFOPRINT2085	IBM InfoPrint 2085 (2785)	QWPIBM2705	*LETTER	X	X	X	X
*INFOPRINT2105	IBM Infoprint 2105 Printer	QWPIBM2705	*LETTER	X	X	X	X
*INFOPRINT2705	IBM Infoprint 2705 Printer	QWPIBM2705	*LETTER	X	X	X	X
*INFOPRINT2706	IBM Infoprint 2105ES Printer	QWPIBM2706	*LETTER	X	X	X	X
*RICOH2035	Ricoh Aficio 2035 Printer	QWPRICOH1	*LETTER	X	X	X	X
*RICOH2045	Ricoh Aficio 2045 Printer	QWPRICOH1	*LETTER	X	X	X	X
*RICOHAP400	Ricoh Aficio AP400 Printer	QWPRICOH6	*LETTER	X	X	X	X
*RICOHAP600N	Ricoh AP600N Printer	QWPRICOH2	*LETTER	X	X	X	X
*RICOHAP900	Ricoh Aficio AP900 Printer Series	QWPRICOH11	*LETTER	X	X	X	X
*RICOHAP3200	Ricoh AP3200 Printer	QWPRICOH3	*LETTER	X	X	X	X
*RICOHAP4510	Ricoh AP4510 Printer	QWPRICOH1	*LETTER	X	X	X	X
*RICOHCL2000	Ricoh Aficio CL2000 Color Printer	QWPRICOH8	*LETTER	X	X	X	X
*RICOHCL3100	Ricoh Aficio CL3000e Color Printer	QWPRICOH8	*LETTER	X	X	X	X
*RICOHCL3100	Ricoh Aficio CL3100N Color Printer	QWPRICOH8	*LETTER	X	X	X	X
*RICOHCL4000	Ricoh Aficio CL4000 Color Series	QWPRICOH10	*LETTER	X	X	X	X
*RICOHCL5000	Ricoh CL5000 Color Printer	QWPRICOH4	*LETTER	X	X	X	X
*RICOHCL7000	Ricoh CL7000 Color Printer	QWPRICOH4	*LETTER	X	X	X	X
*RICOHCL7100	Ricoh Aficio CL7100 Color Printer	QWPRICOH4	*LETTER	X	X	X	X
*RICOHSP9100DN	Ricoh Aficio SP9100DN Printer Series	QWPRICOH12	*LETTER		X	X	X
*WORKIO_BL	Panasonic WORKiO DP-23xx Series Printer Panasonic WORKiO DP-30xx Series Printer	QWPPANDP1	*LETTER	X	X	X	X

Figure 15-6 Excerpt of supported Host Print Transform printers from Support Web site

As this book went into the publishing cycle the latest HPT PTFs were:

- ▶ V5R2: SI25364, SI25365
- ▶ V5R3: SI25430, SI25431
- ▶ V5R4: SI25450, SI25451

Note that new printer support announced during October 2006 includes:

- ▶ Ricoh Aficio SP9100DN Printer Series
- ▶ IBM Infoprint Color 150xx Series printer, IBM Infoprint Color 160xx Series Printer, IBM Infoprint 16xx Series Printer

The *INFOPRINT1600 value is mapped to the same WSCST object as the *INFOPRINT1500. The new Infoprint color models are now mapped to the QWPLEXMARKC object instead of the QWPIPCOLOR object. The enhancements this change includes are:

- ▶ EBCDIC code page 37 is now mapped to ASCII code page 850 and provides a much more accurate mapping of characters than ASCII code page 437.
- ▶ The OUTBIN attribute *DEVD is now mapped to the PCL outbin control for *device default*. This allows users to change the default outbin on the print operations panel to redirect *DEVD output.

Users can force the output to the standard bin by specifying an OUTBIN value of 100.

Unicode updates

We have upgraded the International Components for Unicode (ICU) Layout Engine from Release 2.6 to Release 3.2. This added more support for Indic scripts. There were also bi-directional layout bug fixes. More information about Unicode 3.2 is available at:

<http://www-306.ibm.com/software/globalization/icu/prevreleases.jsp>

Additional TrueType fonts in option 43

Additional TrueType fonts in option 43 are:

- ▶ Monotype Sans Duospace WT IN (indic)
- ▶ Times New Roman WT IN (indic)
- ▶ Monotype Sans Duospace WT TC TW EA (Taiwan variant of Traditional Chinese)

	Traditional text layout	ICU Complex text layout
Bengali	বর্তমান অঙ্কুর বভিনি	বর্তমানে অঙ্কুর বিভিন্ন
Devanagari	अर्जुन विषाद योग	अर्जुन विषाद योग
Gujarati	ભાષામાં સોફ્ટવેરની ઉભૂપ	ભાષામાં સોફ્ટવેરની ઉભૂપ
Gurmukhi	ਸੁਰੂ ਕੀਤੀ ਹੈ। ਅੰਤ ਵਿਚ	ਸੁਰੂ ਕੀਤੀ ਹੈ। ਅੰਤ ਵਿਚ
Kannada	ಪೌರತೌ ದೌಪರದ ಎರಡು	ಪೃತಿ ದ್ವಾರದ ಎರಡು
Malayalam	എല്ല്യാ സ്പതനക്രവ്വം	എല്ലാ സ്വതന്ത്രവ്വം
Tamil	இராந்தா விலகும்	இருந்து விலகும்
Telugu	అత్తయింతపర్వతాపాత్తర	అత్యంతప్రీతిపాత్త

Figure 15-7 Complex text layout examples

15.6.5 DDS printer file enhancements

In this section we discuss DDS printer file enhancements.

Relative position keyword

Use this file level keyword to indicate that when the +n form of positioning a field is used, the fields subsequent to the first field in the line are positioned relative to the end of the previous field. Without the RELPOS keyword, the column number calculated at compile time is generally relative to the beginning of the line.

This file level keyword has no parameters.

The +n positioning means that if you specify a location of a field in a record and the field is not ignored, you can specify the location of subsequent fields within that record by leaving the line number blank and specifying a plus value (+n) for position entry 42 through 44. The plus value indicates the number of spaces to be left between the end of the previous field and the beginning of the field you are defining. The plus value must be in the range of 0 through 99. If you specify a plus value, the line number entry must be blank. As soon as the positions are calculated, the real values are stored and treated as if they were hard-coded. This is a compile time function.

Specify DEVTYPE (*AFPDS) on the CRTPRTF command when RELPOS is specified at the file level. If DEVTYPE is changed to anything other than *AFPDS, the keyword is ignored and a warning message is issued at print time.

Option indicators are not valid for this keyword.

The Print Services Facility (PSF) feature is required for printing DEVTYPE (*AFPDS) files on an IPDS printer. If the printer file FONT parameter is *DEV, or the printer file has a FNTCHRSET or CDEFNT specified at the file level, the file-level font width is not known.

Example 15-1 shows how to specify the RELPOS keyword.

Example 15-1 How to specify the RELPOS keyword

```
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8  
  
A*  
A                RELPOS  
A          R REC1          FONT(1051)  
A*                (1051 specifies Sonoran Serif)  
A          FLD1          5A    1 10  
A          FLD2          5A    +2  
A*
```

FLD2 is placed two spaces to the right of the end of FLD1. Depending on the width of the characters used in FLD1, the start of FLD2 might vary. Without the RELPOS keyword, FLD2 is placed in column 17 of the line.

Position keyword

Use this field-level keyword to define the location of a named field on the page. The format of the keyword is:

```
POSITION(position-down | &position-down-field  
position-across | &position-across-field)
```

The position-down parameter is required and defines the vertical starting point of the field relative to the margins specified on the FRONTMGN or BACKMGN parameter on the CRTPRTF command. Valid values are 0 to 57.790 cm (0 to 22.750 in.).

The position-across parameter is required and defines the horizontal starting point of the field relative to the margins specified on the FRONTMGN or BACKMGN parameter on the CRTPRTF command. Valid values are 0 to 57.790 cm (0 to 22.750 in.).

You can specify the position-down and position-across parameters as constants, program-to-system fields, or a combination of both, as shown in the following examples:

- ▶ POSITION(3.56 6.24)
- ▶ POSITION(&field1 9.625)
- ▶ POSITION(0.5 &field2)
- ▶ POSITION(&field3 &field4)

Field1, field2, field3, and field4 are the names of program-to-system fields. The fields must exist in the same record format as the POSITION keyword and be defined as having length 5 with 3 decimal positions, data type S (zoned decimal), and usage P (program-to-system).

The UOM parameter on the CRTPRTF command determines the units of measure for the position-down and position-across parameter values. If the value specified for a parameter is outside the valid range, it is flagged when the spooled file is created.

An error message is issued at print time if the field does not fit on the page. An error message is issued at create time if line and position values, columns 39 through 44, are also specified.

Because the POSITION keyword allows a field to be positioned anywhere on the page, a new page is not generated by the use of the position keyword. The ENDPAGE keyword should be used to end the current page and proceed to the next page.

If the POSITION keyword is specified for a field, all fields in the record format must also have the POSITION keyword specified. Location entries in positions 39 through 44 are not allowed. You can specify constant fields where the POSITION keyword is also specified. You just need to use a special name (*NONE).

Specify DEVTYPE(*AFPDS) on the CRTPRTF command when POSITION is specified in the file. If DEVTYPE is changed to anything other than *AFPDS, the keyword is ignored and a warning message is issued at print time.

You cannot specify POSITION with the following keywords:

- ▶ SPACEA
- ▶ SPACEB
- ▶ SKIPA
- ▶ SKIPB

Option indicators are valid for this keyword.

Example 15-2 shows how to specify the POSITION keyword.

Example 15-2 Specify the POSITION keyword

		...	1	+	...	2	+	...	3	+	...	4	+	...	5	+	...	6	+	...	7	+	...	8	
	A*																								
	A		R			REC1																			
	A					FLD1			6S		2				POSITION(2.0										
	A*																								
	A					FLD2			42A						POSITION(&FLD2A										
	A					FLD2A			5S		3P														
	A					FLD2B			5S		3P														
	A*																								

The UOM parameter on the Create Printer File (CRTPRTF) command determines the units of measure for the parameter values.

In REC1, FLD1 prints 2.0 units down and 1.983 units across from the margins specified on the FRONTMGN or BACKMGN parameter on the CRTPRTF command.

The application program determines the position of FLD2 by assigning values to program-to-system variables FLD2A and FLD2B.

15.6.6 CHGIPLA modification aids spooled file retention

On previous releases, a damaged job table (sometimes referred to as a Work Control Block Table) would lead to a Clear/Clear/Clear IPL. Customers had no control over this and typically did not realize that all of their jobs and spooled files would be lost.

Starting with V5R4, in order to preserve spooled files, the behavior of the system was modified to retain all spooled files during a Clear/Clear/Clear IPL. The SPLFRCY IPL attribute was created to handle this function.

The shipped IPL attribute for SPLFRCY is *DETACH.

If a Clear/Clear/Clear IPL is performed with SPLFRCY(*DETACH), all spooled files on the system will be detached from their jobs, but will remain in their OUTQS. The result of this would be that users and applications will no longer be able to issue Job commands against spooled files, as spooled files will be detached from jobs. (For example, WRKJOB will show no spooled files.)

However, a user could use the new JOB parameter on WRKSPLF to find spooled files for certain jobs.

Users wishing to remove all spooled files from the system, due to damaged spooling DB files or another situation, can change the IPLA SPLFRCY parm to *REMOVE before IPL. After the IPL, we recommend that you change it back to *DETACH.

15.7 Printing hardware update (selected devices)

In this section we discuss printing hardware updates for selected drives.

15.7.1 RFID: Infoprint 6700

Radio frequency identification (RFID) is a technology that used the radio frequency (RF) spectrum to communicate. RFID is coming into increasing use in the industry as an alternative to bar code or in combination with barcode. The advantage of RFID is that it does not require direct contact or line-of-sight scanning.

An RFID system consists of three components: an antenna and transceiver (often combined into one reader) and a transponder (the tag). The tag is usually a label with the RFID antenna and memory built in. RFID acceptance has accelerated in the last year for three key reasons:

- ▶ Key supply chain companies (that is, Wal-Mart) are requiring it.
- ▶ The cost per tag has come down to a few cents or less.
- ▶ A consolidated RFID standard is finally emerging.

In industrial environments, reliable output management and radio frequency identification (RFID) capability are becoming critical to the supply chain. This is an important application space for IBM, with \$250M already invested in RFID solutions.

To keep your products and information flowing, the new IBM Infoprint 6700 printers support thermal bar code printing and RFID encoding in the same printer footprint. The tag is electronically encoded and printed with human readable characters and bar-codes by the 6700 based on instructions from the application.

By integrating the Infoprint 6700 printers with IBM RFID and middleware architecture, you can implement a complete RFID solution. IBM is the only vendor that can provide a complete end-to-end output management implementation. In addition, with impending changes in RFID technology, the Infoprint 6700 provides an upgrade path to protect the RFID investment.

Relative to i5/OS enablement of the Infoprint 6700, keep in mind that bar code printing and RFID encoding are two separate steps, at two different stations in the printer, utilizing different data streams. To print bar codes, the printer supports either IPDS or thermal data streams. IPDS would use standard DDS coding for development. For RFID, the PGL, ZGL, and STGL data streams are supported.

Note: XGL™ is a text file format containing OpenGL 3D information written in XML. ZGL is a compressed version of XGL file. Both of these formats are used for communication of 3D designs over the Internet, for example, by Autodesk Inventor, Mechanical Desktop, and Autodesk Streamline™.

StereoGL (STGL) provides a uniform interface for stereo 3D displays under OpenGL.

These can be set up as ASCII escape sequences in the application or generated with bar code-RFID design tools such as MarkMagic and Barcode/400.

Mapping Suite supports both RFID and thermal data stream design and production. Thermal data streams supported include ZPL (Zebra), Datamax, IGP, TEC, and Intermec.

15.7.2 MFP devices

IBM Multi-function Printers (MFPs) combine printer, scanner, fax, and copier in a single footprint. They have been around for some time, highly focused on the small-to-medium sized business (SMB) and home office markets, but they have grown up and now have the capability to provide comprehensive departmental support in mid-market and large enterprises.

The key here is device consolidation — replacing rafts of inkjet printers, faxes, and other machines with a single high-performance machine.

The IBM set of MFP-capable devices has expanded to the devices listed in Table 15-2. There are speeds up to 55 ppm, monochrome and color options, and available IPDS support.

Table 15-2 IBM MFP devices

Printer	Machine number	Speed (impressions per minute)	Monthly volume rating	IPDS capable
Infoprint 1540 MFP	4543	35	150,000	Y

Printer	Machine number	Speed (impressions per minute)	Monthly volume rating	IPDS capable
Infoprint 1560 MFP	4543	45	200,000	Y
Infoprint 1570 MFP	4542	50	225,000	Y
Infoprint 1572MFP	4897	50	250,000	Y
Infoprint 1580 MFP	4543	55	300,000	Y
Infoprint 1650 MFP	4544	45	200,000	N
Infoprint 1664 Color MFP	4898	25	120,000	Y

Most customers can use some guidance in implementation of this output convergence technology. Understanding what equipment they have and what their total device and operational costs really are can be very complex. Simply tracking usage and supply costs can be a difficult task. This is often complicated because of the lack of centralized responsibility for various output devices within an organization. Lab Services can provide device consolidation studies that pave the way for efficient implementation of MFPs in an enterprise.

For more information about Lab Services contact Cathy McGraw at cmcgraw@us.ibm.com.

15.7.3 Infoprint 15xx and 16xx workgroup printers

IBM provides a complete line of workgroup laser printers. The current line is composed of the 15xx and 16xx printer series. Table 15-3 shows workgroup printers, speeds, monthly volume ratings, and availability of IPDS. These printers are designed, built, tested, and supported to work with System i. The PCL print mode is supported by Host Print Transform. IPDS is managed by Print Services Facility (PSF). PDF printing is enabled by the Infoprint Server.

Table 15-3 IBM Workgroup printers

Printer	Machine number	Speed (impressions per minute)	Monthly volume rating	IPDS capable
Infoprint 1612	4546	30	15,000	N
Infoprint 1622	4546	32	65,000	N
Infoprint 1532	4536	35	200,000	Y
Infoprint 1552	4537	45	225,000	Y
Infoprint 1572	4538	50	250,000	Y
Infoprint 1585	4539	50	300,000	Y
Infoprint Color 1634	4931	24	65,000	N

Printer	Machine number	Speed (impressions per minute)	Monthly volume rating	IPDS capable
Infoprint Color 1567	4935	32	200,00	Y
Infoprint Color 1654	4929	25	100,000	Y
Infoprint Color 1664	4930	25	120,000	Y

For more information about the IBM Infoprint printers see:

<http://www.printers.ibm.com>



i5/OS-based performance update

In this chapter we provide an overview of i5/OS V5R4 performance characteristics and V5R4 level enhancements to some of the performance tools used to manage and size workloads running under i5/OS.

To get the latest i5/OS-based performance information refer to the Web site:

<http://www.ibm.com/eserver/series/perfmgmt>

16.1 i5/OS V5R4 performance updates

The i5/OS V5R4 implementation included specific performance enhancements. Other performance characteristics are based upon new processor models or new I/O device support.

The most complete set of System i and i5/OS performance information can be accessed from the performance management Web site:

<http://www.ibm.com/eserver/iserries/perfmgmt>

One of the primary performance documents accessible via the Resource Library link is the V5R4 Performance Capabilities Reference manual.

16.2 CPW and MCU ratings

The Commercial Processing workload (CPW) is the internal System i (iSeries) benchmark traditionally used as a relative measure of performance among the System i family of processors.

Though not commonly included in marketing documentation, there is also another internal performance rating comparison called (the number of) Mail and Calendaring Users. The Mail and Calendaring Users (MCUs) is an authorized mixed workload benchmark from the Lotus Corporation. It can be in the System i context as a more CPU intensive than CPW workload for performance rating of System i models.

Compute Intensive Workload (CIW) is another internal System i benchmark that is significantly more CPU intensive than either the CPW or the MCU benchmark. However, this CIW benchmark was discontinued as a basis for rating System i models with V5R2 when the newest 8xx models were announced.

For these benchmarks, all measurements were taken on maximum configurations. So, for CPW and MCU ratings, an actual customer application and hardware configuration may not yield similar performance rating ratios among System i models.

If you have a CPU-intensive application environment consider using the relative CIW rating if available for the older technology systems. For the new technology systems, consider a 20% higher CPW rating for CPU-intensive workloads.

Note that IBM System i5 models are frequently also tested and validated for worldwide benchmark applications such as in JAVA-based SPECjbb2000, SAP environments, Notesbenchmark environments, and Oracle JD Edwards, and other Independent Software Vendor (ISV) workloads. Since the benchmark rankings for these benchmarks often change every few months and are commonly achieved with very special techniques not typically used by many application developers, we do not include any of these benchmark ratings in this book. If you are interested in the rankings with these kinds of benchmarks, consider reviewing the information at these Web sites:

- ▶ SAP:

http://www.sap.com/solutions/benchmark/BW2_results.htm

- ▶ Java-based SPECjbb2000 Benchmark Results (32-bit JVM and 64-bit JVM results):

<http://www.spec.org>

Other documented benchmark results include:

- ▶ Intentia One Box - Movex Workplace (WebSphere), Net Extension (WebSphere), and Movex Business Engine (Java):

<http://www.ibm.com/servers/enable/site/technical/pdfs/a49a.pdf>

- ▶ mySAP™ ERP R/3 Enterprise (C/C++ application in PASE):

<http://www.ibm.com/servers/enable/site/technical/pdfs/a4ae.pdf>

- ▶ Oracle JDE Enterprise One (WebSphere, ILE C/C++, SQL):

<http://www.ibm.com/servers/enable/site/technical/pdfs/a48e.pdf>

The chart in Figure 16-1 shows the CPW and MCU ratings of each of the System i5 models as announced throughout 2006. Notes about some of the information items follow the chart.

Model	Edition Feature ²	Accelerator Feature	Chip Speed MHz	L2/L3 cache per CPU ⁽¹⁾	CPU Range	Processor CPW	5250 OLTP CPW	MCU
9406-595	7583 ⁽⁵⁾	NA	1900	1.9/36MB	32 - 64 ⁽⁸⁾	92000-184000	Per Processor	213K ⁽⁷⁾ - 405K ⁽⁷⁾
9406-595	7487	NA	1900	1.9/36MB	32 - 64 ⁽⁸⁾	92000-184000	Per Processor	213K ⁽⁷⁾ - 405K ⁽⁷⁾
9406-595	7486	NA	1900	1.9/36MB	32 - 64 ⁽⁸⁾	92000-184000	0	213K ⁽⁷⁾ - 405K ⁽⁷⁾
9406-595	7581 ⁽⁵⁾	NA	1900	1.9/36MB	16 - 32	51000-92000	Per Processor	115000 - 213K ⁽⁷⁾
9406-595	7483	NA	1900	1.9/36MB	16 - 32	51000-92000	Per Processor	115000 - 213K ⁽⁷⁾
9406-595	7482	NA	1900	1.9/36MB	16 - 32	51000-92000	0	115000 - 213K ⁽⁷⁾
9406-595	7590 ⁽⁴⁾	NA	1900	1.9/36MB	4 - 32	13600-92000	Per Processor	31500 - 213K ⁽⁷⁾
9406-595	7580 ⁽⁵⁾	NA	1900	1.9/36MB	8 - 16	26700-50500	Per Processor	60500 - 114000
9406-595	7481	NA	1900	1.9/36MB	8 - 16	26700-50500	Per Processor	60500 - 114000
9406-595	7480	NA	1900	1.9/36MB	8 - 16	26700-50500	0	60500 - 114000
9406-570	7760 ⁽⁴⁾	NA	2200	1.9/36MB	2 - 16	8100-58500	Per Processor	18200 - 130000
9406-570	7765 ⁽⁵⁾	NA	2200	1.9/36MB	8 - 16	31100-58500	Per Processor	67500 - 130000
9406-570	7749	NA	2200	1.9/36MB	8 - 16	31100-58500	Per Processor	67500 - 130000
9406-570	7759	NA	2200	1.9/36MB	8 - 16	31100-58500	0	67500 - 130000
9406-570	7764 ⁽⁵⁾	NA	2200	1.9/36MB	4 - 8	16700-31100	Per Processor	35500 - 67500
9406-570	7748	NA	2200	1.9/36MB	4 - 8	16700-31100	Per Processor	35500 - 67500
9406-570	7758	NA	2200	1.9/36MB	4 - 8	16700-31100	0	35500 - 67500
9406-570	7763 ⁽⁵⁾	NA	2200	1.9/36MB	2 - 4	8400-16000	Per Processor	18200 - 34500
9406-570	7747	NA	2200	1.9/36MB	2 - 4	8400-16000	Per Processor	18200 - 34500
9406-570	7757	NA	2200	1.9/36MB	2 - 4	8400-16000	0	18200 - 34500
9406-550	7551 ⁽⁵⁾	NA	1900	1.9/36MB	1 - 4	3800-14000	Per Processor	8200 - 30000
9406-550	7629 ⁽⁶⁾	NA	1900	1.9/36MB	1 - 4	3800-14000	0	8200 - 30000
9406-550	7155	NA	1900	1.9/36MB	1 - 4	3800-14000	Per Processor	8200 - 30000
9406-550	7154	NA	1900	1.9/36MB	1 - 4	3800-14000	0	8200 - 30000
9406-520	7375 ⁽⁵⁾	NA	1900	1.9/36MB	1 - 2	3800-7100	3800-7100	8200 - 15600
9406-520	7736	NA	1900	1.9/36MB	1 - 2	3800-7100	3800-7100	8200 - 15600
9406-520	7785	NA	1900	1.9/36MB	1 - 2	3800-7100	0	8200 - 15600
9406-520	7784	NA	1900	1.9/36MB	1	3800	0	8200
9406-520	7691 ⁽¹⁰⁾	NA	1900	1.9/36MB	1	3800	0	8200
9406-520	7374 ⁽⁵⁾	NA	1900	1.9/36MB	1 ⁽³⁾	2800	2800	6100
9406-520	7735	NA	1900	1.9/36MB	1 ⁽³⁾	2800	2800	6100
9406-520	7373 ⁽⁵⁾	NA	1900	1.9/36MB	1 ⁽³⁾	1200	1200	2600
9406-520	7734	NA	1900	1.9/36MB	1 ⁽³⁾	1200	1200	2600
Value								
9406-520	7352	7357	1900	1.9/36MB	1 ⁽³⁾	1200-3800 ⁹	60	2600 - 8200
9406-520	7350	7355	1900	1.9MB/NA	1 ⁽³⁾	600-3100 ⁹	30	NR - 6600
Express								
9405-520	7152	NA	1900	1.9/36MB	1	3800	60	8200
9405-520	7144	NA	1900	1.9/36MB	1	3800	60	8200
9405-520	7143	7354	1900	1.9/36MB	1 ⁽³⁾	1200-3800 ⁹	60	2600 - 8200 ⁽⁹⁾
9405-520	7148	7687	1900	1.9/36MB	1 ⁽³⁾	1200-3800 ⁹	60	2600 - 8200 ⁽⁹⁾

Figure 16-1 CPW and MCU ratings for System i5 models announced January 2006

Notes:

1. These models have L2 and L3 cache local (physically on the same node) to each pair (2) of processors (*cores*).
2. This is the Edition Feature for the model. This is the feature displayed when you display the system value QPRCFEAT.
3. CPU Range: Partial processor models, offering multiple price/performance points for the entry market.
4. Capacity Backup model.
5. High Availability.
6. Domino Edition.
7. The MCU rating is a projected value.
8. The 64-way is measured as two 32-way partitions since i5/OS does not support a 64-way partition.
9. These models are accelerator models. The base CPW or MCU value is the capacity with the default processor feature. The max CPW or MCU value is the capacity when purchasing the accelerator processor feature.
10. Collaboration Edition (announced May 9, 2006).

NR = Not Recommended: the 600 CPW processor offering is not recommended for Domino.

The chart in Figure 16-2 shows the CPW and MCU ratings of each of the System i5 models, as announced prior to January of 2006. Notes for some of the information items follow the chart.

Model	Chip Speed MHz	L2 cache per CPU ⁽¹⁾	L3 cache per CPU ⁽²⁾	CPU Range	Processor CPW	5250 OLTP CPW	MCU
595-0952 (7485)	1650	1.9 MB	36 MB	32 - 64 ^(*)	86000-165000	12000-165000	196000 ⁽⁷⁾ -375000 ⁽⁷⁾
595-0952 (7484)	1650	1.9 MB	36 MB	32 - 64 ^(*)	86000-165000	0	196000 ⁽⁷⁾ -375000 ⁽⁷⁾
595-0947 (7499)	1650	1.9 MB	36 MB	16 - 32	46000-85000	12000-85000	105000 -194000 ⁽⁷⁾
595-0947 (7498)	1650	1.9 MB	36 MB	16 - 32	46000-85000	0	105000 -194000 ⁽⁷⁾
595-0946 (7497)	1650	1.9 MB	36 MB	8 - 16	24500-45500	12000-45500	54000-104000
595-0946 (7496)	1650	1.9 MB	36 MB	8 - 16	24500-45500	0	54000-104000
570-0926 (7476)	1650	1.9 MB	36 MB	13 - 16	36300-44700	12,000-44,700	83600-102000
570-0926 (7475)	1650	1.9 MB	36 MB	13 - 16	36300-44700	0	83600-102000
570-0926 (7563) ⁵	1650	1.9 MB	36 MB	13 - 16	36300-44700	12000-44,700	83600-102000
570-0928 (7570) ⁴	1650	1.9 MB	36 MB	2 - 16	6350-44700	6,350-44,700	14100-102000
570-0928 (7474)	1650	1.9 MB	36 MB	9 - 12	25500-33400	12,000-33,400	57300-77000
570-0924 (7473)	1650	1.9 MB	36 MB	9 - 12	25500-33400	0	57300-77000
570-0924 (7562) ⁵	1650	1.9 MB	36 MB	9 - 12	25500-33400	12000-44,700	57300-77000
570-0922 (7472)	1650	1.9 MB	36 MB	5 - 8	15200-23500	12,000-23,500	33600-52500
570-0922 (7471)	1650	1.9 MB	36 MB	5 - 8	15200-23500	0	33600-52500
570-0922 (7561) ⁵	1650	1.9 MB	36 MB	5 - 8	15200-23500	12,000-23,500	33600-52500
570-0921 (7495)	1650	1.9 MB	36 MB	2 - 4	6350-12000	12000	14100-26600
570-0921 (7494)	1650	1.9 MB	36 MB	2 - 4	6350-12000	0	14100-26600
570-0921 (7560) ⁵	1650	1.9 MB	36 MB	2 - 4	6350-12000	12000	14100-26600
570-0930 (7491)	1650	1.9 MB	36 MB	1 - 2	3300-6000	6000	7300-13300
570-0930 (7490)	1650	1.9 MB	36 MB	1 - 2	3300-6000	0	7300-13300
570-0930 (7559) ⁵	1650	1.9 MB	36 MB	1 - 2	3300-6000	6,000	7300-13300
570-0920 (7470)	1650	1.9 MB	36 MB	2 - 4	6350-12000	Max	14100-26600
570-0920 (7469)	1650	1.9 MB	36 MB	2 - 4	6350-12000	0	14100-26600
570-0919 (7489)	1650	1.9 MB	36 MB	1 - 2	3300-6000	Max	7300-13300
570-0919 (7488)	1650	1.9 MB	36 MB	1 - 2	3300-6000	0	7300-13300
550-0915 (7530) ⁶	1650	1.9 MB	36 MB	2 - 4	6350-12000	0	14,100-26600
550-0915 (7463)	1650	1.9 MB	36 MB	1 - 4	3300-12000	3,300-12,000	7300-26600
550-0915 (7462)	1650	1.9 MB	36 MB	1 - 4	3300-12000	0	7300-26600
550-0915 (7558)	1650	1.9 MB	36 MB	1 - 4	3300-12000	3,300-12,000	7300-26600
520-0905 (7457)	1650	1.9 MB	36 MB	2	6000	3,300-6000	13300
520-0905 (7456)	1650	1.9 MB	36 MB	2	6000	0	13300
520-0905 (7555) ⁵	1650	1.9 MB	36 MB	2	6000	3,300-6,000	13300
520-0904 (7455)	1650	1.9 MB	36 MB	1	3300	3,300	7300
520-0904 (7454)	1650	1.9 MB	36 MB	1	3300	0	7300
520-0904 (7554) ⁵	1650	1.9 MB	36 MB	1	3300	3,300	7300
520-0903 (7453)	1500	1.9 MB	NA	1	2400	2400	5500
520-0912 (7397)	1500	1.9 MB	NA	1	2400	60	5500
520-0912 (7395)	1500	1.9 MB	NA	1	2400	60	5500
520-0903 (7452)	1500	1.9 MB	NA	1	2400	0	5500
520-0903 (7553) ⁵	1500	1.9 MB	NA	1	2400	2400	5500
520-0902 (7459)	1500	1.9 MB	NA	1 ⁽³⁾	1000	1000	2300
520-0902 (7458)	1500	1.9 MB	NA	1 ⁽³⁾	1000	0	2300
520-0902 (7552) ⁵	1500	1.9 MB	NA	1 ⁽³⁾	1000	1000	2300
520-0901 (7451)	1500	1.9 MB	NA	1 ⁽³⁾	1000	60	2300
520-0900 (7450)	1500	1.9 MB	NA	1 ⁽³⁾	500	30	NA recommended

Figure 16-2 CPW and MCU ratings for System i5 models announced prior to January 2006

Notes:

- 1.9 MB - These models share L2 cache between two processors.
- 36 MB - These models share L3 cache between two processors.
- CPU Range - Partial processor models, offering multiple price/performance points for the entry market.
- Capacity Backup model.
- High Availability model.

6. Domino Edition.
7. The MCU rating is a projected value.
8. The 64-way is measured as two 32-way partitions since i5/OS does not support a 64-way partition. IBM stopped publishing CIW ratings for iSeries after V5R2. We recommend that the eServer Workload Estimator be used for sizing guidance, available at:

<http://www.ibm.com/eserver/iseries/support/estimator>

We have provided another chart for comparing the CPW and MCU ranges between the pre-January 2006 System i5 models and the post-January 2006 System i5 models.

The chart in Figure 16-3 shows an improved CPW and MCU range of the models with the new processor speed of 1.9 GHz. In the chart you see that the CPW ratings have almost doubled, and the MCU has significantly improved for some models.

Model	Processor GHz	Number of Processors Range	CPW Range	MCU Range
595 January 2006	1.9	8 - 64	26,700 – 184,000	60,500 - 405,000
595 Before Jan 2006	1.65	8 - 64	24,500 – 165,000	54,000 – 375,000
570 January 2006	2.2	2 - 16	8100 – 58,500	18,200 – 130,000
570 Before Jan 2006	1.65	1 - 16	3300 – 44,700	7300 – 102,000
550 January 2006	1.9	1 - 4	3800-14000	8200 - 30000
550 Before Jan 2006	1.65	1-4	3300 – 12000	7300 - 26600
520 January 2006	1.9	1 - 2	(1200 Express) 3800-7100	2600 8200 - 15600
520 Before Jan 2006	1.65	1-2	(1000 Express) 3,300-12,000	2300 7300 - 13300

Figure 16-3 CPW and MCU ranges comparison

Accelerator on System i5-520

With the activation of the IBM Accelerator (available on demand) for the System i5-520 Value and Express Editions, significant processing power is gained. A chart on this performance achievement on the i520 is shown in Figure 2-3 on page 10.

16.2.1 Database performance update

In i5/OS V5R4 there are several new performance enhancements to DB2 UDB for iSeries. More queries will now be processed by the new (in V5R2) query optimizer. This should result in improved performance because the SQE optimizer has greater potential to choose a more efficient access plan, as well as more efficient query processing than the CQE optimizer.

Also implemented in i5/OS V5R4 are enhancements to materialized query tables (MQTs) and improved efficiency of partitioned table processing support, which were both new additions in i5/OS V5R3.

DB2 UDB for iSeries now supports the use of Recursive Common Table Expressions (RCTEs), which may provide for more elegant and better performing implementations of recursive processing than the use of stored procedures.

SQE now can choose to use a temporary index for data access.

The query dispatcher controls whether an SQL query will be routed to SQE or to CQE. In i5/OS V5R4 more queries are implemented in SQE. SQL queries with the following attributes, which were routed to CQE in previous releases, may now be routed to SQE in i5/OS V5R4:

- ▶ Sensitive cursor
- ▶ Like/substring predicates
- ▶ ALWCPYDTA(*NO)
- ▶ LOB columns

The following query attributes continue to be processed by the CQE:

- ▶ NLSS/CCSID translation between columns
- ▶ References to DDS logical files
- ▶ Tables with select/omit logical files over them
- ▶ DB2 Multisystem

In general, queries with like and substring predicates, which are newly routed to SQE, see substantial performance improvements in i5/OS V5R4. Figure 16-4 is a chart taken from the 2006 V5R4 Performance Capabilities Reference manual. It shows Rochester laboratory performance test results that demonstrate a wide range performance improvements. That chart also shows the performance ratio of the V5R4 query runtime to the V5R3 query runtime for a set of queries that have like and substring predicates.




Figure 16-4 Query performance: like and substring predicates

Each vertical line (not shown) in the chart represents both a V5R3 and a V5R4 test result for the same test query. Due to the large number of queries summarized in each graph, you see

a shaded curve summarizing the results rather than distinct bars for each V5R4-V5R4 query comparison.

When the ratio is 1.0 the query takes the same amount of time to execute on both releases. A fraction of 0.5 (that is, less than 1) would mean that the query runs half as long — twice as fast on V5R4 as V5R3 (for example, on V5R4 for 1 minute, while V5R3 takes 2 minutes). Similarly, a value of 2.0 would mean that the query runs twice as long on V5R4 as V5R3 (for example, on V5R4 for 2 minutes, and on V5R3 for 1 minute).

The queries and data vary greatly in this set of queries, including a wide range of like and substring predicates and various file sizes. The large performance gains for the like and substring queries are seen, with a large percentage of the queries seeing up to a 10X reduction in query runtime.

In another performance chart taken from the August 2006 V5R4 Performance Capabilities Reference manual shown in Figure 16-5, the queries with references to LOB columns, routed to SQE, in general, see substantial performance improvements in i5/OS V5R4.

LOB support on V5R4 performance is much improved since these queries now route to SQE instead of CQE on V5R3. Very few queries are degraded (with ratios above 1.0). The LOBS results show a few more queries that are degraded (above 1.0 ratio) with V5R4 over V5R3, but in general queries are improved slightly or up to 2–5x faster on V5R4.

A large percentage of the queries see up to a 5X reduction in query runtime.




Figure 16-5 LOB performance comparison

16.2.2 Journaling

To reduce the abnormal IPL duration consider using CHGJRN Journal Recovery Count (JRNRCYCNT).

The maximum number of objects associated with a single journal can be set to either 250,000 or 10,000,000 using the Journal Object Limit (JRNOBJLMT) parameter on the Create Journal (CRTJRN) or Change Journal (CHGJRN). The value *MAX10M can only be specified for the Journal Object Limit (JRNOBJLMT) parameter if the Receiver Size Option (RCVSIZOPT) parameter has one of the receiver maximum values specified, or if RCVSIZOPT is *SYSDFT. *MAX250K is the default value. Once you have chosen the *MAX10M option, you cannot switch back to the *MAX250K value. Increasing the quantity of objects associated with a single journal may increase your IPL time, independent ASP vary on time, or disaster recovery time.

As a general rule, if the number of actively changing objects is likely to be greater than 5,000, consider journaling some of these objects to a separate journal. The larger the number of actively changing objects for a given journal at system termination, the longer it will take to recover the journal at IPL or vary on of an independent ASP. If the new JRNOBJLMT(*MAX10M) is specified for your journal, you cannot restore the journal to a release prior to V5R4.

16.2.3 Reorganize table (physical file member)

The reorganization of a physical file member (table) function removes deleted records from (compresses) one member of a physical file in the database, and it optionally reorganizes that member. If a keyed file is identified in the Key file (KEYFILE) parameter of a reorganize file function, the system reorganizes the member by changing the physical sequence of the records in storage to either match the keyed sequence of the physical file member's access path, or to match the access path of a logical file member that is defined over the physical file.

Reorganization can decrease file processing time for applications reading sequentially through a keyed physical file or through a keyed logical file.

This function has been available for over 10 years. With i5/OS V5R3 there were important functional and performance improvements made. This is one of the important V5R3 capabilities we include in this book.

Information in this book about this function is a combination of information from the following sources:

- ▶ iSeries Information Center. Search with *ALWCANCEL(*YES)*.
- ▶ V5R3 Database presentation in the ITSO V5R3 Technical Overview presentation set located at the IBM System i technical support Web site:
<http://www.ibm.com/systems/support/i/>
Select the **Technical Library** link and then **System i Technical Overviews** → **V5R3**.
- ▶ An article written by IBMers Amy Anderson and Fernando Echeveste appearing in the IBM Systems Magazine, i5 Edition, published by MSP Techmedia. The article is titled "Revisiting Table Reorganization to Reduce Paging Rates." The following is a direct link to this article:
<http://www.ibmssystemsmag.com/ME2/Audiences/dirmod.asp?sid=&nm=&type=Publishing&mod=Publications%3A%3AArticle&mid=8F3A7027421841978F18BE895F87F791&AudID=1E8FEE745A284521B6CFB3FD70B49099&tier=4&id=F2BEF8D29DFF46F89DBCADAC04545013#Fig2>

Recall that you can invoke the reorganize table (physical file member) through either:

- ▶ The Reorganize Physical File Member (RGZPFM) command.
- ▶ The iSeries Navigator Database interface. Select a schema (library), open the library, open the tables, and right-click a table. Select **Data** → **Reorganize**.

The new reorganize capabilities are activated with new parameter `ALWCANCEL(*YES)`, as follows:

- ▶ `ALWCANCEL(*NO)`: This is the default and uses the traditional reorganize algorithm and rules for processing. A full copy of the data might be made, so you need up to two times the amount of space. This option cannot be canceled (suspended) and cannot fully run in parallel. It requires exclusive use of the file.
- ▶ `ALWCANCEL(*YES)`: This invokes the new algorithm and rules for processing. The data rows are moved within the file so that a full copy of the data is not required. The file must be journaled, however, so storage is necessary for the journal entries.

You can use the journal receiver threshold to minimize the amount of storage used in a specific journal receiver. This option can be canceled (suspended) and restarted.

If you cancel the reorganize function in progress (`ALWCANCEL(*YES)` specified), the reorganize will be partially complete. Subsequently, another reorganize with the same parameters may be able to continue from where the last (suspended) reorganize ended. If the number of changes that have occurred since the reorganize was canceled (suspended) is too large, the reorganize may be restarted rather than continued.

The reorganize can run in parallel if the additional charge DB2 for i5/OS Symmetric Multiprocessing (SMP) feature is installed.

When you specify that the Reorganize Physical File Member (RGZPFM) command can be canceled or suspended, you are telling DB2 for i5/OS to use the SMP-capable algorithm. DB2 will use the SMP setting assigned to the user who issued the RGZPFM command. SMP lets the system use more than one processor and varying amounts of memory for a single job.

When the SMP system value is set to `*MAX`, jobs running under the user who issued the command will try to use all of the processors and all of the memory allocated to the partition. If SMP is set to `*OPTIMIZE`, the command will try to take its *fair share* of partition resources. The *fair share* is the amount of system resources divided by the maximum number of active users in the job's main storage pool. When SMP is set to `*NONE`, the command will use only one processor, regardless of how many are in the partition.

Verify that SMP is installed by looking for it in the list of licensed programs installed in your partition (system), either through iSeries* Navigator or through the Display Licensed Program (DSPLICPGM) command. You can change its setting through the QQRVDEGREE system value, the CHGQRYA CL command, or Change Query Attributes from the Series Navigator interface.

As with any use of SMP parallelism, you can potentially use all of the system resources to reorganize a single table. But as with any parallelization capability, it does not mean that you always should. If you are running the reorganization as a background task, you would not want to use all of the system resources. You can explicitly choose the number of processors to use.

If you are reorganizing tables concurrent with batch applications, `*OPTIMIZE` might make a lot of sense, since batch operations usually have a relatively low number of concurrent users. Most daytime operations have hundreds or even thousands of concurrent users, so turning SMP off completely by setting the system value to `*NONE` makes the most sense. Even when SMP is turned off, you can still reorganize tables in the background.

This option requires exclusive use for only a few seconds after the reorganize is complete to return storage to the system. If the exclusive lock cannot be acquired, a warning message is sent to the job log indicating that space could not be recovered. To recover the space, you can issue the reorganize again when no concurrent users are accessing the file. The reorganize operation then immediately attempts to recover the space before starting the reorganize. If concurrent data changes have occurred since the initial reorganize, only a portion of the space might be recovered.

Note: The iSeries Navigator interface for specifying ALWCANCEL(*YES) is checking the “Allow reorganization to be suspended” setting.

When using the new ALWCANCEL(*YES) algorithm, keep in mind:

- ▶ If the LOCK parameter specifies *EXCLRD or *SHRUPD, the result of the reorganize is not guaranteed to be exact, since concurrent users may be locking rows or changing rows in the file. For example, if another user has row 43 locked, the reorganize will not be able to move it, so it will not necessarily be in the correct position at the end of the reorganize. In many cases this is fine. However, some applications depend on exact positions. In those cases, LOCK(*EXCL) should be used. If you specify LOCK(*EXCL), the lock is kept for the duration. If you specify LOCK(*EXCLRD) or LOCK(*SHRUPD), you keep that lock for the duration and you additionally need an exclusive lock for a very brief period.

iSeries Information Center has much information about the lock and other reorganize file member (table) functions.

- ▶ The rebuild access paths (RBDACCPH) parameter specifies whether to rebuild or maintain any valid access paths (other than an access path specified as the KEYFILE or a MAINT(*REBLD) access path) over the member. Parameter values are:
 - *YES: Access paths will be rebuilt synchronously at the end of the reorganize operation.
 - *OPTIMIZE: Access paths will either be rebuilt asynchronously at the end of the reorganize operation or maintained during the reorganize, based on which method will result in the access paths being rebuilt the fastest.
 - *NO: Access paths will be maintained during the reorganize. If *NO is specified, ALWCANCEL(*YES) must also be specified
- ▶ The RI (index created as the result of specifying a Referential Integrity (RI) constraint between a parent table and at least one dependent table) and unique indexes are always maintained regardless of the index option.

The type of reorganize you decide to use will depend on several factors. For example:

- ▶ Is your goal simply to recover space
- ▶ Is the sequence of the rows important?
- ▶ Is it important that the reorganize can be canceled (suspended)?
- ▶ Is it important to allow concurrent access to the file?

The development lab did some testing on their own internal benchmark applications and came up with the table shown in Figure 16-6 to help evaluate which option is most appropriate for you, based on the factors shown. The blue shaded entries are the characteristics of a key file (KEYFILE) option that make its choice particularly desirable.

	ALWCANCEL(*NO)		ALWCANCEL(*YES)		
	KEYFILE (*NONE)	KEYFILE (*FILE or keyfile)	KEYFILE (*RPLDLTRCD)	KEYFILE (*NONE)	KEYFILE (*FILE or keyfile)
Cancel and restart	No	No	Yes	Yes	Yes
Concurrent Access	No	No	Yes	Yes	Yes
Parallel processing	Only index rebuilds	Only index rebuilds	Data movement and index rebuilds	Data movement and index rebuilds	Data movement and index rebuilds
Non-parallel performance	Very fast	Fast	Very fast	Slower	Slowest
Temporary storage	Double data storage	Double data storage	Journal receiver storage	Journal receiver storage	Journal receiver storage
LIFO KEYFILE index processing	N/A	Duplicates reversed	N/A	N/A	Duplicate ordering preserved
Index processing (non-KEYFILE)	Synchronous or asynchronous rebuilds	Synchronous or asynchronous rebuilds	Maintain indexes or synchronous or asynchronous rebuilds	Maintain indexes or synchronous or asynchronous rebuilds	Maintain indexes or synchronous or asynchronous rebuilds
Final row position exact	Yes	Yes	Only if LOCK(*EXCL) and not restarted	Only if LOCK(*EXCL) and not restarted	Only if LOCK(*EXCL) and not restarted
Amount of CPU & I/O used	Smallest	Next smallest	Smallest	More	Most
Variable length segment reorganize	Good	Good	Worse	Worse	Worse
Allows referential integrity parents and FILE LINK CONTROL DataLinks	Yes	Yes	No	No	No
Allows QTEMP & Database Cross Ref Files	Yes	Yes	No	No	No
HABP replication cost	Minimal - one journal entry	Minimal - one journal entry	More - journal entries for all rows moved	Most - journal entries for all rows moved	Most - journal entries for all rows moved

Figure 16-6 Considerations for using ALWCANCEL (*NO) or ALWCANCEL(*YES)

Remember that using the new algorithm requires that the tables are journaled. The support does not apply to tables with data links, a table that is a parent in a referential constraint, tables in QTEMP, or the DB2 for i5/OS database cross reference files.

Figure 16-7 shows an example of the iSeries Navigator Database → Schema → Tables Data → Reorganize window for table ITMFIL in schema (library) PFREXP. With journaling active for ITMFIL you can see the Allow reorganization to be suspended settings activated.

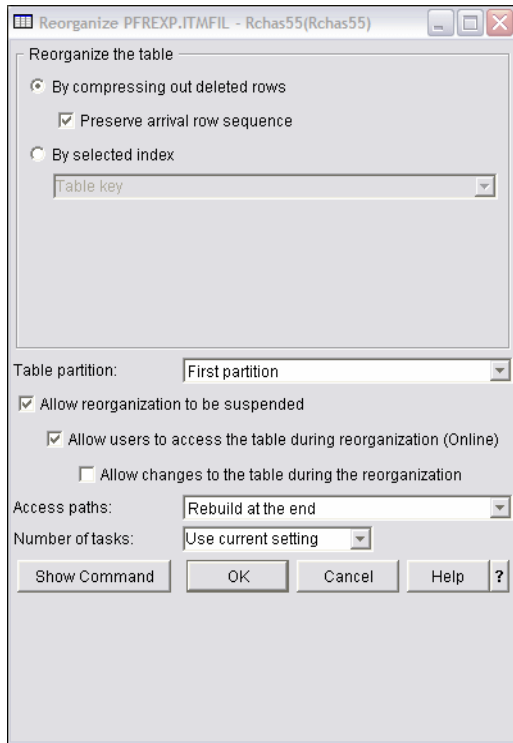


Figure 16-7 Allow reorganization to be suspended (ALWCANCEL(*YES))

You can use the iSeries Navigator interface to see the status of a reorganization. Figure 16-8 shows an example of the iSeries Navigator interface showing the completed status of the reorganize function.

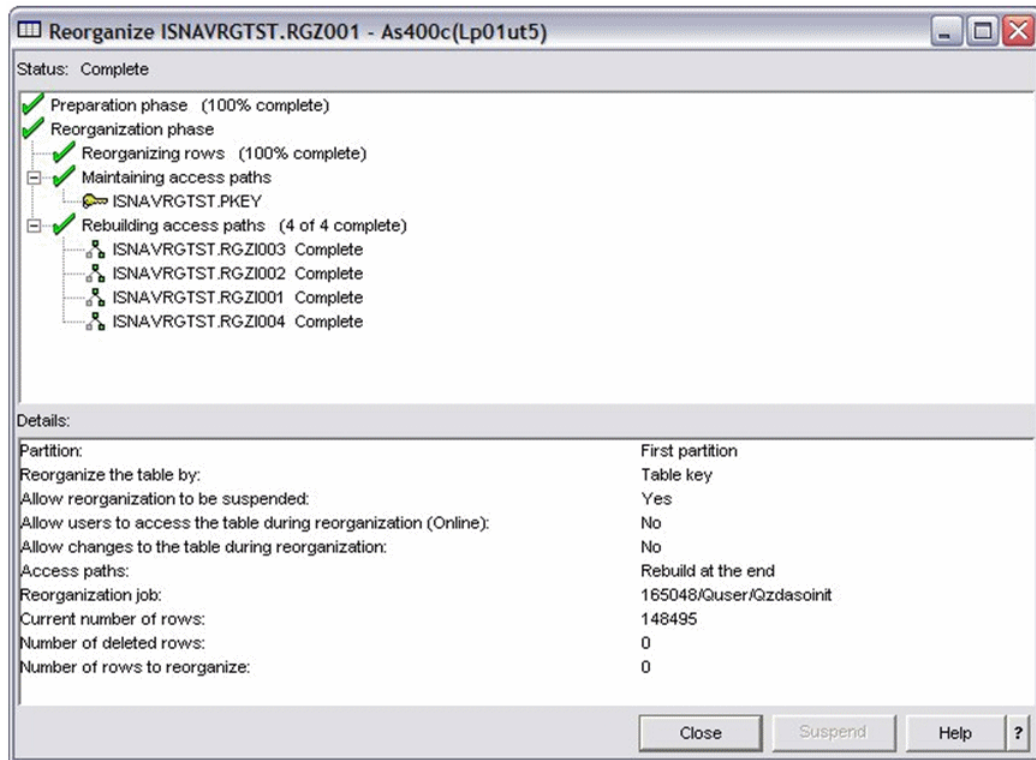


Figure 16-8 Completed status reorganize file member

16.2.4 Disk hardware performance - smaller cache disk controllers

New disks and disk controllers (IOAs) were announced during 2006 and 2007 that have improved performance over the corresponding disk hardware previously available. The table shown in Figure 16-9, is taken from the January 2007 Performance Capabilities Reference manual. It shows disk drive speeds, maximum disk drive interface speed supported by each controller, and the write cache sizes of the disk controllers.

CCIN Codes	Approximate Size (GB)	RPM	Seek Time (ms)		Latency (ms)	Max Drive Interface Speed (MB/s) when mounted in a given enclosure		
			Read	Write		5074/5079	5094/5294	5786/5787
6718	18	10K	4.9	5.9	3	80	80	NA
6719	35	10K	4.7	5.3	3	80	160	NA
4326	35	15K	3.6	4.0	2	Not Supported	160	320
4327	70	15K	3.6	4.0	2	Not Supported	160	320
4328	140	15K	3.6	4.0	2	Not Supported	160	320
CCIN Codes	(IOA) Feature Codes	Cache non-compressed / up to compressed		Min/Max # of drives in a RAID set	Max Drive Interface Speed supported #1 (MB/s)			
5702	5705, 5712, 5715, 0624	NA		NA	160			
5703	5703	40 MB		3/18	320			
2757	5581, 2757, 5591	235 MB / up to 757		3/18	160			
2780	5580, 2780, 5590	235 MB write/up to 757 256 MB read/up to 1GB		3/18	320			
5709 Write cache card for built in IOA	5709, 5726, 9509	16 MB		3/8	NA			
573D Write cache card for built in IOA	5727, 5728, 9510	40 MB		3/8	NA			
571A	5736, 5775, 0647	NA		NA	320			
571B	5737, 5776, 0648	90 MB		3/18 RAID-5 4/18 RAID-6	320			
571E/574F	5738, 5777, 5582, 5583	390 MB write/up to 1.5GB 415 MB read/up to 1.6GB		3/18 RAID-5 4/18 RAID-6	320			
571F/575B	5739, 5778, 5781, 5782, 5799, 5800	390 MB write/up to 1.5GB 415 MB read/up to 1.6 GB		3/18 RAID-5 4/18 RAID-6	320			

Figure 16-9 Disk drive models and disk controller (IOA) base performance comparisons

Note that the actual drive interface speed (MBps) is the minimum value of the maximum supported speeds of (1) the drive, (2) the enclosure, and (3) the IOA. Also note that the minimum value for the various (1) drive and (2) enclosure combinations are identified in the table.

The follow-on disk controllers to the #2757 and #2780 disk controllers were announced in February 2007. These are dual mode IOAs with larger write and read caches. Their feature numbers and descriptions are summarized in 16.2.5, "Fastest disk controllers (adapters) performance" on page 419. For more information about these new disk controller capacities, see the January 2007 System i Performance Capabilities Reference manual.

The following performance figures contain performance comparison charts from the January 2007 Performance Capabilities manual. The first chart (Figure 16-10) includes the 5709 and 573D, which represent the built-in IOAs with cache within the 520/550/570 CECs. The 5709 had eight 15 K 35 GB DASD units and the 573D had eight 15 K 70 GB DASD, the maximum DASD allowed on the built-in IOAs.




Figure 16-10 Lower speed disk controllers compared to the #2780 high speed disk controller

The other IOAs used ten 15 K 35 GB DASD units. Again, this is all for relative comparison purposes, as the 571B only supports 10 DASD units in a 5094 enclosure and a maximum of 12 DASD units in a 5095 enclosure. The 2757 and 2780 can support up to 18 DASD units with the same performance characteristics as they display with the 10 DASD units, so when considering the right IOA for your environment remember to take into account your capacity needs along with this performance information.

Figure 16-10 shows in this set of tests that the #2780 IOA has the highest throughput rate (I/Os per second (Ops/Sec)). Therefore there is no need include higher operations per second capacity disk IOAs announced February 2007.

The second figure, Figure 16-11 on page 420, includes the higher speed capacity disk controllers announced in February 2007.

The #5737 (IOP required) and #5776 (IOP-less) have 90 MB write cache and support RAID5 or RAID6. Their I/O operations per second rates are identified by CCIN number 571B in this figure. This orderable feature to CCIN value cross-reference can be verified with Figure 16-9 on page 417. CCIN number 573D represents the orderable features #5727, #5728, and #9510, which have a 40 MB write cache.

For 571B, you can see a disk response time up to approximately 10 milliseconds. RAID5, RAID6, and IOP/IOP-less configurations all have similar disk I/O rates. The 2780 has greater I/Os per second capacity at the 10 millisecond response time value.

The next topic shows operations per second rates for the fastest IOAs supported under i5/OS.

See the January 2007 or later Performance Capabilities Reference manual for more complete performance test results and considerations. This can be found at:

<http://www.ibm.com/eserver/series/perfmgmt>

16.2.5 Fastest disk controllers (adapters) performance

The 2757 and 2780 can support up to 18 DASD units with the same performance characteristics as they display with the 10 DASD units, so when you are considering the right IOA for your environment, remember to take into account your capacity needs along with this performance information.

The new (as of February 2007) highest speed disk controllers (several feature numbers including #5738 and #5777) with 1.6 GB read cache and 1.5 GB write cache support up to 20 disk drives and offer even higher disk I/Os per second capacity.

The #2757 and #2780 can optionally have an auxiliary write cache (#5580, #5581) of 757 MB. The #5738 and #5777 can optionally have an auxiliary write cache (#5582, #5583) of 1.5 GB.

Attention: The cache sizes cited in most documentation are the *up to* values, which assume a significant amount of data compression. The following list shows the actual cache size as well as the *up to* values for the fastest disk IOAs supported on System i configurations:

- ▶ 2757:
 - Write cache: 235 MB with up to 757 MB
 - Read cache: 0
- ▶ 2780:
 - Write cache: 235 MB with up to 757 MB
 - Read cache: 256 MB with up to 1 GB
- ▶ CCIN 571B (#5737, #5776, #0648):
 - Write cache: 90 MB
 - Read cache: 0
- ▶ CCIN 571E/574F(auxiliary write cache) (#5738, #5777, #5582, #5583):
 - Write cache: 390 MB with up to 1.5 GB
 - Read cache: 415 MB with up to 1.6 GB
- ▶ CCIN 571F/575B (built-in auxiliary write cache) (#5739, #5778, #5781, #5782, #5799, 5800):
 - Write cache: 390 MB with up to 1.5 GB
 - Read cache: 415 MB with up to 1.6 GB

When using the System i Service Tools interface, if you look at the 2757, 2780, or newer large write and read cache disk controllers you see a value that represents the actual cache size, not the larger size value cited in marketing documents. The marketing value assumes a significant level of data compression. For example, for a #2780 IOA you see a value of 235 MB for write cache, not the 757 MB size often discussed in marketing documentation.

Figure 16-11 shows disk I/O operations per second performance comparisons for these fastest disk controllers. The values shown are the CCIN values for these orderable features.




Figure 16-11 2780 - 5738, disk operations per second comparison

You can see the higher operations per second rates supported by the February 2007 disk IOAs with the up to 1.5 GB write cache and 1.6 GB read cache.

See the January 2007 or later System i Performance Capabilities Reference manual for more information.

16.2.6 Virtual I/O restructuring

For V5R4 internal implementation of virtual I/O support has been restructured. The chart in Figure 16-12 shows what was done and presents early Rochester performance test results.

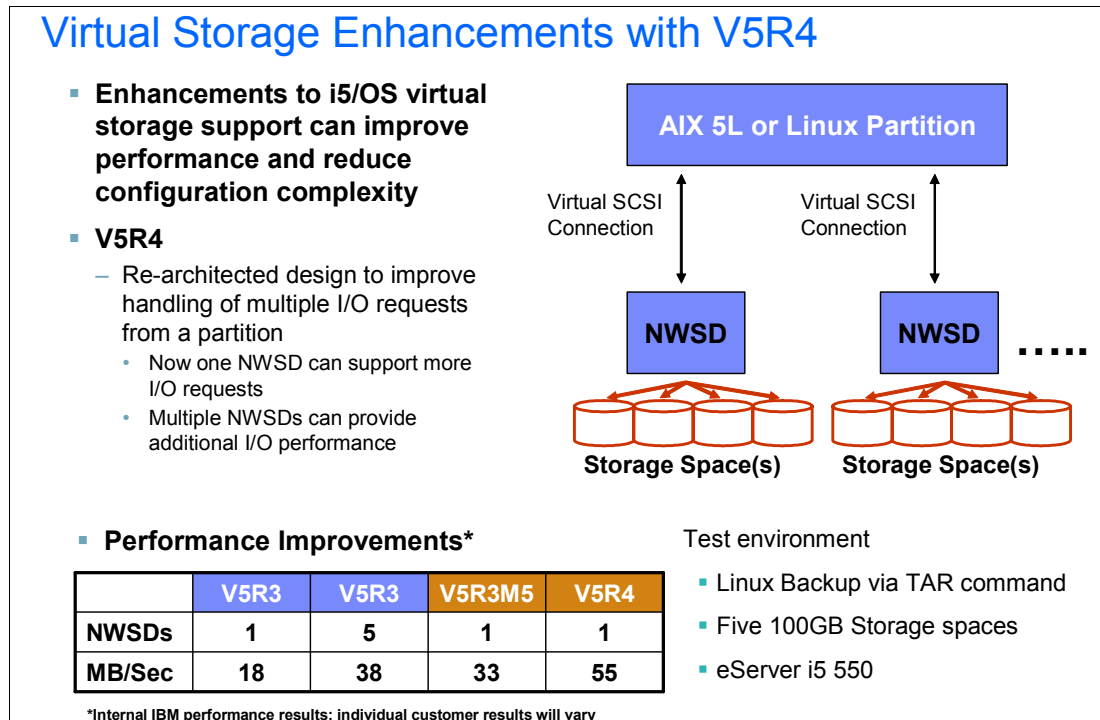


Figure 16-12 Virtual storage enhancements with V5R4

At this time you are cautioned that each environment will have different results when comparing V5R3 to V5R4.

16.2.7 Improved NetServer performance

In i5/OS V5R4 NetServer has been changed to improve performance. NetServer performance enhancements and charts are discussed in 10.8, “iSeries NetServer changes” on page 237.

16.2.8 Java and WebSphere performance

Most WebSphere environments running i5/OS V5R4 will see approximately 5% better performance over environments in V5R3. This is attributed to improvements in the Just-in-Time compiler (JIT) as well as optimizations in database. The Trade6 Benchmark showed even higher performance gains.

Query optimization on the LIKE clause whereby invocations of LIKE are sent down the SQE path and not the CQE path yielded performance improvements of 7–18% in addition to the gains mentioned above.

Since an often used query in the Trade 6 Benchmark uses the LIKE clause, more so than typical WebSphere environments would, most users are not expected to see this 7–18% additional improvement.

In general, i5/OS V5R4 has 3–10% improvement compared to V5R3 running the same Java programs. These improvements are primarily based on Java JIT improvements. In V5R3 Java performance improvements are available in V5R3 Java PTF group level 5.

Starting with V5R3, JIT is always recommended for best performance of 32-bit (no DE) or 64-bit. On the other hand, the Direct Execution (DE) is the default for 64-bit JVMs, unless specified otherwise by setting the `java.compiler` property or specifying in `ADDENVVAR` to `jitc`, as shown in the command:

```
ADDENVVAR ENVVAR (JAVA_COMPILER) VALUE(jitc) JAVA CLASS(Test)
```

Or in QSHELL command:

```
java -Djava.compiler=jitc Test
```

Even if running JIT, the JVM will have to create a Java program object (with optimization level `*INTERPRET`) the first time a particular Java class is used on the system, if one does not already exist.

The creation of this program object is much faster than creating a full DE program, but it may still take some time. But running `CRTJVAPGM` with `OPTIMIZE(*INTERPRET)` will create this program ahead of time, making the first startup faster.

By default, WebSphere Application Server runs in JIT mode, but if you have Java Servlets running only in Direct Execution (DE), you may have to change the WAS instance or re-create your Java program to be able to run it in JIT.

IBM Technology for Java Virtual Machine

i5/OS V5R4 delivers two JVM and Software Development Kit (SDK) environments — the classic JVM running in 64-bit mode and the 32-bit IBM Technology for JVM and the Java2 SDK, Standard Edition.

The 32-bit implements a standard IBM JVM 1.5 runtime environment and SDK and is specifically focused for applications on smaller systems, with less impact on memory utilization. In addition, it can be implemented mostly without any code changes for the majority of the applications and requires less storage and heap size than the 64-bit version.

Garbage Collection is similar to most other platforms for the 32-bit JVM, while it remains unchanged for the 64-bit JVM.

WebSphere Application Server

In general, WAS V5.1 and V6.0 show 5% better performance on V5R4 compared to V5R3. This performance gain is based primarily on the improved JIT in V5R4 and V5R4 database optimizations.

16.2.9 Object conversion

In almost all new releases there is some object conversion that takes place. In general, there are more object conversions in V5R4 than in V5R3. This is because a general conversion is done during first V5R4 IPL. In addition, conversion can also take place upon first use via the background LIC tasks or by the use of the Start Object Conversion command.

SQL program object conversion

The internal representation of SQL statements in `*PGM`, `*SRVPGM`, `*MODULE`, and `*SQLPKG` objects changed in V5R4. `*PGM`, `*SRVPGM`, and `*SQLPKG` objects that were created prior to V5R4 will be converted automatically to the new format on first use. You can

also initiate object conversions using the Start Object Conversion (STROBJCVN) command. *MODULE objects that were created prior to V5R4 cannot be permanently converted, and a conversion will be performed each time the *MODULE is bound into a program or service program.

The amount of time to convert a single program object is small, less than 160 milliseconds of CPU time on a model 890 system for a program approaching the system limit on number and size of SQL statements. On a model 840 system, it is less than 350 milliseconds. The conversion time is proportional to the number of SQL statements in the object and to the cumulative size of the SQL statements. If large numbers of objects need to be converted, the conversion time could be noticeable.

Hardware Storage Protection object conversion

The operating system creates objects using a hardware storage protection attribute. This is further discussed in “Enhanced Hardware Storage Protection” on page 137.

Database conversion recommendations:

- ▶ Understand critical application DB files and libraries on V5R2 or V5R3.
- ▶ Use program QSYS/QDBFIMIC from V5R2 or V5R3 PTFs on critical DB files and libraries to get the number of objects being converted.
- ▶ Determine whether first use conversion is acceptable. If not, perform DSPxxx on the library or the file within library to force conversion when you want to.
- ▶ The following rates have been extrapolated from memo to users test result examples for estimated time to convert:
 - 830, approximately 5512 CPW 6 processor partition, in a 6402 MB storage pool: ~ 87 objects/second
 - 825, 6600 CPW 6 processor partition, in a 3271 MB storage pool: ~ 103 objects/second
 - 570, 3300 CPW 1 processor partition, in a 512 MB storage pool: ~ 164 objects/second

16.2.10 LPAR considerations

The following should be considered to get maximum throughput from each partition.

- ▶ Dedicated processors equal best performance.

Two dedicated 4-processor partitions yield more total CPW than a single 8-processor partition.
- ▶ For shared processor partitions, with more than one full processor capacity assigned to the partition, best performance should result in most cases if you do the following:
 - Assigned processors Greater Than or Equal (GTE) to 1:
 - Set number of virtual processors up to the next whole processor value above the assigned processor capacity (for example, 3.5 becomes 4) or up to 2x the assigned capacity when there are at least that number of processors activated on the system. That is, 2x could be 4 when there are two active processors on the system.
 - Using uncapped partitions can be very successful, but you must understand the operating environment with multiple partitions active concurrently. For example, an uncapped partition may get improved performance when there is additional processing unit capacity available (use more than 100% utilization) but get less performance at another time when other partition work is active and no additional processor capacity is available.

Use different values for the Sharing Mode - Weight parameter to prioritize processor capacity among all active partitions.

- Assigned processors less than 1 (base micropartitioning):

Set the number of virtual processors to no greater than the number than 2x assigned capacity.

- ▶ Less than expected partition performance may be realized if all of the following are true:
 - The number of virtual processors value approaches 10x (maximum) the assigned shared processor capacity. This leads to making the system less efficient by requiring it to manage more active threads than are necessary.
 - The partition has less than one full processor capacity assigned to it (for example, .75) and the active workloads consume significant CPU utilization. This can be further degraded if dozens of such threads are active at the same time.

An example of this would be .6 processor partition with many threads active, each for longer than 6 milliseconds.

For more complete information see the white paper *i5/OS LPAR Performance on POWER4 and POWER5 Systems* PDF, which can be found on the performance management Web site, using the Resource Library link.

16.2.11 Performance updated on integrated xSeries servers

The V5R4 System i Performance Capabilities Reference manual, August 2006 or later, contains updated performance information about integrated xSeries servers. This section includes information excerpted from this book.

For more information go to:

<http://www.ibm.com/eserver/series/perfmgmt>

Select the **Reference Library** link.

In this section we discuss the following regarding xSeries integration:

- ▶ Storage pool rules of thumb
- ▶ Disk operation rates rules of thumb

Note that rules of thumb are starting places. They may not apply in all customer environments.

Attachments cover xSeries servers connected via:

- ▶ Integrated xSeries Servers (IXS) hardware feature
- ▶ Integrated xSeries Servers and BladeCenter servers connected over the HSL loop via the Integrated xSeries Server Adapter (IXA)
- ▶ Integrated xSeries Servers connected over new in V5R4 iSCSI adapter and Ethernet cabling.

Figure 16-13 reminds us of several connection options available via the iSCSI support. In the upper right graphic we see two integrated xSeries servers connecting over a single i5/OS partition's Host Bus Adapter. In the lower right graphic we see three BladeCenter servers, each connecting to the same i5/OS partition's three Host Bus Adapters.

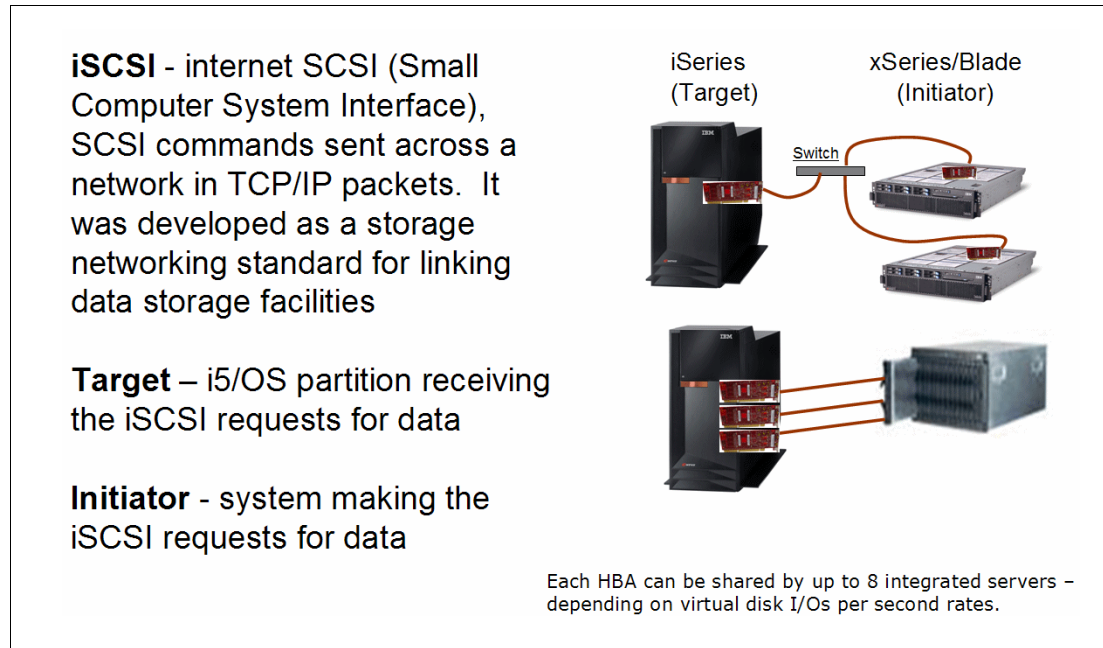


Figure 16-13 iSCSI attachment examples

Note the text saying that up to eight integrated servers can connect over a single HBA, depending on the aggregated total of virtual I/Os per second that will be performed.

i5/OS storage pool rules of thumb

Figure 16-14 gives the starting values for storage pool sizes for each active Network Server (Description - NWSD), whether connected via IXS, IXA, or iSCSI. It also shows the additional starting storage pool values when using iSCSI connections. So, if iSCSI is being used and five servers are active, you multiply the values in the *per NWSD column* by 5 and add the storage in the *per HBA column* if using a single HBA.

	For each HBA on the partition	For each NWSD
Machine Pool	21 MB	1 MB
*BASE Pool	1 MB	0.5 MB
IBM Subsystem QFPHIS¹ Private Pool	0.5 MB	4.0 MB² plus 1 MB for NWSD 2 - n (²10.0 MB)

¹ QFPHIS is supported if PTF V5R4 PTF SI23027 is applied.

² 10.0 MB is good starting point. 4.0 MB is minimum.

Figure 16-14 Integrated xSeries server storage pool rules of thumb

Of course you would have decided how many HBAs were required by using the values shown in Figure 16-15.

Note that in Figure 16-15 the disk I/Os per second rules of thumb are modified by how much processor power (CPWs) you have per 1 K of disk I/Os per second.

	CPWs per 1K I/Os per second *
iSCSI linked disks	190
IXS/IXA static or dynamically linked disks with write caching enable	130
IXS/IXA shared or quorum linked disks or write caching disabled	155

* While the disk I/O activity driven by the IXS/IXA or iSCSI is not strictly a "CPW" type load, the CPW estimate is still a useful metric to estimate the amount of iSeries CPU required for a load. You can use the values above to estimate the CPW requirements if you know the expected I/O rate.

Figure 16-15 Integrated xSeries server disk I/Os per second rules of thumb

For more complete information go to the System i Performance Capabilities Reference manual.

16.3 Performance tools updates

i5/OS V5R4 offers many enhancements to the i5/OS-based performance tools. This section discusses these enhancements in detail.

16.3.1 iSeries Navigator Monitor

Enhancements to the iSeries Navigator Monitor are added in i5/OS V5R4. Now it is possible for the System monitors to be shared among users. When you have a V5R4 GUI and a V5R4 Central system, you can share system monitors and system events. You do this in the same way that you share job monitors, message monitors, and file monitors. You specify the sharing level on the Sharing tab, which is located on the Properties window of the monitor.

In addition, you can now coordinate all graphs to the same time interval on the View Monitor panel.

Another enhancement is that in the Average Line Utilization part of the system monitor, you now have the option to exclude communication lines. You can now exclude heavily utilized communication lines, such as fax lines, from the system monitor graph. For example, when you have two communication line utilization metrics, the average of all of the communication lines is plotted. Thus, if there are one or more lines that you do not want to include in the

average, such as a line that is heavily loaded because of fax traffic, you can optionally exclude these lines. For instructions on how to do this, use the online help for system monitors.

Here are several other enhancements in the areas of system monitor and graph history in i5/OS V5R4 iSeries Navigator:

- ▶ **Save:** You can now save a screen capture of the Graph History or the Systems Monitor window (or just the selected graph) to your client workstation drive.
- ▶ **Print:** You can print your graphs from the Graph History or System Monitor window. You can print the entire window, or just the selected graph.
- ▶ **Coordinate:** The system monitor has the ability to link all of the displayed graphs together. When this option is selected, all of the graphs that make up the monitor will be coordinated in terms of the time-slice shown, and the scroll position in the graphs. Thereafter, if you scroll to one position in the graph, all of the graphs in the monitor will simultaneously scroll to that same position.
- ▶ **Organize:** It is possible now to move the graphs in the System Monitor window into any visual configuration and size that you want. When you close the Systems Monitor window, the size and positions of the graphs are saved. The next time you open the window, the graphs will display in that configuration.
- ▶ **Visualize:** Change the colors of the lines on the graph by using a menu option on the Graph History and System Monitor windows, **View** → **Set Monitor Colors**. Prior to V5R4 the only way to change the graph line colors were to go to the main iSeries Navigator window and select **Properties**.

16.3.2 Performance Tools for i5/OS 5722-PT1

i5/OS brings new improvements to the licensed program 5722-PT1 Performance Tools.

Job trace

One notable change is that the job trace reports have been changed to take advantage of the Job Trace functionality in the Start Trace (STRTRC), End Trace (ENDTRC), and Print Trace (PRTRC) commands. This also includes changes in the Performance Tool reports and the addition of a new section on the Component report of the Domino server activity statistics. In summary, here are the enhancements to the Job trace reports:

- ▶ In V5R4, the STRJOBTRC command issues STRTRC MAXSTG(maxstg) TRCFULL(*STOPTRC) to start the job tracing function to collect performance statistics for the specified job. The Start Trace(STRTRC) command is more flexible and less intrusive than the Trace Job (TRCJOB) command. It allows tracing across multiple jobs and shows more in-depth details about the module flow.

The default value of Maximum storage (MAXSTG) in the Start Job Trace (STRJOBTRC) command is changed from 1024 K to 10000 K bytes. It takes the default value from STRTRC.

In V5R4, there is new Thread ID (THDID) parameter in STRJOBTRC, as shown in Figure 16-16. You can specify a list of up to twenty threads whose calls and returns are included in the trace. Only trace records for the specified thread identifiers are included.

```

Start Job Trace (STRJOBTRC)

Type choices, press Enter.

Maximum storage . . . . . 10000      1-4000000 K
Job name . . . . . *              Name, *
User . . . . .                  Name
Number . . . . .                000000-999999
Thread ID to include . . . . . *ALL   Thread ID, *ALL
      + for more values

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 16-16 STRJOBTRC command parameter changes

- ▶ In V5R4 the command PRTJOBTRC processes the DB file created by STRTRC. Also, the System model parameter is no longer required in this command, as shown in Figure 16-17.

```

Print Job Trace (PRTJOBTRC)

Type choices, press Enter.

Data base file member . . . . . QAJOBTRC      Name
Data base file library . . . . . QPFRDATA     Name
Report type . . . . . *BOTH                 *BOTH, *DETAIL, *SUMMARY
Report title . . . . . *BLANK

Starting sequence number . . . . *FIRST    1-999999, *FIRST
Ending sequence number . . . . . *LAST      Number, *LAST
Transaction ending program . . . . QT3REQIO   Name, QT3REQIO, *BATCH
Transaction starting program . . . QWSGET     Name
System model code . . . . . *CUR           Character value, *CUR
Job name . . . . . PRTJOBTRC              Name, PRTJOBTRC, *MBR
Job description . . . . . QPFRJOB          Name, *NONE
Library . . . . . *LIBL                   Name, *LIBL, *CURLIB

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display

```

Figure 16-17 PRTJOBTRC command parameter changes

- ▶ The MAXSTG parameter in STRJOBTRC accepts a maximum value of 10,000 KB.
- ▶ There is a new parameter to select thread IDs (up to 20).

16.3.3 iSeries Navigator performance tools plug-in

The performance tools plug-in for iSeries Navigator contains the following enhancements:

- ▶ The All Jobs view of the Display Performance Data window contains two new columns that show the IP address most recently associated with the job and the remote port number. Refer to Figure 16-18 for a window screen capture sample of how it looks in the iSeries Navigator interface.

Job	User	Number	Type	IP Address	Remote port...	Disk I/O ...	CPI
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5544		459
Qrwtssvr	Quser	996798	DDM Server	9.5.80.209	5127		260
Qrwtssvr	Quser	996799	DDM Server	9.5.80.209	5128		2174
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5589		38
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5282		93
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5485		99
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5482		355
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5410		2
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5608		8
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5417		37
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5913		97
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5263		401
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5271		644
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.80.209	5462		2
Qpadev006l	Rgarvey	087633	Pass through	9.5.64.162	6082		503
Qypsjsvr	Qypsjsvr	005574	Batch	9.5.211.23	1418		2
Qhttp00075	Qtcp	087524	Batch	9.5.177.154	34564		97
Qhttp00075	Qtcp	087655	Batch	9.5.177.154	36242		14
Qhttp00076	Qtcp	087472	Batch	9.5.177.154	32962		95
Qhttp00076	Qtcp	087506	Batch	9.5.177.154	34010		62

Figure 16-18 Display Performance Data window in iSeries Navigator

- ▶ You can print graphs that help with analyzing performance in the Performance Tools graphical user interface (GUI).

To print the performance data graphs, follow these steps:

1. In the iSeries Navigator window, expand **Configuration and Service**.
2. Right-click **Collection Services**, select **Performance Tools**, and select **Performance Data**.
3. Select the performance data that you want to print. Click **Display** to open the Display Performance Data window.
4. From the menu bar, select **Graphs** → **Print**.
5. On the Print dialog, you can select which graphs to print and whether to print the graphs one per page, two per page, or four per page. When you have made your selections, click **OK** to print the selected graphs.

16.4 Changes to job trace and job analysis reports

In this section we discuss changes to job trace and job analysis reports.

The Job Trace Information report (QPPTTRCD)

The Job Trace Information report shows the job trace data that is collected with the STRJOBTRC and ENDJOBTRC commands. Changes in this report are shown in Figure 16-19.

The screenshot shows the QPPTTRCD report with three callouts: 'Microsecond resolution' pointing to the TIME column, 'Thread ID' pointing to the THREAD column, and 'Increased' pointing to the ENTRY, EXIT, and INV columns. The table below is a transcription of the data shown in the screenshot.

TIME	THREAD	SEQNR	FUNCTION	PROGRAM	LIBRARY	ENTRY	EXIT	INV	CPU	DB	NON-DB
21:39:17.894381	00000001	00034836	XCTL	QDMROUTE	QSYS	0002BE	0002BE	010	.007481		3
21:39:17.894385	00000001	00034837	XCTL	QDBEXDFI	QSYS	001036	000001	010	.000024		
21:39:17.894428	00000001	00034838	CALL	QLIMROIR	QSYS	0001DE	000001	011	.002890		
21:39:17.894458	00000001	00034839	RETURN	QLIMROIR	QSYS	0001DF	000133	010	.001967		
21:39:18.649196	00000001	00017431	RETURN	QDBEXDFI	QSYS	001037	00024A	009	.321531		127
21:39:18.661727	00000001	00086950	CALL	QWCLOBJL	QSYS	00108B	000001	010	.008884		2
21:39:18.674400	00000001	00121754	CALL	QWCSCVTR	QSYS	00074C	000001	011	.035423		4
21:39:18.674420	00000001	00121755	RETURN	QWCSCVTR	QSYS	00074D	000126	010	.000689		
21:39:18.694359	00000001	00121756	CALL	QDBOBJLK	QSYS	000388	000001	011	.006419		3
21:39:20.546996	00000001	00034840	RETURN	QDBOBJLK	QSYS	000389	000056	010	1.200703	168	166
21:39:20.547229	00000001	00034841	RETURN	QWCLOBJL	QSYS	00108C	000087	009	.011773		
21:39:20.547242	00000001	00034842	CALL	QDMCOPEN	QSYS	0010AB	000001	010	.000428		

Figure 16-19 QPPTTRCD report changes

16.4.1 The Trace Analysis Summary report (QPPTTRC1)

The Job Trace Analysis Summary report shows the number and type of I/O operations, such as database reads, non-database reads, and writes, that occurs for each transaction. In this report, the time resolution is increased to microsecond scale. Changes in this report are shown in Figure 16-20.

The screenshot shows the QPPTTRC1 report with a callout 'Microsecond resolution' pointing to the SECONDS column. The table below is a transcription of the data shown in the screenshot.

	SECONDS	CPU SECONDS	DB READS	NON-DB RDS	WRITES	WAITS	SEQUENCE
WAIT-ACT	3.585887						104359
ACTIVE	61.197166	11.844141	643	837	1		105381
WAIT-ACT	61.197166	11.844141	649	1780	1012	131	52327
ACTIVE	5.283278	17.413898	112	190	36		71092
WAIT-ACT	5.283278	17.413898	761	1970	1048	133	123056
ACTIVE	.000011	17.598178		11			105574
WAIT-ACT	.000011	17.598178	761	1981	1048	133	105577

Figure 16-20 QPPTTRC1 report changes

16.4.2 The Trace Analysis I/O Summary report (QPPTTRC2)

The Job Trace Analysis I/O Summary report shows the number of IBM-supplied database modules, such as GETDR and GETSQ, used during the transaction, and the number of full and shared file opens and closes, the number of subfile operations, and the number of messages that occurred in the transaction.

16.5 Changes in the System report

Here we discuss the various changes in each of the sections of the System report.

Change to storage pool utilization: System report

The Storage Pool Utilization section of the System Report helps you set the storage pool size and activity level.

As shown in Figure 16-21, changes in the System Report include:

- ▶ Values for columns DB and Non DB are expanded one space.
- ▶ The values under the Size column are changed to be expressed in megabytes.

System Report								
Storage Pool Utilization								
Member	: Q284120043		Model/Serial . . .	: 595/02-55B0C		Main storage		
Library	: QMPGDATA		System name	: RCHASKOW		Version/Rele		
Partition ID . . .	: 001		Feature Code	: 7487-8966		Int Threshold		
Virtual Processors:	8		Processor Units . . .	: 8.0				
Pool	Expert	Size	Act	CPU	Number	Average	-----	DB
ID	Cache	(MB)	Lvl	Util	Tns	Response	-----	Fault
----	----	----	----	----	----	----	----	----
*01	0	2,476	0	.5	0	.00		.0
*02	3	12,462	1,447	20.3	0	.00		2.6
*03	3	4,117	477	16.2	499	.32		.1
04	3	15	6	.0	0	.00		.0
Total		19,070		37.1	499			2.8
Average						.31		

Figure 16-21 Storage pool utilization: System report

16.5.1 Changes to the disk utilization section of the System report

There are also changes in the Disk utilization section of the System Report, as shown in Figure 16-22.

Disk Utilization													
System report (RCHTWINS)													
Member :	Q271090001	Model/Serial . . :	570/10-3484C	Main storage . . . :	52.0 GB	Started :							
Library :	PTBIGCVT	System name . . . :	RCHTWINS	Version/Release . . :	5/ 3.0	Stopped :							
Partition ID . . :	001	Feature Code . . . :	7476-8971	Int Threshold . . . :	100.00 %								
Virtual Processors:	12	Processor Units . . :	12.0										
Unit	Name	Type	Size (M)	IOP	IOP	Util	Dsk CPU	--Percent--	Op Per	K Per	- Average Time	Per I/O --	
Unit	Name	Type	(M)	Util	Name	Util	Full	Util	Second	I/O	Service	Wait	Response
0106	DD071	4326	30,769	.0	CMB13	.0	44.8	.0	2.96	6.8	.0000	.0001	.0001
0107	DD069	4326	35,165	.0	CMB12	.0	44.8	.0	.26	14.4	.0000	.0015	.0015
0108	DD124	4326	35,165	.0	CMB19	.1	44.8	.0	1.78		.0005		.0005
0109	DD070	4326	35,165	.1	CMB14	.0	44.8	.0					.0002
0110	DD022	4326	30,769	.0	CMB17	.0	44.8	.1					.0008
0111	DD040	4326	35,165	.0	CMB17	.0	44.8	.0	.36	15.2	.0000	.0019	.0019
Total for ASP ID:			1	3,327,443									
Average							44.8	.0	1.44	7.8	.0000	.0004	.0004
ASP ID/ASP Rsc Name:			2/										
0112	DD120	4326	30,769	.0	CMB15	.1	8.7	.0	.00	.0	.0000	.0000	.0000

Figure 16-22 Disk utilization section: System report

Changes to this section of the System Report include:

- ▶ The ASP Rsc Name and ASP ID columns have been removed.
- ▶ A label has been added at the beginning of each ASP/IASP section that indicates the ASP ID and ASP Rsc Name. The ASP Rsc Name is printed only when the DSASPN (ASP resource name) field in the QAPMDISK database file contains data.
- ▶ Totals and averages for each ASP/IASP section are printed at the end of each group of ASP/IASP.
- ▶ Totals and averages for all disk units are printed at the end of the Disk Utilization section, as before.

16.5.2 Changes to the workload section of the system report

The Workload section of the system report displays the interactive and non-interactive workload of the system (Figure 16-23).

System Report		
Workload		
Member	Q284133410 Model/Serial	595/02-55B0C Main storage
Library	QMPGDATA System name	RCHASKOW Version/Release
Partition ID : 001	Feature Code	7487-8966 Int Threshold
Virtual Processors:	8	Processor Units : 8.0
Average CPU Utilization		36.7
CPU 1 Utilization		37.6
CPU 2 Utilization		32.8
CPU 3 Utilization		35.1
CPU 4 Utilization		35.4
CPU 5 Utilization		37.0
CPU 6 Utilization		37.6
CPU 7 Utilization		38.2
CPU 8 Utilization		39.4
Total CPU Utilization (Interactive Feature)0
Time exceeding Int CPU Threshold (in seconds)		0
Total CPU Utilization (Database Capability)		4.0

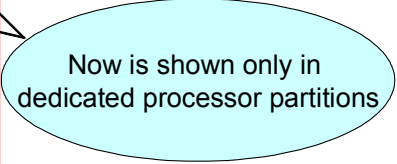


Figure 16-23 Workload section: System report

The change to the workload section of the System Report is that this report shows individual CPU utilization for all processors in dedicated processor partitions. In shared processor partitions, individual CPU utilization rows are not printed.

16.5.3 HTTP Server summary

The HTTP Server Summary section of the System Report includes summary data at the server instance level for the IBM HTTP Server (powered by Apache).

The columns Non-SSL Inbound Connections, SSL Inbound Connections, Requests Received, and Responses Sent are shown at a rate of hits per second, as shown in Figure 16-24.

System Report 101105 0:0
 HTTP Server Summary Page

Member . . . : MBR2 Model/Serial . . : 595/02-55B0C Main storage . . . : 20.2 GB Started : 07/05/05 12:2
 Library . . . : NONEXIST System name . . : RCHASKOW Version/Release . . : 5/ 4.0 Stopped : 07/06/05 00:0
 Partition ID : 001 Feature Code . . : 7487-8966 Int Threshold . . : 100.00 %
 Virtual Processors: 8 Processor Units : 8.0

Server name	Server job user	Server job number	Server start date/time	Threads		Inbound Connections		Requests	Responses
				Active	Idle	Non-SSL	SSL	received	sent
ADMIN	QTMHHTTP	998028	07/05/05 15:12	1	5	.43	.00	.42	.42
ADMIN	QTMHHTTP	998124	07/05/05 15:32	0	5	.21	.00	.21	.21
ADMIN	QTMHHTTP	998139	07/05/05 15:37	0	5	.01	.00	.01	.01
SMART_RPTS	QTMHHTTP	997217	07/05/05 10:22	0	40	.00	.00	.00	.00
SMART_1113	QTMHHTTP	997223	07/05/05 10:22	0	40	.00	.00	.00	.00
SMART_1114	QTMHHTTP	997226	07/05/05 10:22	0	40	.00	.00	.00	.00
SMART_1115	QTMHHTTP	997231	07/05/05 10:22	0	40	.00	.00	.00	.00
V5R4A1	QTMHHTTP	998242	07/05/05 15:52	0	40	.00	.00	.00	.00

Figure 16-24 HTTP server summary: System report

16.6 Changes in the Component Report

Like the System Report, there are also changes in the sections of the Component Report in i5/OS V5R4.

16.6.1 HTTP Server activity

The HTTP Server Activity section of the Component report includes detailed information about various HTTP Server request types such as CGI, WebSphere, or Proxy requests.

Changes in the report, as shown in Figure 16-25, are:

- ▶ A new column, Hits/Second, has been added to this report.
- ▶ A row that contains the average requests in hits per second has been added to the Total subsection of this report.

Component Report								
HTTP Server Activity								
Member . . . : MBR2	Model/Serial . . : 595/02-55B0C	Main storage . . . : 20.2 GB	Starte					
Library . . . : NONEXIST	System name . . : RCHASKOW	Version/Release . . : 5/ 4.0	Stoppe					
Partition ID : 001	Feature Code . . : 7487-8966	Int Threshold . . : 100.00 %						
Virtual Processors: 8	Processor Units : 8.0							
Server : 998028/QTMHHTTP/ADMIN								
----- Responses -----				KB		KB		
Itv	Req	Requests		Pct	Transmitted	Received	Hits	
End	type	Received	Sent	Error	/Second	/Second	/Second	
15:15	SR	39	39	0	.00	1	0	.32
15:15	CG	2	2	0	.00	0	0	.01
15:15	UM	10	10	0	.00	0	0	.08
15:20	SR	24	24	0	.00	0	0	.08
15:20	UM	37	37	0	.00	1	0	.12
15:25	SR	100	100	0	.00	0	0	.33
15:25	CG	19	19	0	.00	1	0	.06
15:25	UM	206	206	0	.00	2	0	.68

Figure 16-25 HTTP Server activity - Component report

16.6.2 Domino server statistics

In V5R4 there is a new Domino Server Statistics section in Component Report, as shown in Figure 16-26. The Domino section of the Component report includes detailed metrics for Domino server statistics for iSeries systems.

Component Report												12/15/05 18:03:20
Domino Server Activity												
DOM ONLY JC												
Member	Q348073157	Model/Serial . .	270/10-4RT9M	Main storage . . .	8192.0 MB	Started	12/14/05					
07:31:5		Library	QMPGDATA	System name . . .	:RCHASM27	Version/Release . .	5/ 4.0	Stopped	12/15/05			
00:00:0		Partition ID . .	: 000	Feature Code . . .	:22AB-2253-1520	Int Threshold . . .	: 3.50 %					
Virtual Processors:	2	Processor Units . .	: 2.0									
Server	012441/QNOTES/SERVER											
				Peak	----- Mail -----		----- Database -----		----- Name lookup -----			
Itv	Tns		CPU	Concur	Pending	Waiting	Cache	Cache	Cache	Cache	URLs	
End	/Hour	Users	Util	Users	Outbound	Inbound	Hits	Lookups	Hits	Lookups	Rcv/Sec	
07:35	2,546	41	3.00	261	0	0	510	1665	0	0	0	
07:40	12	3	0.00	3	0	0	18	157	0	0	0	
07:45	1,744	41	2.25	261	0	0	296	1281	0	0	0	
07:50	31	2	.01	3	0	0	20	167	0	0	0	

Figure 16-26 Domino server activity: Component report

16.7 WRKSYSACT

Running the WRKSYSACT command optionally generates the QAITMON output file.

The QAITMON file has the following changes in V5R4:

- ▶ The fields for individual CPU utilization have been removed.
- ▶ The following new fields have been added:
 - CPU Used and CPU Available - CONFTOT
 - Minimum and Maximum Individual CPU Time Used - MINCPU and MAXCPU

Figure 16-27 is an example of V5R4 WRKSYSACT — when your partition is configured to use a shared processor pool and have a maximum virtual processor capacity greater than one. You can change the number of virtual processors dynamically through the HMC interface. For WRKSYSACT to detect this change you must restart WRKSYSACT after the HMC change has completed.

```

                                Work with System Activity
                                01/07/07 13:40:11 RCHAS55B
Automatic refresh in seconds . . . . . 5
Elapsed time . . . . . : 00:00:02   Average CPU util . . . . . : 24.8
Number of CPUs . . . . . : 2         Maximum CPU util . . . . . : 25.7
Overall DB CPU util . . . . . : 8.7   Minimum CPU util . . . . . : 24.0
                                         Current processing capacity: 1.00

Authorization Type . . . :
Type options, press Enter.
  1=Monitor job  5=Work with job

                                Total  Total  DB
                                CPU   Sync  Async  CPU
Opt  Task      User      Number Thread  Pty  Util  I/O   I/O   Util
-----
JOB0301  AS0301    011557 00000026 50 23.3 1895 2418 8.7
QPADEV0004  JC00K    011556 00000007  1  .1   7    0    .0
JO-EVALUAT          0    .0   3    98    .0
SMPOL001          99    .0   0    25    .0

                                Bottom
F3=Exit  F10=Update list  F11=View 2  F12=Cancel  F19=Automatic refresh
F24=More keys

```

Figure 16-27 WRKSYSACT V5R4 for an uncapped partition with more than one processing unit possible

Note the following:

- ▶ Fields Maximum CPU util and Minimum CPU util are not shown when the partition has a maximum processor capacity of 1 or less.
- ▶ The Authorization Type field is intended for a future capability. Disregard any associated online help text.

16.8 Display performance data

The Display Performance Data (DSPPFRDTA) command can now show the IP address associated with the job.

The Display by Jobs and Display Job Detail screens show the IP address that is most recently associated with the job and the remote port number that the network connection uses.

To display a new IP address and port number column in Display Job Screen, press F11.

Figure 16-28 shows the Display Jobs screen.

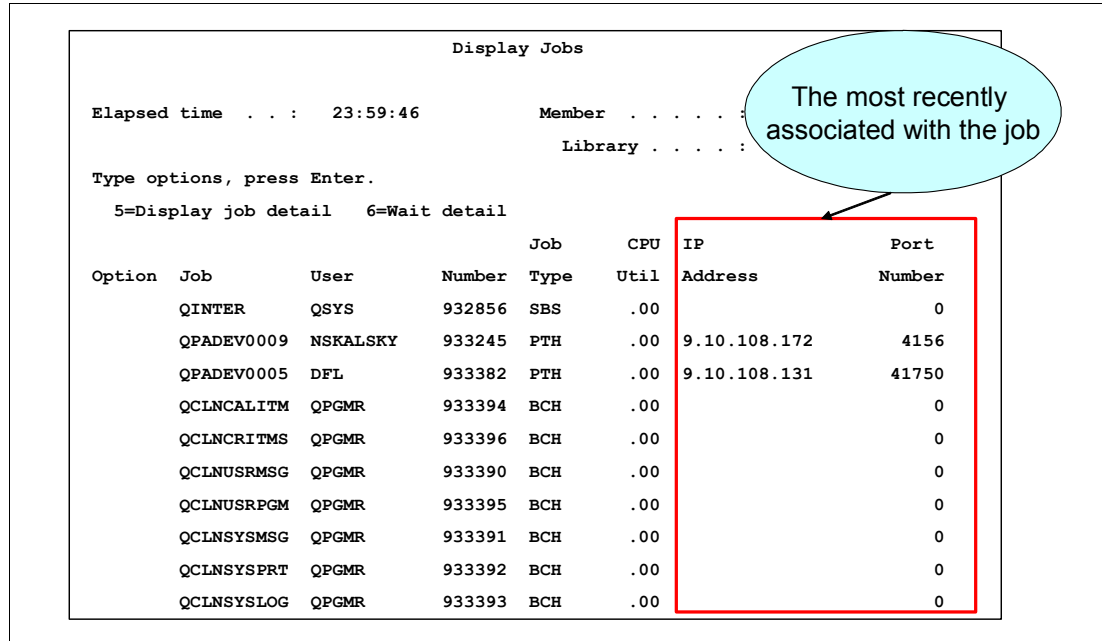


Figure 16-28 DSPPFRDTA: Display Jobs screen

16.9 Job Watcher

The Job Watcher has been updated to run on V5R4 while it includes the new function for V5R3, and V5R2 users as well.

The enhanced version includes new graph and legend support as well as user-defined query and graphs stored in the local PC.

To summarize the graph and legend enhancements:

- ▶ The legend shows fields already in the graph and available fields that can be included in a graph refresh.
- ▶ Legend fields can be dragged into, out of, and within the various areas of the graph. The fields in a stacked bar graph can be reordered if desired.
- ▶ Bar and line colors can also be set.
- ▶ X-axis and secondary Y-axis:
 - Redesigned.
 - Up to 30 lines can now be specified on the secondary Y-axis page.
 - Up to 3 fields can be used to build the X-axis label.
- ▶ Users can define the fields to be shown in the graph tip window (flyover) in addition to the fields for the bars and lines in the graph.

Figure 16-29 shows a graph and legend page sample.




Figure 16-29 Job Watcher: Graph and legend page sample

Note: Many of the graphics shown with Job Watcher contain many different colors, as illustrated in Figure 16-29. The colors can only be appreciated when viewing the graphs in color.

Enhancements summary for defined query and graphs:

- ▶ Major change: User-defined queries and graphs are now stored in a local database (MS mdb file) of the user's choice on the client PC:
 - Noticeable performance improvement
 - Now able to write a query or graph over data on one system and apply it to any system without having to move files around from one server to the other
- ▶ To avoid having to redo all query and graph definitions, the import interface allows a user to import definitions from either:
 - Server of their choice (from the iDoctor UDQ/UDG files in QUSRSYS)
 - Another iDoctor user-defined definitions database
- ▶ To use your own query or graph, right-click the **Job Watcher** component icon and select the **menu** option.
- ▶ User-defined reports: Access the import interface or set the current user-defined definitions database.

Figure 16-30 shows in the iSeries Navigator how to access user-defined queries and graphs that reside in the server or on the local workstation.

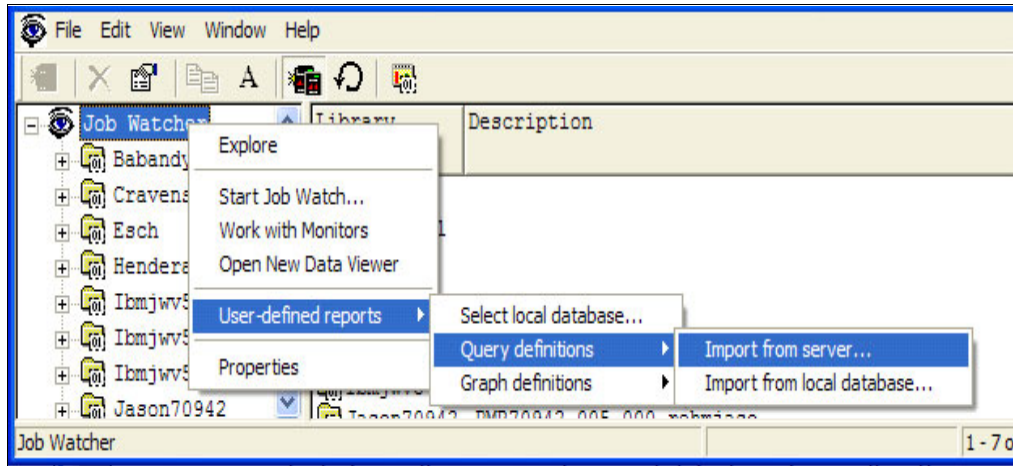


Figure 16-30 User query and graph access sample

16.10 IBM Systems Workload Estimator

Important: To assist you in sizing and capacity planning processes, consider using the IBM Redbook *Sizing IBM i5/OS Work on IBM System i5 Partitions*, SG24-6656, which contains examples of using the sizing capabilities of:

- ▶ IBM Systems Workload Estimator
- ▶ Midrange Performance Group (MPG) Performance Navigator
- ▶ BMC PATROL for iSeries - Predict

The IBM Workload Estimator (WLE) has been updated to support micro-partitions (shared processors) on the IBM eServer pSeries brand of systems around October of 2005, and now with the capability to support measured performance data from a POWER4 or POWER5 pSeries system (AIX 5.2 or 5.3).

Many of the screens within the Estimator have been changed to use the word *core* instead of *processor* when referring to the actual computing engine of the system. For example, the two processors per card configuration in the 520, 550, and 570 models are now termed Dual Core Modules (DCMs).

The WLE has also been updated to reflect the performance characteristics of IBM WebSphere Commerce for iSeries Version 5.6, as well as adding support for sizing the Learning component for IBM Workplace on i5/OS.

Also, WLE has been enhanced to support the ability to specify the target storage protection type when sizing with Performance Management (PM) eServer and existing workloads. This new feature allows you to size the disk capacity needed when changing the desired type of disk protection, like changing from RAID to Mirrored protection.

The System i5 feature #5709 I/O controller is now supported. In addition, a new storage attachment type of External has been added. This attachment type should be used if the intent of the estimation is to use an external storage subsystem (like the IBM TotalStorage ESS and DS6000 and DS8000™ models). Selecting this type of external storage attachment

instructs the Estimator to skip the checks for the available number of disk arms (bays) during the system selection process.

Users of PM pSeries can export historical performance data to the Estimator to size upgrades, server consolidations, or new workloads on pSeries systems. For more information about PM pSeries see:

<https://srm.raleigh.ibm.com/pmweb/en/pSeries/pseries.jsp>

There are also updates to the AIX 5.2 (APAR IY73893) and AIX 5.3 (APAR IY73894), which makes it possible to provide automatic collection of historical pSeries performance reports for use by the Estimator. This report provides an alternative to the PM pSeries. An Estimator-compliant XML report named `[hostname]_aixwle_weekly.xml` will be maintained in the systems `/etc/perf/` directory. Users should use the Restore Saved Estimation option to import this data into the Estimator for analysis. Refer to the local documentation file `/usr/lpp/perfagent/README.perfagent.tools` for more detailed information.

The Virtual I/O workload has been added to i5/OS on iSeries to account for the resources required to host Virtual LAN and Virtual Disk activity on behalf of an AIX or Linux workload. AIX and Linux workloads have been updated to allow you to specify the i5/OS partition that will host the virtual I/O. Processor, memory, and disk requirements are now estimated in the i5/OS partition for handling the I/O requests from the specified AIX and Linux workloads.

The File Serving workload for AIX and Linux has been updated to incorporate a *slider bar* as the mechanism of specifying file serving complexity. Users of this workload now have the ability to specify clients with multiple levels of complexity in a single workload.

The Domino workload has been updated with the latest sizing guidelines for Domino 7. The Workplace workload has been updated with the latest guidelines for messaging and brand new support for instant messaging.

The support added in Version 2005.3 for sizing external storage subsystems has been enhanced. A Download option has been added, which will generate a Disk Magic configuration file and download it on the Workload Estimator client. This configuration file can be opened from the Disk Magic tool and used to size external storage subsystem requirements for the workloads sized by the Workload Estimator. More information about Disk Magic can be found at:

<http://www.intellimagic.net>

The Estimator has been updated to use new IBM system brand names. References to pSeries have been replaced by System p5 where appropriate. References to iSeries have been replaced by System i5 where appropriate. The official name of the Estimator has now been changed to IBM Systems Workload Estimator.

The Estimator has been updated to include the latest information about the System i5 models that were recently announced, as well as the version/release of the i5/OS. All estimations by WLE are now based on the i5/OS V5R4 as the default operating system.

The Estimator now supports RAID-6 as a storage protection option. RAID-5 remains the default protection option.

The new Estimator has been updated to support new sizings with the System i5 feature #5737 I/O controller. This is one of the disk controllers that can run either under an IOP or without an IOP.

In addition, the Estimator now has several new enhancements for sizing disk requirements. Now an adjustable target disk busy rate in the range of 5–40% can be specified when sizing a

system. Users of the existing and generic workloads can specify the disk busy rate for the current system and size to a different target busy rate.

Additional improvements have been made to use an estimated read/write ratio when changing disk protection types so that sizing an upgrade from RAID to Mirrored or other combinations are more accurate.

Note: Changes to the disk sizing algorithms could result in a different number of disk arms being recommended than in previous versions of the Estimator. The changes reflect the use of additional performance factors that the Estimator previously did not consider. Also, an intermittent problem existed in the previous version of the Estimator, which could cause the number of disk arms recommended to be incorrectly high. This problem is corrected in the Estimator as of October 2006.

The Traditional workload has been renamed *Traditional OLTP* to better describe how this workload is used. Adjustments to memory and disk estimations have been made to better reflect OLTP environments of System i5 users.

The Estimator has consolidated the existing workloads from each system into one common workload. As a result, the screen navigation of this workload has changed slightly. Additional support has been added to allow upgrade sizings of existing Linux workloads on the System i5 or iSeries.

WLE now allows you to add new workloads or consolidate workloads from multiple systems onto a system that is no longer sold by IBM. After completing a System i5 estimation, a new option called *Retired Sizing* is available on the Estimation Results screen that lets you determine whether the workloads can run on a system that is currently owned, but is no longer sold. Supported systems include selected 8xx and 5xx models.

The Workload Developer component has also been updated with many enhancements. Developer is included in the download version of Workload Estimator.

As this book was going into the publication cycle at the end of December 2006, the following WLE enhancements became available:

- ▶ New IBM System models supported include:
 - Support for the most recent IBM System p models was added.
 - Support for the new IBM System i Telephony edition models was added.
- ▶ New System x and BladeCenter QUAD CORE models
Support for the latest System x and BladeCenter was added.
- ▶ i5/OS WebSphere workload integrates PMI data
The WebSphere Application Server workload for i5/OS was updated to optionally integrate with WebSphere Performance Monitoring Infrastructure (PMI) performance data. This enhancement allows you to more closely tailor the WebSphere workload to your application.
- ▶ File Serving workload updates for AIX and Linux
This workload has been updated with new performance information so that it is more representative of real world file serving.
- ▶ IBM Workplace workload updates for i5/OS
The workload has been updated with the most recent workplace performance information.
- ▶ Miscellaneous improvements to the Workload Developer

The Workload Developer component has also been updated with many enhancements. Be sure to read the What's New information found in the Help menu of the Developer. Remember that the Developer is included in the download version of Workload Estimator.

16.11 IBM Disk Magic

IBM Disk Magic for Windows is an easy-to-use external disks sizing tool that runs on a PC running a Microsoft Windows operating system. IBM Disk Magic for Windows is a product of IntelliMagic, licensed exclusively to IBM and IBM Business Partners. It may be used only by IBMers on IBM equipment or IBM Business Partners on IBM Business Partner equipment. Disk Magic cannot be left with customers for their own use.

The Disk performance modeling capability within Disk Magic for Windows is intended to provide a fast, straightforward estimate of likely performance of an IBM or non-IBM disk system. Disk Magic for Windows does not accurately model sequential (batch) workloads.

In Figure 16-31 you see the initial Disk Magic window with sizing for iSeries systems by using selected columns of data on the Performance Tools for iSeries reports. You can use the iSeries Navigator Printer Output folder (not shown here) to copy the files on the i5/OS output queue to a PC file. Each file must have the .txt file type for Disk Magic to find the file and extract data from it. Always review the text on the initial windows as Disk Magic reads each performance tools report. There may be specific actions called out to you.

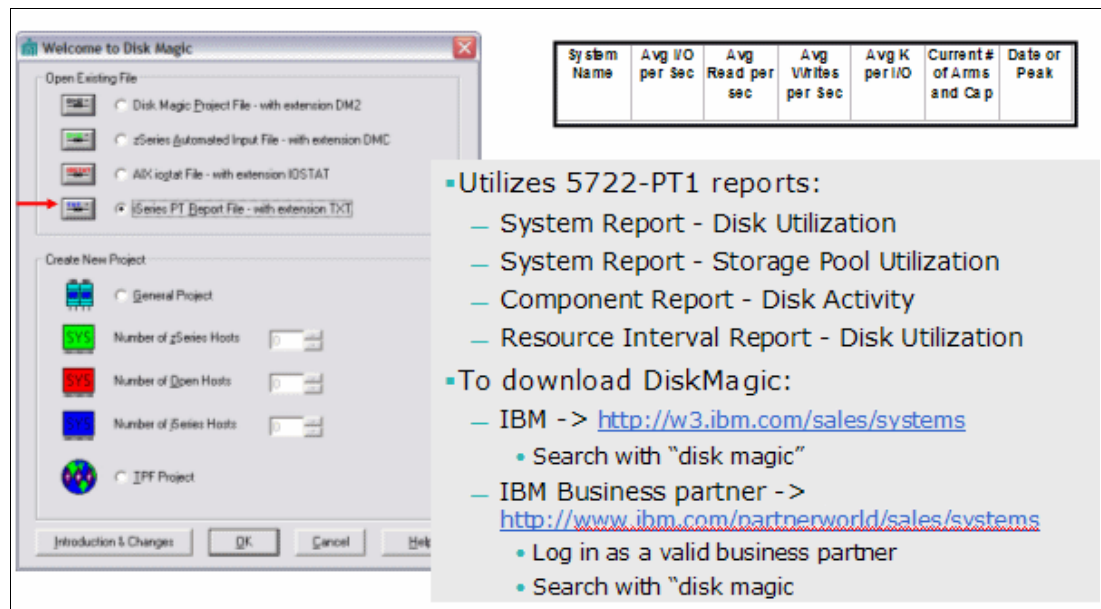


Figure 16-31 Disk Magic tool

16.12 Performance Management for System i5

Performance Management for System i5, formerly called Performance Management (PM) for iSeries and PM pSeries, gather non-proprietary performance data from your System i5 partition or an IBM System p partition and sends that summarized performance data to IBM on a daily, weekly, or specified schedule.

The data collected is converted into reports, tables, and graphs that show your partition's growth (trends) and performance calculations. You can set up PM for System i5 to send summarized performance data for each active partition to the IBM performance management Web site to include a system-wide view of each partition and all partitions together.

The PM for System i5 Web-based service automates the collection and summarization of your performance and capacity data. With this service, you are able to view graphical reports via a standard Web browser to identify how well the partitions and systems in your enterprise are performing.

This function is available for the OS/400, AIX, and Linux (PM for System p5 for Linux) at a no-charge and full-service (with fee) options for both the System i and System p family of systems.

PM for System i5 and PM for System p5 support highlights include:

- ▶ Input into IBM eServer Workload Estimator (WLE).
- ▶ Multiple partition support is enabled by the use of the IBM Director V5.10 server and agent products. The IBM Director Server must be active in an i5/OS partition and IBM Director agents in other i5/OS, AIX 5L, or POWER Linux partitions. Currently, only summary level CPU performance information is available for a Linux or AIX partition.

This support functions with i5/OS V5R3 or later and IBM Director V4.1 or V5.1 or later.

The IBM Director steps to set this up are documented in the iSeries Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp?topic=/rzahx/rzahxcollectdatapart.htm>

We show a figure (Figure 16-33 on page 447) later in this topic depicting this configuration.

- ▶ Cross-LPAR support.
- ▶ Manage the performance of a system with multiple LPARs.
- ▶ Summary and detailed information is available on individual LPARs and all LPARs together for a total system view.
- ▶ The Summary view gives you a snapshot, allowing you to view the dynamics of each partition relative to the available capacities in other partitions.
- ▶ Detailed information provides a series of graphs for further investigation and review. This includes areas such as:
 - Identifying possible need to utilize CoD processors
 - Tracking customer actual *interactive* utilization
 - Reporting disk arm utilization
- ▶ Consolidated view of all partitions.
- ▶ Graphs that make it much easier to understand where potential utilization thresholds may exist and the additional capacity that may be available to address the resource requirement capabilities.

We provide figures showing some of the PM for System i5 capabilities listed above. Since this is an overview book, we do not go into details on all of the figures shown. These figures are best viewed in color.

Figure 16-32 shows an example of no-charge summary information for a single system. The yellow shading indicates that a metric is approaching a guideline value (marginal), and a red shading indicates a metric at or above a guideline value.

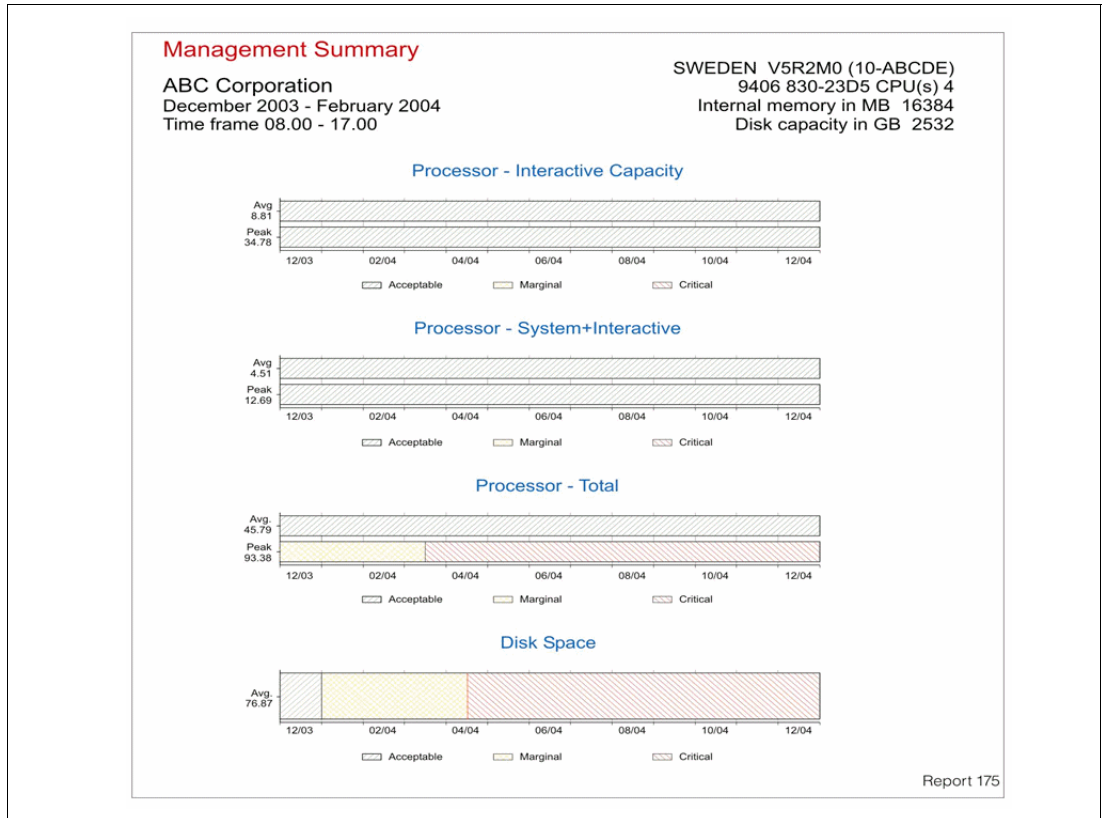


Figure 16-32 No-charge summary graphic example

Figure 16-33 shows a high-level representation of setting up PM for System i5 and IBM Director to collect performance data from multiple i5/OS, AIX, and Linux partitions to send all performance data to the PM for System i5 server. Then you can interface to the PM for System i5 Web site and show data for all partitions on a system.

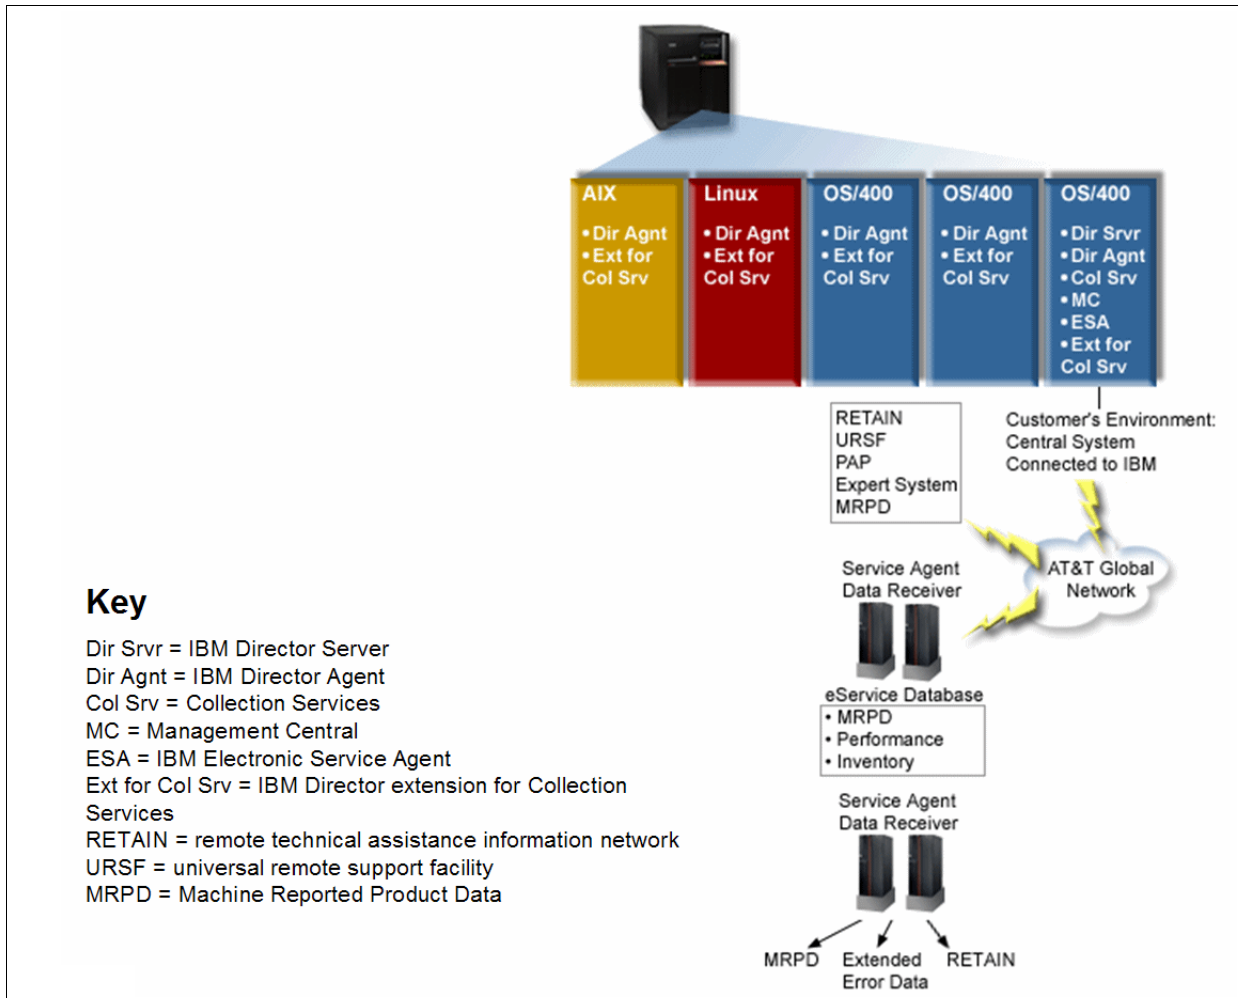


Figure 16-33 Setting up PM for System i5 and IBM Director high level overview

The following is a summary of the main set up steps:

1. Configure your IP network properly for all partitions on the same physical system.
2. Verify supported operating systems on each partition for which you want to collect performance data:
 - a. i5/OS, Version 5 Release 3 or later
 - b. AIX 5L Version 5.3 or later
 - c. Red Hat Enterprise Linux AS, Version 4.0, for IBM PowerPC^(R) or later
 - d. SUSE LINUX Enterprise Server 9 for IBM System p or IBM System i platforms

Consult the Web page listed below for additional information, There are some limitations in the range of metrics collected for an AIX or Linux partition.

3. Apply all Collection Services fixes to the partition that will act as your management server.
4. Install IBM Director for i5/OS Server on the i5/OS partition that you want to act as the management server.

5. Install the appropriate operating system IBM Director Agent on the partitions that you want to be managed by IBM Director Server.
6. Install IBM Director Console on the system (workstation) that you want to function as your IBM Director management console. Supported operating systems are AIX 5L, Windows, or Linux.
7. Install the Director extensions for Collection Services on the i5/OS management partition and distribute the Collection Services files from the management partition to the i5/OS partitions from which you plan to collect performance data.
8. On each Linux and AIX partition, install the Director extension for Collection Services.
9. In IBM Director Console, update the collection inventory on each partition by right-clicking the partition (managed system) and selecting **Perform Inventory Collection**.
10. Activate PM iSeries on the management server partition, which automates the start of Collection Services and then creates the database files during collection.

For more information about PM for System i5 and PM for System p5, refer to:

- ▶ Web sites:
 - Primary:
<http://www.ibm.com/systems/i/pmsystemi5>
 - For IBM director set up:
<http://publib.boulder.ibm.com/infocenter/eserver/v5r3/ic2924/index.htm?info/rzahx/rzahxcollectdatapart.htm>
- ▶ IBM Redbook: *A Systems Management Guide to Performance Management for i5 and p5 Systems*, SG24-7122

Figure 16-34 shows an example of one of the views of multiple partition PM for System i5 performance data. Again, the yellow color indicates the marginal or approaching a guideline value, and red indicates at or over a guideline value. The processor units shown in the upper partition window are duplicated in the Cross LPAR Summary window at the bottom of each bar graph.



Figure 16-34 Multiple partition PM for System i5 example

16.12.1 PMiSeriesIS tool

This tool is used by authorized IBMers and Business Partners to view and analyze collected PM for iSeries performance data. This tool securely connects to the Web via Virtual Private Network (VPN) as it runs on the PC client.

At a high level PMiSeriesIS enables:

- ▶ IBMers: Automatic access to accounts in their country
 - Can subdivide into smaller personal groups.
 - Access to worldwide data is available. This means that a representative can see customer systems that are installed worldwide.
- ▶ Business Partners: Access to customer data authorized by end-user only
 - Hardcopy customer data release form.
 - Customer electronic authorization process. This option is presented when viewing their graph on the Web

A customer explicitly authorizes access to PM iSeries performance data by system or partition to IBMers and Business Partners for:

- ▶ Recognizing resource utilization trends
- ▶ Capacity planning assistance
- ▶ Impact of adding additional workloads
- ▶ Upgrading or replacing a system
- ▶ System (Server) Consolidation

The PMiSeriesIS can query performance data that can be reviewed or used as input to the IBM Workload Estimator (WLE) tool. This performance data includes the following:

- ▶ Configuration information
- ▶ Contact information
- ▶ Upgrade history
- ▶ Utilization and growth detail
- ▶ Customer graphs

An example of customer charts and graphs is shown in Figure 16-35. This can be printed or sent via mail to the appropriate customer representatives.

Customer Name	Shift	CPU I	CPU T	DSK	# months	CPU T,a	CPU T,p	CPU H,a	CPU H,p	CPU I,a	CPU I,p	DSK used	RSP time
My Corp - Atlanta, Ga	1	12	0	0	12	72.75	99.51	34.59	52.52	25.84	45.60	87.00	0.63
My Corp - New York, NY	1	12	0	12	13	91.68	98.40	54.28	44.04	26.28	36.87	77.40	0.70
My Corp - Chicago, IL	1	12	0	1	7	81.68	98.04	48.44	56.57	32.28	37.82	78.65	0.62
My Corp - Paris, Fr	1	12	0	12	13	89.30	97.85	35.45	43.91	26.63	60.06	49.01	0.47
My Corp - Hong Kong	1	12	0	0	13	90.24	97.34	22.45	69.97	47.46	41.40	90.51	0.74
My Corp - UK	1	12	0	0	8	70.67	96.88	23.86	43.18	24.06	42.33	84.70	0.48
My Corp - Australia	1	12	0	12	13	71.42	96.80	41.55	55.82	32.49	55.46	73.30	0.63
My Corp - St. Paul, Mn	1	12	0	12	13	72.39	96.67	31.64	62.73	38.42	26.26	79.10	1.24
My Corp - San Deigo, Ca	1	12	0	2	13	49.26	96.62	30.11	31.52	18.10	26.94	71.90	0.24
My Corp - Seattle, Wa	1	12	0	0	13	76.87	96.20	24.61	45.14	14.87	78.78	81.87	1.27
My Corp - Vancouver, BC	1	0	0	5	4	66.99	96.11	22.84	74.82	39.79	23.51	60.31	1.14
My Corp - Germany	1	3	0	12	2	70.44	95.39	14.88	26.32	19.58	38.25	56.90	0.42
My Corp - Brazil	1	2	0	0	2	90.27	94.76	40.31	50.95	26.78	46.03	90.39	0.66

Figure 16-35 iPMSeriesIS performance data example

Since this is an overview book we do not go into details on the example shown. However, we note the following for the example chart:

- ▶ The Customer name field shows one customer with multiple world-wide locations (partitions or systems).
- ▶ The number of months (# months column) over which the averages in the succeeding columns were computed.
- ▶ CPU T,a: Total CPU utilization average for first shift (Shift column).
- ▶ CPU T,p: Highest peak CPU utilization average for first shift.
- ▶ CPU H,a: High priority jobs CPU utilization average for first shift.
- ▶ CPU H,p: High priority jobs highest peak CPU utilization average for first shift.

- ▶ CPU I,a: Interactive jobs CPU utilization average for first shift (Shift column).
- ▶ CPU I,p: Interactive jobs highest peak CPU utilization average for first shift.
- ▶ DSK used; Percent of disk storage used at the end of the reported period.

Note the table cells in red. This indicates that the value shown is higher than suggested guideline values.

16.12.2 Integrating PM for System i5 and the IBM System Workload Estimator

PM for System i5 data integration with the IBM Work Load Estimator (WLE) provides the following capabilities:

- ▶ Size needed upgrade or replacement system.
- ▶ *Pass* all customer PM configuration, utilization, and growth data to WLE with a click of the mouse.
- ▶ Variable Time Horizons — up to 36 months.
- ▶ Size traditional upgrades, server consolidations, Linux, Domino, WebSphere, add on applications to the customer's existing workload.
- ▶ Ability to pass all partitions, regardless of OS type (Linux, i5/OS, OS/400, AIX 5L) to WLE with a single pass.
- ▶ Save, print, or create PDF of recommendation.
- ▶ Input to LPAR sizing.

Authorized users of standard PM for System i5 and authorized IBMers and business partners using the PMiSeriesS tool can link to the Workload Estimator.

16.13 Web Performance tools

During October 2006 several *Web enablement* enhancements were announced. Two of them, the Web Performance Monitor and the Web Performance Advisor, are useful tools for non experts in Web application performance measurement and performance analysis.

For more information about these Web monitor and advisor tools refer to Chapter 14, "i5/OS-based Web enablement enhancements" on page 337.

16.14 Non-IBM sizing tools

There are other non-IBM tools that offer printed and graphical interfaces on collected i5/OS performance data and perform sizing exercises to determine the next System i configuration and sizing functions. In this section we only discuss the sizing function products.

16.14.1 Midrange Performance Group (MPG) Performance Navigator

Performance Navigator is a graphical performance reporting and sizing tool running on the PC. Versions 11, 11.1, and now 11.2 can be used to give you:

- ▶ Performance analysis graphs and reports
- ▶ Extensive single system and LPAR What If (sizing/capacity planning) capabilities
 - Quick changes = quick graphs

- Does not include transaction throughput or response time predictions

The latest version and enhanced version include:

- ▶ The January 2006 announced systems and disk models.
- ▶ New What If capabilities:
 - To show the impact of enabling January 2006 System i5 520 Accelerator.
 - To replicate existing workload (includes DASD).
 - Has a new button to “Set to Current Performance” (discard in-progress what-if changes).
 - Buttons to round processor units up or down.
 - Automatically calculate Total Software/SWMA License (i5/OS and AIX).
 - Room for growth calculator showing annualized CPU growth%.
 - Configuration section enhanced for DASD.
 - Server consolidating/LPAR support for i5/OS (CPW and MCU sizing algorithms), AIX, Linux, workloads.
- ▶ Set graphs to default scaling to 100% utilization.
- ▶ CPU graphs show uncapped CPU capacity (5xx hardware).
- ▶ New Service Level graphs for key performance metrics, including high priority jobs CPU utilization.
- ▶ New Job Wait graphs.
- ▶ New historical memory graphs compare Total Memory versus Total Faults.
- ▶ Additional Current Day graphs include TCP/IP, Journeying, Lock Conflicts, Group faulting, disk I/O rates, and more).
- ▶ Support of UNIX Capacity Planning on iSeries (requires Power Navigator License).

For more information about the Performance Navigator tool visit their Web site at:

<http://www.mpginc.com>

Important: To assist you in sizing and capacity planning processes, consider using the IBM Redbook *Sizing IBM i5/OS Work on IBM System i5 Partitions*, SG24-6656. This book contains examples of using the sizing capabilities of:

- ▶ IBM Systems Workload Estimator
- ▶ Midrange Performance Group (MPG) Performance Navigator
- ▶ BMC Software (BMC) PATROL for iSeries - Predict

16.14.2 BMC PATROL for iSeries - Predict

The BMC PATROL for iSeries - Predict is a capacity planning tool for the System i5 models. This tool uses the existing performance data gathered by the performance collection services and commands within the i5/OS partition.

BMC Patrol Predict has extensive capacity planning capabilities that include transactions throughput and response time predictions. On the other hand, this tool has limited graphics and restricted disk configuration.

Latest enhancements to this tool include:

- ▶ Ability to show and include in calculations the number of processors assigned to the partition, as well as the maximum processors available to the partition and total processors potentially available on entire system.
- ▶ Includes impact of L2 cache rebuild.
- ▶ Transaction CPU utilization can be made globally.
- ▶ For capped or uncapped partitions, processor utilization predictions are capped at 100%. Uncapped partition note added to output.
- ▶ Disk and RAID usability enhancements including deletion of selected disks and RAID sets and group copy/rename, and changing between RAID and non-RAID.
- ▶ New Model Summary (under Computer Reports) and Period-Model Summary reports show model-wide predicted values on a single report. Included in the reports are:
 - CPU model, memory size, and processor information.
 - More detailed disk information.
 - Statistics are broken out by total, interactive, and non-interactive for CPU utilization, queue length, throughput, and average response time.
 - All values in the reports that are either above a guideline/threshold or show the impact of cutback are highlighted in bold.
- ▶ New response time calibration (measured versus created model).
- ▶ Predicted partition CPU utilization now models the impact of cross-processor overhead. For example, two 4-processor partitions will show more total capacity than a single 8-processor partition.

For more information about this product, visit their Web site at:

<http://www.bmc.com>

Important: To assist you in sizing and capacity planning processes, consider using the redbook *Sizing IBM i5/OS Work on IBM System i5 Partitions*, SG24-6656, which contains examples of using the sizing capabilities of:

- ▶ IBM Systems Workload Estimator
- ▶ Midrange Performance Group (MPG) Performance Navigator
- ▶ BMC Software (BMC) PATROL for iSeries - Predict



i5/OS and IBM Director

In this chapter, we provide an overview of IBM Director 5.10 and 5.20 capabilities from an i5/OS viewpoint. IBM Director 4.1 was the initial version that supported multiple operating systems, including running under i5/OS (V5R3). IBM Director 5.10 and 5.20 run on either i5/OS V5R3 or V5R4.

The contents of this chapter focus on IBM Director 5.10. From an i5/OS viewpoint, IBM Director 5.20 enhancements are primarily in the ease of installation and ordering areas. These 5.20 enhancements were announced during November 2006. An i5/OS view of these enhancements is summarized at the end of this chapter.

17.1 IBM Director 5.10

This provides heterogeneous multi-system management so that administrators can view and track the configuration of different platforms in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

In Versions 4.1 and 4.2 the product was called IBM Director Multiplatform to illustrate its capabilities across multiple hardware platforms and operating systems. In V5.10 these multiplatform capabilities are expanded, but the word *multiplatform* has been dropped from the product's name.

IBM Director enables managing i5/OS, AIX, Linux, and Windows operating systems, and eServer HW platforms from a single GUI:

- ▶ Inventory
- ▶ Monitoring
- ▶ Alerting
- ▶ Group management
- ▶ Automated reactions
- ▶ Process management
- ▶ Remote session
- ▶ File transfer
- ▶ Real time diagnostics

IBM Director for i5/OS V5.1 runs under i5/OS V5R3 or V5R4, as a Director server or agent. Supported agent software includes one or more of the following:

- ▶ Linux on POWER: Red Hat Enterprise Linux (RHEL) 4.0 or SUSE Linux Enterprise Server (SLES) 9 and 10
- ▶ AIX V5L 5.2 or 5.3
- ▶ Windows x86 versions supported with integrated System x solutions (IxS, IxA, iSCSI)
- ▶ Linux RHEL 4.0 or SLES 9 with System x and IXA hardware and iSCSI

Important: During November 2006, IBM Director Version 5.20 was announced. From the viewpoint of IBM Director Server running on i5/OS, V5.20 enhancements are essentially in the packaging and ease of installation. Therefore, in this chapter we describe IBM Director V5.10 on i5/OS capabilities and then summarize the i5/OS-oriented V5.20 enhancements.

17.2 IBM Director overview

The base IBM Director file transfer function is *real time* only. It is not included in the base scheduling function of IBM Director. To have full automated scheduling and tracking of functions such as file or directory transmission, you must install the additional cost extension Software Distribution (Premium Edition). We discuss IBM Director extensions in 17.6, “IBM Director 5.10 extensions” on page 471.

In the IBM Director console window shown in Figure 17-1 on page 457 you see three panes of folders, objects, and tasks you can perform:

- ▶ Groups
- ▶ Managed objects (new term used in V5.10)
- ▶ Tasks

The three panes shown are the *traditional look* of the console. In 17.4, “IBM Director V5.10 Console” on page 461, we discuss some V5.10 customizing, which you can utilize to tailor the look and feel of the IBM Director console interface to what you want.

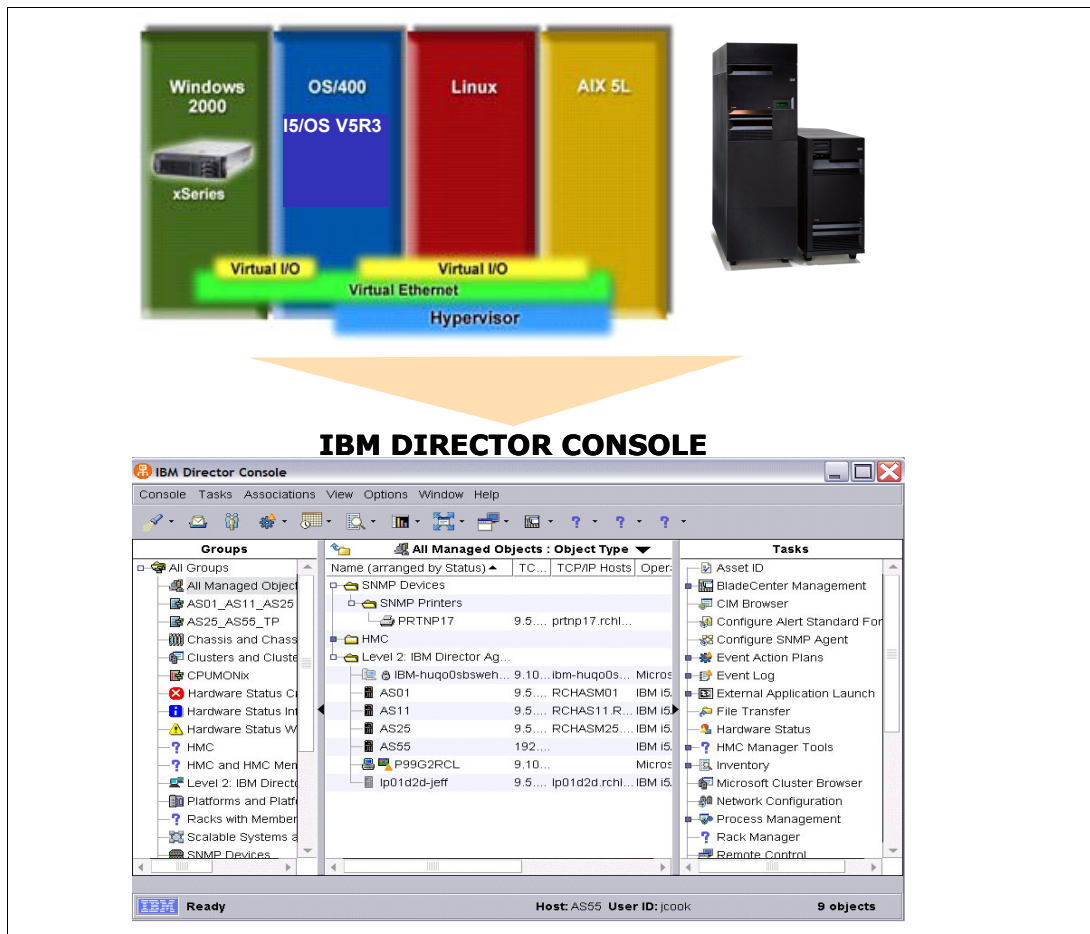


Figure 17-1 Director overview

In the Managed Objects center pane we show an example of the following objects in our network, which can be managed:

- ▶ SNMP devices, such as LAN printer PRTNP17
- ▶ i5/OS partitions, such as AS01, AS11, AS25, and AS55
- ▶ PC workstation running a supported Windows operating system, such as P99GRCL
- ▶ Hardware Management Console (HMC), basic functions

For information beyond the content of this chapter, there is lots of good IBM Director information in IBM Redbooks and the IBM Systems Software Information Center. See 17.9, “References” on page 481.

17.2.1 IBM Director topology

The IBM Director has several major components. To use the Graphical User Interface (GUI) for IBM Director server functions (the most commonly used interface), you need at least one IBM Director console within the network of managed nodes. You can have multiple systems with the IBM Director console installed, but this requires planned, coordinated, and secured use by more than one IBM Director console.

The IBM Director Console is a Java-based implementation. A Management Console is connected via network (normally using TCP/IP) protocol with the management server and all managed systems. The IBM Director console can operate under AIX, Linux, or a Windows operating system.

The management server is installed in a central location. This server is the *manager* from the IBM Director point of view for all IBM Director information collected from managed systems. Most functions are requested by the IBM Director console to the IBM Director server. The specific function could be performed by the server for a specified managed node.

The server contains all collected inventory information. The IBM Director console connects to the server to define actions and report inventory and events that have been defined.

Note: Use of the i5/OS V5R4 support of integrated xSeries hardware running supported operating systems and connected over the new integrated Small Computer System Interface (iSCSI) requires IBM Director Server installed on the serving i5/OS partition. IBM Director performs activation and deactivation of the xSeries server directed by i5/OS support. For these i5/OS-controlled functions, no IBM Director Console interface is required.

For more information about iSCSI support see 5.2, "System x and BladeCenter support via iSCSI" on page 121.

Agents

IBM Director 5.10 provides three different levels of managed systems. The agents levels (0, 1, 2) have different software requirements and provide a set of functions that correspond to the level number:

- ▶ Level 0 is a base function that needs no IBM Director software installed.

Note: For an IBM Director Agent set of functions, i5/OS is supported as a Level 2 endpoint with limited support for L0. With SSH Portable Utilities installed (IBM Portable Utilities for i5/OS, 5733SC1), the Director Agent may be distributed from an IBM Director to a L0 i5/OS partition (without IBM Director installed) to update it to a Level 2 agent.

The 5733-SC1 LPO contains the OpenSSH, OpenSSL, and zlib open source packages, which were moved to i5/OS using the i5/OS PASE runtime environment. 5733-SC1 requires i5/OS V5R3 or later, and also requires that i5/OS Option 33 (i5/OS PASE - Portable Solutions Application Environment) be installed.

For more information on IBM Portable Utilities for i5/OS, 5733SC1 refer to 6.7, "Portable Utilities for i5/OS" on page 144.

- ▶ Level 1, which requires a base level of IBM Director agent software to be installed.
- ▶ Level 2 installs the full IBM Director agent software, which enables the highest IBM Director agent functions.

Each level agent can be upgraded to the next higher level agent up to level 2. See Figure 17-3 on page 459 for a simple network graphic.

Managed systems, which are also referred to as managed nodes, can be workstations or full systems. Nodes can be as simple as SNMP devices (such as printers) to mainframes. In Figure 17-2 you can see that several different operating systems are supported.

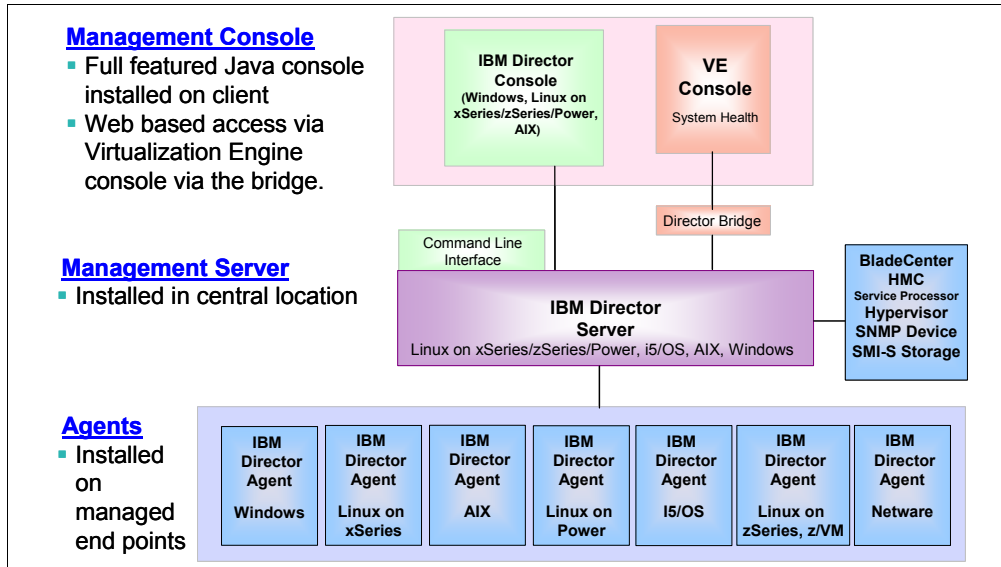


Figure 17-2 IBM software components

We refer to this as the group of managed SNMP devices, HMC devices, and IBM Director agents managed from a single IBM Director server as the network.

In the context of being managed, IBM Director specifically groups managed nodes into *classes*, such as agent levels, SNMP devices, and HMC devices.

Figure 17-3 shows a simple topology map.

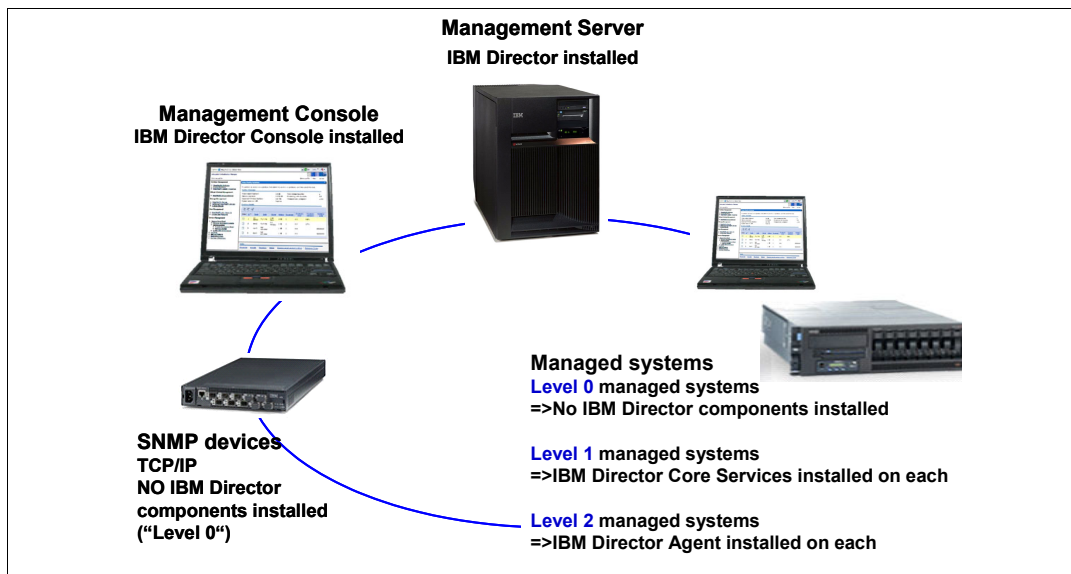


Figure 17-3 Simple topology map

17.3 Reducing complexity, delivering value

IBM Director product has been available for several years prior to Versions 5.10 and 5.20. Version 4.n brought the multiple operating systems support. Previous to Version 5.10, improvements allowed you to get started faster, and accomplish more in a shorter period of time. This helped to reduce training costs for IT staff.

The open design extends functionality to meet your specific needs. You were able to use a single tool to manage different server platforms and non-IBM systems. You could also take advantage of the latest industry standards.

The IBM Director integrated tools help reduce costs with a consistent, single point of management, which complements your existing IT investment.

Release 5.10 continues these enhancements by further reducing complexity. The main IBM Director 5.10 objectives were:

- ▶ Simplification to improve usability (project simplify).
- ▶ New Main Console Interface. Move to familiar paradigm view (Explorer, Web).
- ▶ Simpler wizard interface for setting up monitors and notifications.
- ▶ Reworked online documentation.
- ▶ CI-162 accessibility compliance.
- ▶ Lighter Footprint® (deployment, support for Level 0, 1, 2 systems).

Our Partners in Management program provides customers with more choice and management capabilities:

- ▶ Platform Management for eServer
- ▶ Cross-series hardware management
 - Added management servers for AIX, power Linux, and zSeries® Linux
 - Added agent support for zSeries Linux
- ▶ Support for the industry standard Storage Management Initiative Specification (SMI-S) for discovering and managing storage devices and others like CIM, SNMP, and SLP
- ▶ Improved hardware alerting
- ▶ z/VM® Center Management Extension to manage zSeries Linux virtual machines
- ▶ HMC Director Extension for pSeries and iSeries (available 2Q 2006)
- ▶ Broad platform support as a key component of IBM Virtualization Engine

The integrated toolset offers an improved Developers Kit to facilitate third-party integration:

- ▶ Improved performance and usability for Director's command line
- ▶ Tivoli Upward Integration: seamless integration of TEC/Director
- ▶ MS Operation Manager (MOM) upward integration
- ▶ Enabling TEC event action in Director Management Server
- ▶ Expressing SDK and formalized SDK
- ▶ Additional extensions
- ▶ Solutions enablement for both internal and Independent Software Vendor (ISV)/Independent Hardware Vendor (IHV)

Table 17-1 lists the feature improvement matrix.

Table 17-1 Generalized IBM Director 5.10 feature improvement matrix

Feature	Description	Benefits
Streamlined user interface	Clean, customizable console. User decides which tasks and what information is displayed.	Easy to get started and improved productivity by accomplishing more in a shorter period of time.
Broader eServer platform coverage	New management servers for AIX, Linux on Power, and Linux on zSeries. New agents for Linux on zSeries.	Reduces number of platform-specific management tools and reduces the skills required to manage your eServer environment.
Reduced agent footprint	Three levels of agents including agent less or zero footprint (Level 0), core hardware services (Level 1), and full IBM Director Agent (Level 2).	Ease of deployment and discovery saves time and resources. More customer choice regarding the level of hardware management desired.
Software health check (for x Series)	Right-click Update to perform software health check. IBM Director creates a dynamic group of systems requiring that update.	Saves time trying to determine which systems are in need of an update. Helps keep systems running optimally with the latest system software.
Software Developers Kit (SDK)	Tools and documentation to facilitate developing and integrated extension to IBM Director.	Facilitates the creation of an IBM Director ecosystem of extensions to provide customers more choice and capability from a single, integrated IBM Director console.
Event action plan wizard	Easy-to-use, step-by-step guide to let IBM Director know what you want to monitor and what action you want to take in response to an alert.	Reduces the time it takes to be productive and enables you to accomplish more in a shorter period of time.
BladeCenter chassis configuration manager	Provides integration point for BladeCenter subsystem configurations, allowing you to create and read configurations for chassis component devices and broadcast configurations to multiple chassis.	Increased productivity and time savings by making it easier to configure BladeCenter chassis and subsystems.
Virtual server deployment and Server complex	Creation and grouping of virtual servers under z/VM and deployment of Linux systems.	Fast, easy, and repeatable deployment of new virtual servers under z/VM.
Microsoft Operations Manager (MOM) integration	No charge module that surfaces IBM Director hardware management information into Microsoft Operations Manager.	Protects and augments investment in MOM deployment and training with complementary systems management information.
Automated inventory change control	Monitors changes in inventory and can alert when there is a change.	More accurate and timely asset management by providing notification of missing systems or components or systems that have been upgraded.

17.4 IBM Director V5.10 Console

Figure 17-4 on page 462 shows an enlarged traditional view of the IBM Director Console so you can see in greater detail the content of the three different window panes. Compare this to

Figure 17-5 on page 463, which shows you a window example of the new in V5.10 customizable IBM Director Console interface.

You see the three panes of folders, objects, and tasks you can perform, as shown in Figure 17-4:

- ▶ Groups
- ▶ Managed objects (new term used in V5.10)
- ▶ Tasks

In the left pane you see IBM director groups. There are many default groups, and you can create your own groups. The groups with ASnn in the name are user-defined groups.

All managed objects are shown as selected and control what is shown in the middle pane.

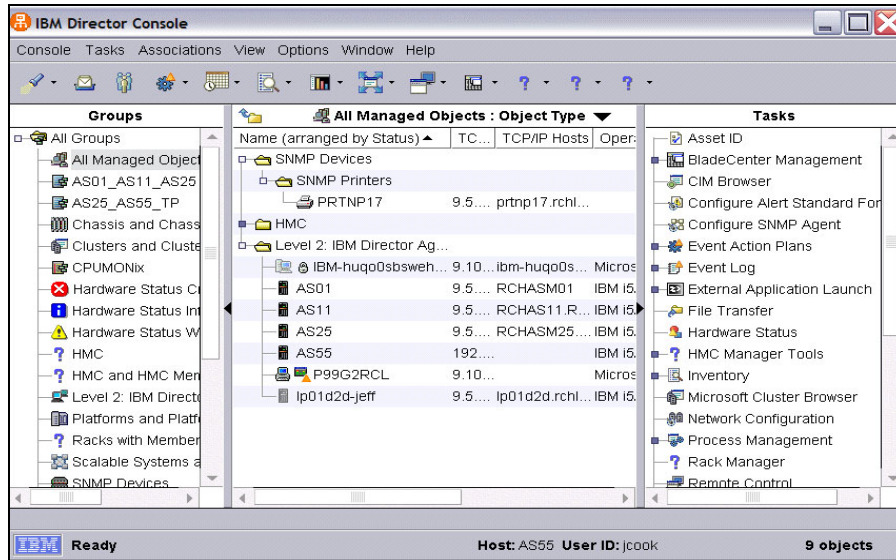


Figure 17-4 Traditional view

In the managed objects central pane, the objects are grouped according to IBM Director type, for example, SNMP devices, Level 2 Agents (Windows XP workstation and i5/OS partition agents), and a Hardware Management Console (HMC).

In the right pane you see a large subset of the many tasks that can be performed either on the groups shown or on the individual managed objects shown in this all objects pane.

With IBM Director 5.10 Console, you can have the same window view of previous releases or a more customized one. In this example you are using a new view. What in the previous release was the middle pane now fills an entire window. With 5.10 you can:

- ▶ Customize columns, sortable by column heading.
- ▶ Navigate between groups by drop-down.
- ▶ Use an optional group pane.
- ▶ Show tasks all in menu bar.
- ▶ Show common tasks in the toolbar, which is fully customizable.
- ▶ Use the optional task pane.
- ▶ See task errors before launching.

You can also have the older pre-V5.10 three-pane window by using the left mouse and the arrow heads highlighted in green, to widen or narrow the associated pane.

Figure 17-5 shows the new interface.

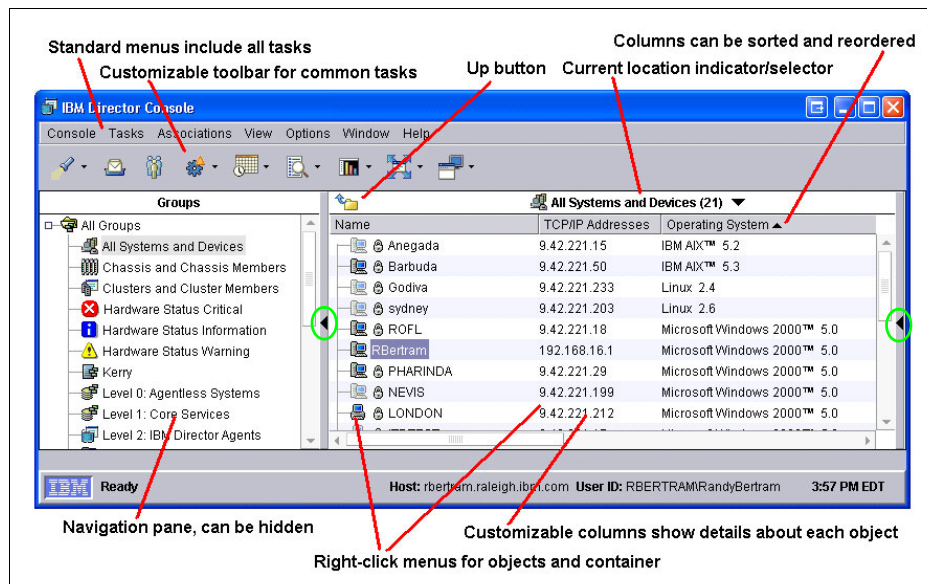


Figure 17-5 New user interface

17.4.1 Starting and stopping

Starting the IBM Director Server under i5/OS includes starting the Java Virtual Machine (JVM) and other setup functions. This could take a few minutes even on a 3000 CPW partition.

- ▶ Job QCPMGTSVR is the primary i5/OS IBM Director Server or Agent job, depending on which was installed.
- ▶ Job QYCMCIMOM is the open industry Common Information Module job under i5/OS responsible for translating information received by or sent by QCPMGTSVR between i5/OS and other operating systems. This job exists on i5/OS V5R3 and V5R4 solely for the use of IBM Director functions. In releases after V5R4, other system functions may use the functions available through this job.

The Common Information Model (CIM) is a standard developed by a consortium of major hardware and software vendors (including IBM) called the *Distributed Management Task Force (DMTF)* as part of the Web Based Enterprise Management (WBEM) initiative. WBEM includes a set of standards and technologies that provide management solutions for a distributed network environment. Interoperability is a major focus of WBEM, and using WBEM technologies can help you develop a single set of management applications for a diverse set of resources.

CIM provides a model for describing and accessing data across an enterprise. CIM comprises both a specification and a schema. The specification defines the details for integration with other management models, whereas the schema provides the actual model descriptions.

CIM on IBM i5/OS V5R3 or V5R4 includes:

- Instrumentation for server resources on the system called *providers*.

The providers, which are based on a subset of the standardized CIM classes, gather data on a system. CIM clients can work with this data by accessing the providers through the Common Information Model Object Manager (CIMOM).

- An open source implementation of the CIMOM called *Pegasus* (Version 2.5) that manages communication between clients and providers. The CIMOM also provides several management functions, including security, and a set of commands that provide configuration and management functions to administrators.
- A schema (CIM schema Version 2.9) that defines an information model for representing systems management functions.
- ▶ Job QYMAP2MIB provides mapping of information contained within a Management Information Base (MIB), which is an architected data structure that defines what is obtainable from a device and what can be controlled (turned off, on, and so on).

A MIB is implemented according to the industry-standard basic Simple Network Management Protocol (SNMP). SNMP is a widely used network monitoring and control protocol. Data is passed from SNMP agents, which are hardware or software processes reporting activity in each network device (hub, router, bridge, and so on) to the workstation console used to oversee the network. The agents returns information contained in the MIB. Originating in the UNIX community, SNMP has become widely used on all major platforms.

The IBM Director server can manage SNMP devices, based upon the MIB information exchanged. This is at a level less than an IBM Director Agent Level 1 or Level 2. Operating systems can provide this base SNMP support as well as, typically, network printer devices.

17.4.2 Discovery of managed objects

You have the facility to configure discovery settings for all levels of IBM Director systems, SNMP devices, BladeCenter Chassis, and Physical Platforms. Advantages include:

- ▶ Automatically secure newly discovered agents.
- ▶ Automatically discover agents that contact the Director Server.
- ▶ Support for SNMPv1, SNMPv2c, and SNMP v3 devices.
- ▶ Enable renaming of BladeCenter Chassis and physical platforms based on available parameters.

Managed objects types that can be discovered:

- ▶ BladeCenter chassis
 - Represents entire BladeCenter chassis.
 - Discovery, management, and some inventory performed out-of-band via management module.
 - BladeCenter Management and BladeCenter Configuration are customized versions of Server Management and Management Processor Configuration.
- ▶ Physical or virtual platforms (Blade servers or virtual machines)
 - Represents the physical server enclosure that can host an operating system, for example, discovery of HS20/HS20 processor blades or xSeries x440 server prior to operating system install
 - Used by the RDM extension to deploy an operating system to physical platforms or operating system containers
- ▶ Director systems (deployed operating systems)
 - Represents the operating system running within a physical platform.

- Today, this is discovered and managed through the agent installed on the operating system.
- ▶ **SNMP devices**
 - Represents SNMP devices (switch modules, printers, and so on)
 - Support for SNMP Version 1, 2, 3
- ▶ **HMC devices (2Q 2006)**

Table 17-2 shows Inventory collected for each supported platform.

Table 17-2 Inventory collected by platform

Director inventory query	Linux/ Windows xSeries	Power Linux iSeries	Power Linux pSeries	Linux on zSeries	AIX, I5/OS
Base Inventory via operating system: <ul style="list-style-type: none"> ▶ Disk and logical drive ▶ Installed memory, processor ▶ Installed patches and packages ▶ Network adapter and IP address ▶ Operating system/system information ▶ Physical enclosure 	Yes	Yes	Yes	Yes	Yes
SNMP Agent configuration	Yes	Yes	No	Yes	No
Serial number information	Yes	Yes	Yes	Yes	Yes
Asset ID/personalized data	Yes	Yes	Yes	Yes	Yes (No Asset ID)
Lease/warranty Information	Yes	N/A	N/A	N/A	N/A
FRU service numbers	Yes	Yes	Yes	Yes	Yes
Memory modules	Yes	No	No	No	No
Firmware	Yes	Yes	Yes	Yes	No
BIOS details	Yes	N/A	N/A	N/A	N/A
Processor details	Yes	Yes	Yes	Yes	No
System board configuration	Yes	Yes	Yes	Yes	Yes
Cache	Yes	No	No	No	No
Port connectors	Yes	Yes	Yes	Yes	No
Partition, geographic information	Yes	Yes	Yes	Yes	Yes

The shaded areas in Table 17-2 are new with 5.10.

17.5 Controlling your system

IBM Director offers several functions that help manage and control supported managed systems and SNMP devices. We describe the major control function and associated components in this section:

- ▶ Event action plans
- ▶ Job scheduler
- ▶ Process management
- ▶ Resource monitors

Most of these resources and controls are applicable over multiple operating systems and hardware platforms. However, some detailed action plan, process, or resource may be unique to a particular operating system or hardware platform.

17.5.1 Event action plans

You use event action plans to specify actions that are performed in response to events that are generated by a managed object. Managed objects include, but are not limited to, Level 0, Level 1, and Level 2 managed systems; SNMP devices; BladeCenter management modules; platforms; and switches.

An event action plan is composed of two types of components:

- ▶ One or more event filters, which specify event types and any related parameters
- ▶ One or more event actions, which occur in response to filtered events

You can apply an event action plan to an individual managed object, several managed objects, or a group of managed objects.

By creating event action plans and applying them to specific managed objects, you can be notified by e-mail or pager, for example, when a specified threshold is reached or a specified event occurs. Or you can configure an event action plan to start a program on a managed object and change a managed-object variable when a specific event occurs. You can use any event, including process-monitor and resource-monitor events, to build an event action plan.

An event action plan includes specifying:

- ▶ Event filters
- ▶ Event actions
- ▶ Logging options

IBM Director provides many predefined *filters*. Some of the event filter capabilities worth noting include the *simple event*. This is the default event filter. These predefined filters are useful for creating generic event action plans.

Several of these filters use the severity of the event as criteria to allow events to pass through. The *critical events* filter is concerned only with events of critical severity. The *fatal events* filter is concerned only with events of fatal severity, and so on.

The *all events* filter passes through any events that occur on any managed system. Carefully consider the use of the supplied simple event filters, as many of these are quite broad in what events will be caught, potentially leading to numerous alerts being generated. For example, there could be a critical CPU utilization event, and a critical storage event, and if both of those systems have an event action plan that filters for critical events, then the defined action would occur for both.

Duplication events allow for duplicate events to be ignored, in addition to the options given by the simple event filters.

Exclusion events allow you to exclude selected event types, in addition to the options given by the simple event filters.

A *Threshold event* allows you to select of an interval or count threshold that must be met in addition to the options given by the simple event filters.

The *Event Action Plan Builder wizard* simplifies the process of creating an event action plan and applying it to a group or system. The wizard starts with builder event filters for classes of events, such as CPU utilization, and helps you build a streamlined event action plane. You can also use this wizard to create and later modify event action plans built with wizard. Figure 17-6 shows the Event Action Plan wizard.

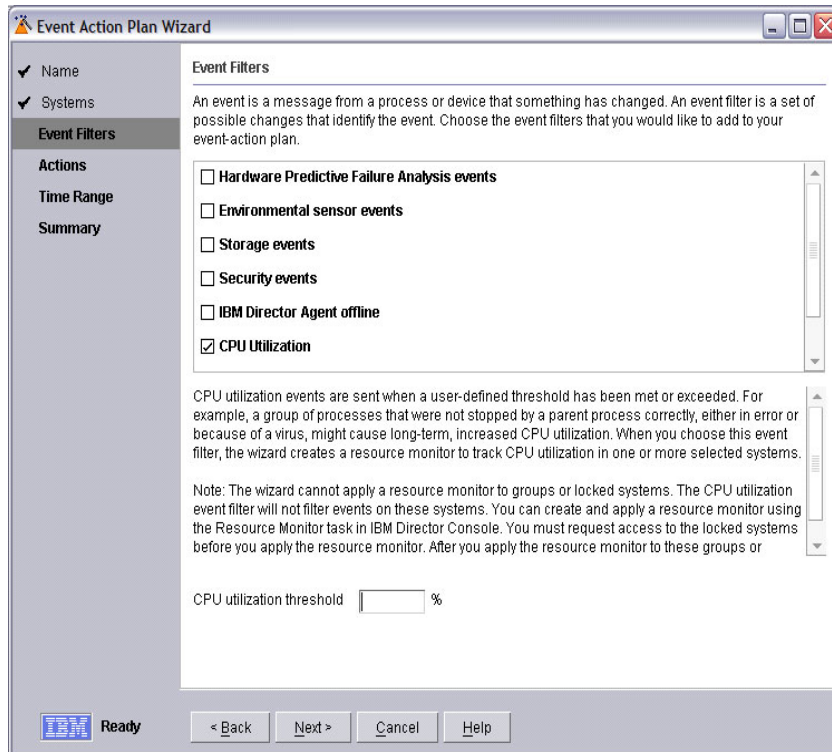


Figure 17-6 Event Action Plan wizard

The Event Action Plan wizard is ideal for getting you started with defining event action plans to get you started with automating control of supported events. However, the wizard does not support the entire set of event action plan options you can choose to use.

Figure 17-7 gives an example of *logging of all events*, the default. You can see the event associated with the *ticker tape* message shown on the IBM Director Console.

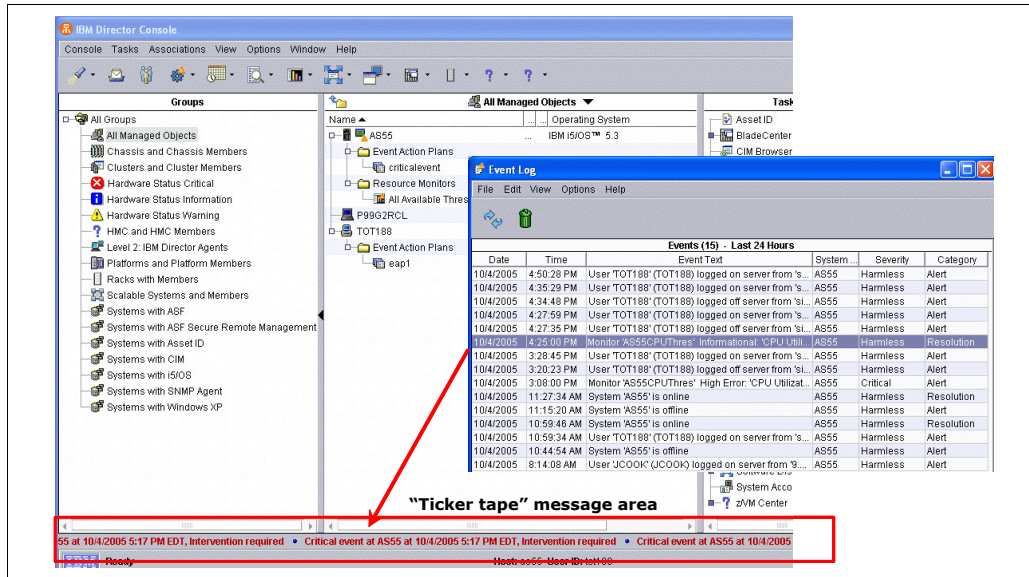


Figure 17-7 Event action plan example

In this example we displayed the detailed event log entries so you can associate the specific event with the ticker tape message action we specified.

17.5.2 Job Scheduler

You can use the Job Scheduler option included in the base IBM Director support, to run a single noninteractive task or set of noninteractive tasks at a scheduled time, once or repetitively. (Only noninteractive tasks, which are defined as tasks that do not require any user input or interaction, can be scheduled.)

You can specify an exact date and time you want the task to be started, or you can schedule a task to repeat automatically at a specified interval. Scheduled tasks are referred to as jobs. See Figure 17-8.

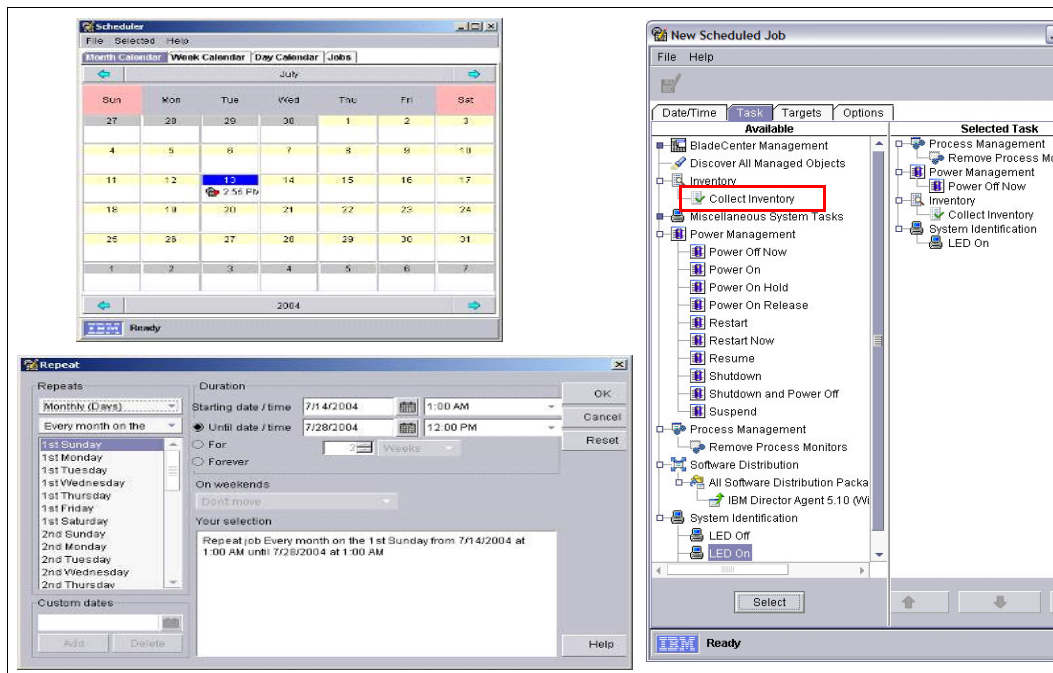


Figure 17-8 Job Scheduler windows

The window on the right shows the task option where you specify the tasks you want to schedule. Though not selected in our example, one of the tasks you would typically schedule repetitively would be the collect inventory task.

The lower left window shows the Repeat window showing when to schedule the task. The upper left window shows the calendar. You can right-click a day to start the definition. You can view your scheduled jobs a in monthly, weekly, or daily calendar. Jobs can be run daily, weekly, monthly, or on a custom basis. The alerts can be sent based on system and job success or failure. You can view an execution history of previously scheduled jobs, and re-execute previously scheduled jobs.

17.5.3 Process Management

The Process Management window (Figure 17-9) allows you to remotely start and stop applications and view processes and services. The tabs and options that are available vary depending on the operating system that you are using.

You use this feature to view active applications, including memory usage and CPU time. The ability to remotely start and stop processes, and be notified when their status changes, runs across multiple operating systems.

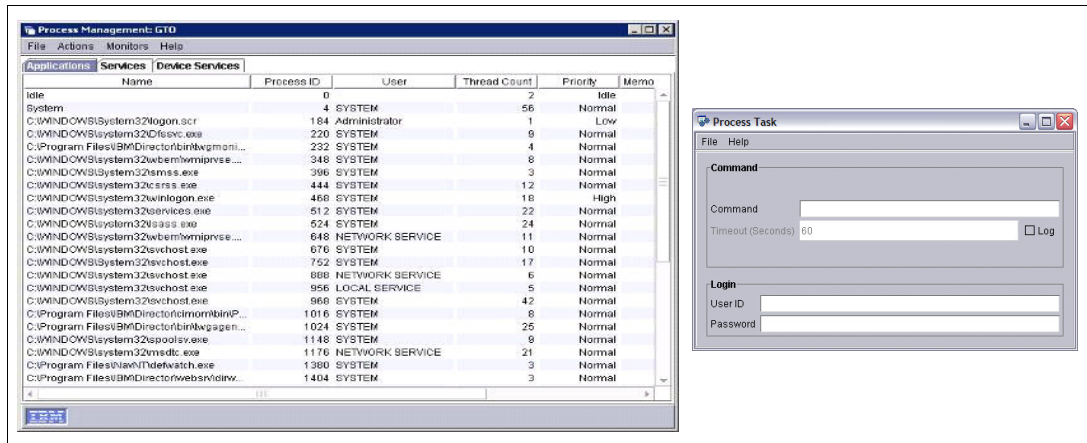


Figure 17-9 Process management window

For Windows, IBM Director includes support for WIN32 services and device drivers. You can create custom processes, which can be run or scheduled on a group or individual system or launched as an action in an event action plan.

17.5.4 Resource monitors

You can use the Resource Monitors task (Figure 17-10) to view statistics about critical system resources, such as processor, disk, and memory usage. With resource monitors, you also can set thresholds to detect potential problems with managed systems or devices. When a threshold is met or exceeded, an event is generated. You create event action plans to respond to resource-monitor events. You can apply resource monitors to individual managed systems and devices and to groups.

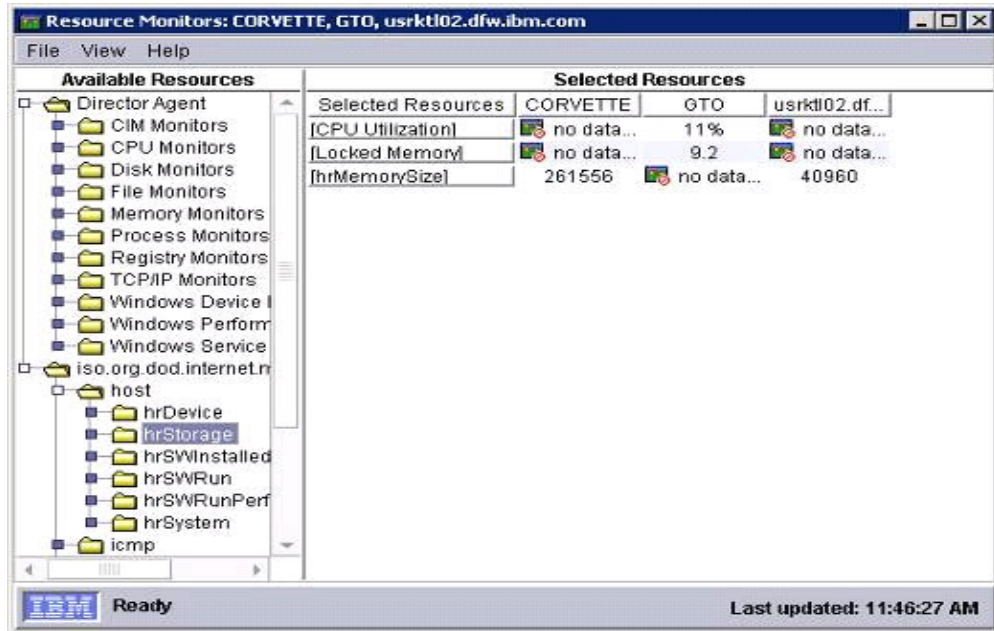


Figure 17-10 Resource monitors

Alerts are based on user-defined system or group thresholds. You can record data for graphing and analysis at a later time.

The resource monitors supports both Director systems (level 2 on all operating systems) and SNMP devices. It has the ability to monitor files and directories for changes. The CPU Utilization resource monitor can be integrated with the Event Action Plan Wizard.

17.6 IBM Director 5.10 extensions

A large set of functions come with the IBM Director commonly ordered with an associated operating system. These can be considered *base functions*. There are two extensions that can be optionally installed as part of the base functions of IBM Director. The other extensions must be downloaded or additionally ordered.

A common term used to collectively refer to these extended functions is *extension task functions*.

Many of the extensions have functions unique to one or two of the supported hardware platforms or operating systems. The Software Distribution Premium Edition has the most cross-operating system functions among these extensions.

The following provides summaries of these extended functions. For more information go to:

http://www-03.ibm.com/servers/eserver/xseries/systems_management/ibm_director/

Many of these extensions are not available for i5/OS or AIX operating systems. A notation is made when the extension applies to i5/OS or AIX.

17.6.1 IBM Director 5.10 base extensions

In this section we discuss IBM Director 5.10 base extensions.

BladeCenter Management

The BladeCenter Management extension is a pair of tools that work together that are focused on managing the IBM eServer BladeCenter infrastructure. BladeCenter Configuration Manager allows you to create and edit profiles that include configuration settings for the components of BladeCenter chassis, such as Ethernet and Fibre Channel switches, and management modules. These profiles can easily and quickly be applied to a BladeCenter chassis in your environment. The Network Device Manager starts vendor software to manage your switches or other network devices.

This capability is critical to support the iSeries capability *iSCSI attachment of xSeries servers* product previewed January 2006. iSCSI represents an international standard — Internet Small Computer System Interface. To find the latest on this iSeries support go to the Web site:

<http://www.ibm.com/eserver/iseries/integratedxseries>

Rack Manager

The Rack Manager is an extension that allows you to group your equipment in rack suites. Using Rack Manager, you can create virtual racks by associating equipment, such as managed systems and devices, networking devices, power devices, and monitors, with a rack in order to visually represent an existing rack in your environment.

You can use Rack Manager to view hardware-status alerts that occur on managed systems or devices in a rack. If a rack component has a hardware-status alert, the rack component is outlined in red, blue, or yellow, depending on the severity level.

Note: BladeCenter Management and Rack Manager are included in the base installation of IBM Director Server on i5/OS.

Also included in the base installation for IBM Director on i5/OS are the following extensions listed in 17.6.2, “Free download extensions” on page 472, and 17.6.4, “Software Distribution (Premium Edition)” on page 474.

- ▶ zVM extension
- ▶ Software Distribution Premium Edition

However, to enable use of the associated functions you must obtain additional authorization (license).

The IBM Director Server also includes the HMC extension. To perform the function supported for the HMC on the server, the PC workstation with the IBM Director Console installed also needs to install a corresponding HMC component.

17.6.2 Free download extensions

In this section we discuss free download extensions.

Electronic Service Agent

This allows IBM Director to interface with IBM support to open support tickets in the case of hardware failures (including predictive failure alerts). For more information go to:

<http://www.ibm.com/support/electronic/>

This is supported on a limited set of operating systems.

ServeRAID Manager

The ServeRAID Manager task manages IBM ServeRAID controllers or adapters that are installed locally or remotely on servers. You can view information that is related to controllers, arrays, logical drives, hot-spare drives, and hard disk drives. Also, you can view configuration settings and events and locate defunct hard disk drives.

IBM ServeRAID controllers are typically installed on IBM xSeries systems.

Web-based access

This enables a browser-based view of managed system information, changing alert standard format (ASF) alerts, changing system settings and configurations, and more. When you install Web-based Access on a managed system, you can access IBM Director Agent and view real-time asset and health information about the managed system from a Web browser.

This feature is supported only on Windows 32-bit operating systems.

System availability

This analyzes the availability of a managed system or group. You can view statistics about managed-system uptime and downtime through reports and graphical representations. This feature is supported only on Windows and Linux operating systems on xSeries platforms.

IBM Real Time Diagnostics

Real Time Diagnostics is systems management software from IBM that runs on your IBM eServer xSeries system while the operating system is active. It can help prevent downtime, thereby increasing your system's availability. It runs with Windows and Linux operating systems.

z/VM Center

This provisions Linux systems on virtual hardware that is based on real System z9™ and zSeries hardware and the z/VM hypervisor. To be able to use z/VM Center, you must purchase the IBM Director Extensions V5.10 feature of IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries V2.1.

UpdateXpress

UpdateXpress can be used to help reduce the cost of computing by providing a method to update server firmware and firmware of supported options contained within the server on most of your xSeries products. If you have purchased an IBM xSeries server, UpdateXpress is available for download at no additional charge.

HMC Manager Tools

This is a suite of tools that integrates many Hardware Management Console (HMC) functions into IBM Director. With this extension installed, IBM Director can discover and authenticate with an HMC that is managing IBM eServer iSeries or pSeries systems (i5/OS, AIX, POWER Linux, and 5/OS IXS/IXA/iSCSI integrated xSeries servers). This presents all processor

enclosures (sometimes referred to as Central Electronic Complexes (CECs)) and Logical Partitions (LPARs) that are running on these servers.

In this manner, the IBM Director console can be used as a single point of management, showing topology and hardware status, as well as providing power and basic hardware control for these advanced server solutions. If more extensive management is required, IBM Director can perform an in-context launch of the HMC directly.

These HMC tools are to be made available originally during early 2006. At the time this presentation was written the tools were not yet available. You cannot perform any IBM Director to HMC functions until the set of tools becomes available during 2006.

17.6.3 Fee-based extensions

In this section we discuss fee-based extensions.

Capacity manager

This provides proactive performance monitoring and bottleneck identification and resolution. This applies to a Windows operating system.

Remote deployment manager

This enables automated, remote system configuration and bare-metal OS deployments. This applies to a Windows operating system.

Application Workload Manager

Application Workload Manager (AWM) extends the systems management capabilities of IBM Director to provide a comprehensive solution for the allocation, management, and control of system resources on IBM xSeries and IBM BladeCenter servers, running a Windows or UNIX operating system. This product is provided by Aurema, Inc. (<http://www.aurema.com>).

17.6.4 Software Distribution (Premium Edition)

The Software Distribution task included in the base IBM Director package includes only IBM Update Assistant, which enables you to distribute to managed systems only software packages created by IBM, such as IBM Director Agent and Console upgrades and packages shipped on the UpdateXpress CD. IBM Update Assistant integrates the use of the UpdateXpress CD to provide an automated method of pushing BIOS, firmware, and device driver updates out to managed xSeries servers.

Software Distribution Premium edition is a fee-based extension to IBM Director that enhances the base functionality described above for the *base* IBM Update Assistant, by enabling you to create packages that can be distributed and installed on any managed system. Using this tool, you can create standard software packages meaningful to the targeted operating system and distribute them throughout your organization. Installation can be scheduled and unattended, enabling you to perform software upgrades on large numbers of managed systems overnight, while end users are out of the office.

This is ordered, for i5/OS, as IBM Director Software Distribution Premium Edition for i5/OS, 5722-SWD. Charges are per processor.

This Premium Edition offers:

- ▶ Integration with IBM Director (for example, 5722-DR1) for consistent, single point of management

- ▶ Intelligent wizards that simplify and streamline the building of a software distribution package from applications that use Microsoft Windows Installer packages, Linux RPM packages, AIX InstallP packages, i5/OS Restore Library, Licensed Program, and Restore Object packages
- ▶ Custom Package Editor for clients who want the flexibility to build and distribute their own applications
- ▶ Convenient, remote drag-and-drop distribution of software to individual systems or groups of systems
- ▶ Ability to schedule a specific time and date to distribute the software, giving you control to minimize network traffic during peak usage
- ▶ Distribution choices of:
 - Streaming copies of the software package to the target system
 - Redirecting large software packages through network distribution servers to alleviate bottlenecks
- ▶ Support for certain non-IBM hardware, helping to leverage your existing Information Technology

As listed above, the Software Distribution Premium edition includes the Custom Package Editor with which you can construct and schedule a unique package distribution of your own.

One *generic example* of a software package is the capability to package several files or a full directory and distribute this package in a scheduled manner. Contrast this with the base IBM Director File Transfer function where you can select multiple files but the transmission occurs in *real time* only.

The figures that follow show a subset of the windows needed to select the Custom Package Editor - Customer Builder function, with which we selected some i5/OS Integrated File System (IFS) files and scheduled a distribution of this package to an IBM Director Agent Level 2 Windows operating system managed system.

We use the IBM Director console to create the package. In Figure 17-11 on page 476 we show two windows to create and edit the custom package.

Create Custom Package/General window

This window shows the General and Files (to be selected) tabs, and the tabs for unique operating system sets of options (AIX, Linux, Windows 2000/XP/2005, and i5/OS).

On the General tab you can specify the package name, send the package directly (checked) or to a file server, and see some system memory and disk space requirements.

Edit custom package

In the Files tab we have already displayed an i5/OS IFS directory on the IBM Director server i5 partition, selected files, and added them to the package (right pane). See Figure 17-11.

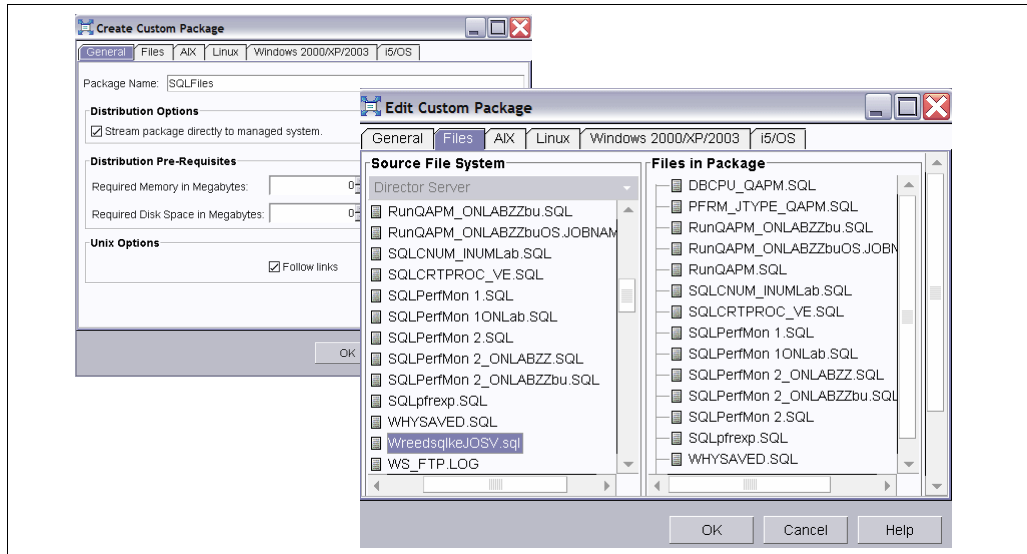


Figure 17-11 Create and edit custom package

In Figure 17-12 we show two windows in the next sequence of steps.

Create custom package

We add i5/OS-specific package capabilities, including naming the destination directory path and the capability to run a function on the target systems before or after the distribution.

Create package for delayed delivery

Create package on the server for later distribution or place the package on a file distribution server (Figure 17-12). In our example scenario we placed the package on the server's disk storage. Though not shown in our screen examples, the package is saved under the name SQLfiles.

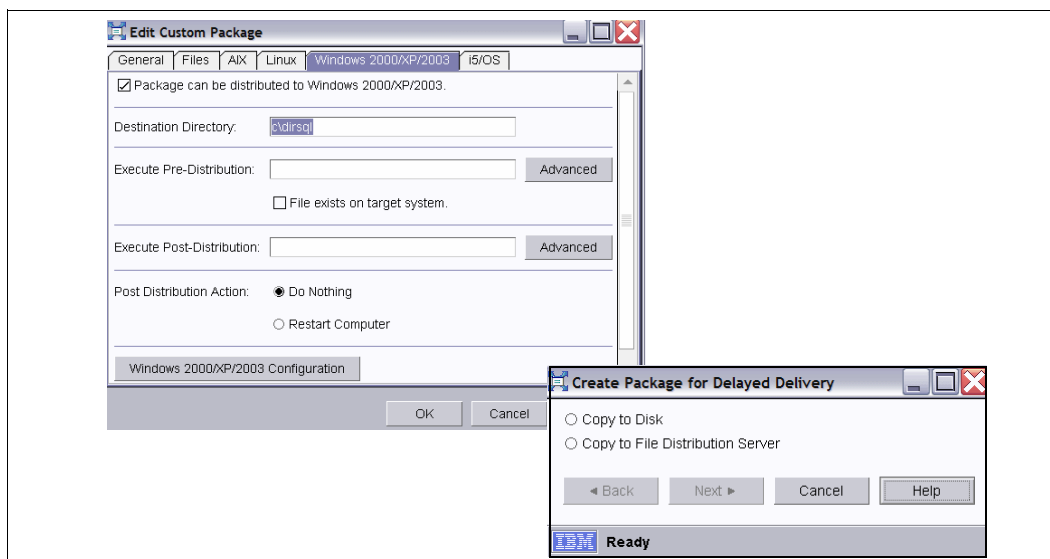


Figure 17-12 Create package

In Figure 17-13 we show an example of the corresponding i5/OS target system options window. This way you can see the i5/OS options to run before and after functions either in native mode (i5/OS) or in QShell.

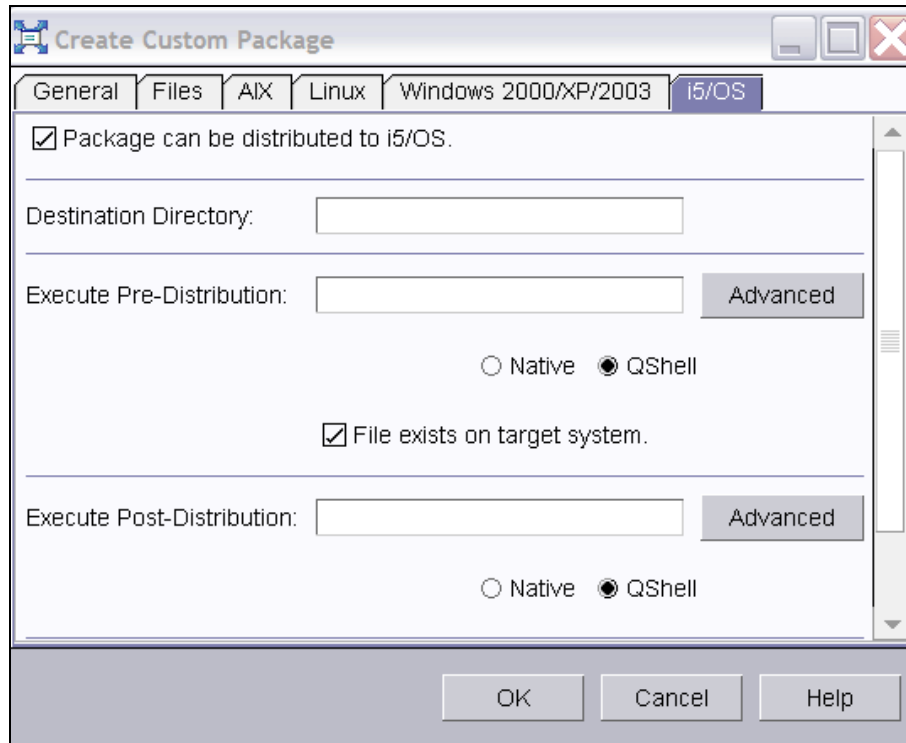


Figure 17-13 i5/OS target system options window

Figure 17-14 on page 478 shows screen examples of the creation of a new job to distribute the package of files and to review the execution results. We have already created and saved the named package of files (SQLfiles) on the IBM Director server's disk storage and selected it to be distributed as part of the job being defined in these windows.

We show two of the windows used to define the scheduled job to distribute the package.

Define a new scheduled job

In this IBM Director task, we do not show the window tabs for selecting the time and date and any repetitive scheduling or the selection of target systems.

The target systems would be among the managed objects shown in the main IBM Director console window central pane. An example of this IBM Director console window is shown in Figure 17-1 on page 457.

In this Options window we show the execution options, how long to keep the execution history of this scheduled job/task, and the level of detail to maintain for events associated with the distribution job/task.

Scheduler

The Job Scheduler support shown here is part of base IBM Director server functionality. This function is generally covered in 17.5.2, "Job Scheduler" on page 468.

This lower right window can show currently scheduled jobs as well as jobs already run. From this window you can make changes to the next scheduled distribution properties, as well as

review the recorded events from a previous distribution, depending on the length of time you specified earlier to retain execution history.

In the window captured in Figure 17-14 we are starting to select to review the execution history of a distribution that ran on November 11, 2005.

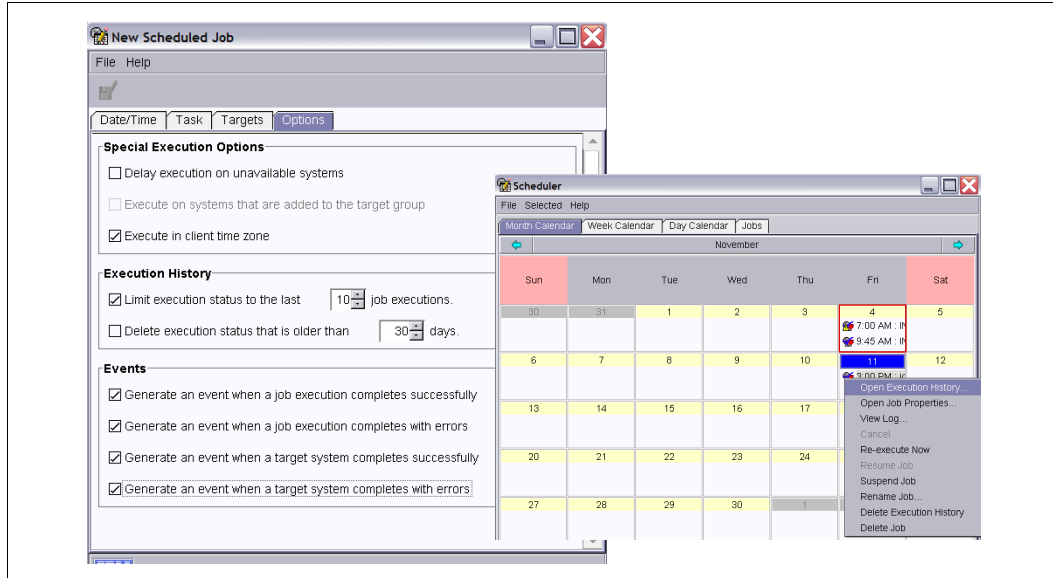


Figure 17-14 Job scheduling

Figure 17-15 shows the selection to view detailed execution history of the distribution of the SQLfiles package. The scheduled job is named jcwindir.

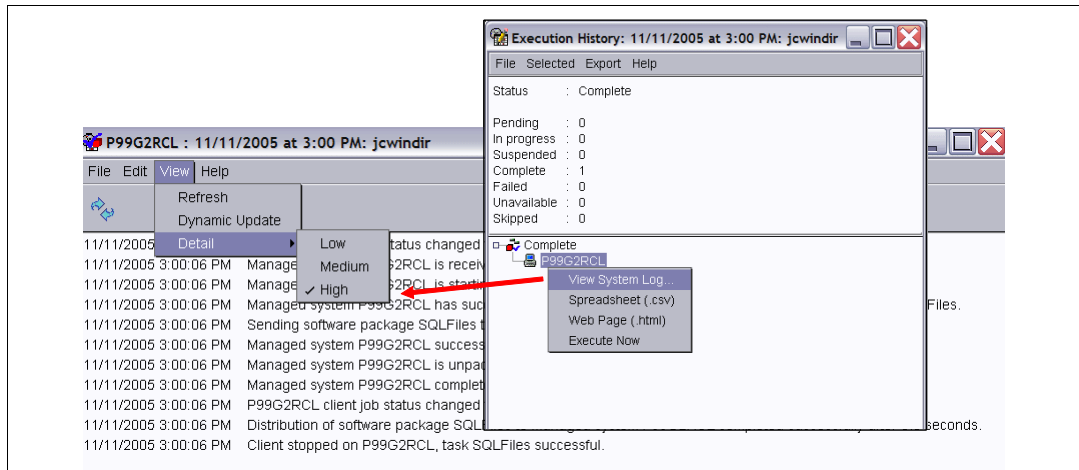


Figure 17-15 Job creation

In the earlier slide showing the new scheduled job, we selected to record all available event classifications so our execution history is, as shown, very detailed.

Software Distribution Premium Edition works with IBM Director V5.10 and V5.20 to support the creation and distribution of software packages across mixed operating system environments.

17.7 IBM Director for i5/OS Version 5.20

As stated earlier in this chapter, Version 5.20 was announced during November 2006. V5.20 extensions were announced across i5/OS and other operating systems in areas like security, high availability, and update management.

For information about IBM Director V5.20 enhancements you can view the following announcement letters:

- ▶ IBM Director for i5/OS, V5.20: Announcement Letter 206-287
- ▶ System x announcement IBM Director Version 5.20: Announcement Letter 206-290
- ▶ IBM Director for AIX 5L and Linux on POWER, V5.20: Announcement Letter 206-286
- ▶ IBM Virtualization Engine and Infrastructure for Services Linux on System z™, V2.2: Announcement Letter 206-294

This section summarizes the Version 5.20 content for IBM Director on i5/OS, which is primarily focused on ease of installation under i5/OS.

- ▶ Simplified installation, optimized for i5/OS

IBM Director for i5/OS in earlier versions (V4.nn and V5.10) uses the Virtualization Engine-based installation process. This is known as IBM Director with Virtualization Engine Console, 5733-DIR.

Of the total CDs available when packaged as part of the Virtualization Engine, the last three contain IBM Director software. The first of these three CDs is placed into a PC workstation LAN-connected to the IBM System i model partition. The IBM Director installation wizard installs software onto the i5/OS partition from the PC workstation.

The IBM Director V5.20 installation process uses a separate CD that is now inserted into the CD-ROM drive owned by the i5/OS partition. The i5/OS Restore Licensed Program (RSTLICPGM) command for 5722-DR1 can be used to install from the CD-ROM drive. With the RSTLICPGM process, either 5722DR1 (server and agent) or 5722DA1 (agent only) can be installed to the i5/OS partition.

Effective February 14, 2007, licensed program 5733-DIR will be removed from marketing. Effective April 30, 2009, IBM will withdraw support for 5733-DIR under the IBM International Program License Agreement.

- ▶ Continued support for new iSCSI supported BladeCenter models and features

New models and features of IBM BladeCenter configurations have become available since IBM Director 5.10 became available. IBM Director V5.20 supports these newer models and features.

Code downloads for this enhanced support are available for IBM Director 5.10.

- ▶ Continued support of updated and new HMC features and functions as they become available

Code downloads for this enhanced support are available for IBM Director 5.10.

- ▶ Easier setup to collect Linux or AIX 5L 5.3 or later basic performance data and include this information with i5/OS Collection Services data

Starting with i5/OS V5R3 i5/OS Collection Services running on one i5/OS partition can be set up to retrieve base performance data from a partition running a supported Linux distribution or AIX 5L V5.3 with IBM Director Agent Version 4.n or 5.10. IBM Director Server Version 4.n or 5.10 and later must be active in the i5/OS partition.

The objective of this multiple partition collection of performance data for i5/OS, AIX 5L, and POWER Linux is to forward summarized performance data to the IBM service offering

PM for System i5 (formerly called PM iSeries). Then under this offering an IBM representative, business partner, or customer can see performance trends for all partitions active on a system. For more summary-level information about this PM for System i5 capability refer to 16.12, “Performance Management for System i5” on page 444.

With the availability of IBM Director for i5/OS V 5.20, setting up this performance data collection is much easier. With previous versions of Director, you had to manually install Director support by using FTP to send and install the IBM Director software files to the Linux and AIX partitions. Even on i5/OS you have to copy some files from the Collection Services IFS directory to the IBM Director directory. IBM Director 5.20 has all of the necessary files already integrated.

The IBM Director steps to perform the manual process are documented at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/index.jsp?topic=/rzahx/rzahxcollectdatapart.htm>

With Director 5.20, steps 8 through 11 are no longer necessary.

- ▶ Easier user ID management across the operating systems supported by IBM Director 5.20

User accounts can be defined on an LDAP server, allowing user administration to be centralized for multiple management servers. This means that you do not have to have a matching user ID (profile) on each system. Rather, the user ID is in the LDAP list, and this is sufficient to access the IBM Director Server on each system or partition within the LDAP domain.

- ▶ Other V5.20 extensions across i5/OS and other operating systems in areas that include security, higher availability, and update management. For example:
 - User accounts can be defined on an LDAP server. This facilitates centralized user administration for multiple management servers. You can access the IBM Director server using your LDAP entry rather than a user profile on the operating system.
 - New update management capability. The Update Manager task provides update management on managed servers, including creating user profiles, downloading updates, comparing updates on managed servers, and generating reports on the activity in these areas.

You can find more details on IBM Director at the IBM Systems Software Information Center at:

<http://publib.boulder.ibm.com/eserver/swic.html>

A direct link to a list of the PDFs for IBM Director at the IBM Support Web site is:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?brandind=5000016&indocid=MIGR-61788>

You can also find specific V5.20 information at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/dirinfo_5.20

17.8 IBM Director and Tivoli

Figure 17-16 is a high-level summary view of IBM Director capabilities and the capabilities of several Tivoli products. In general, Tivoli products are more oriented to the application and middleware levels. IBM Director is more oriented toward the operating system and hardware levels. Of course, there is overlap.

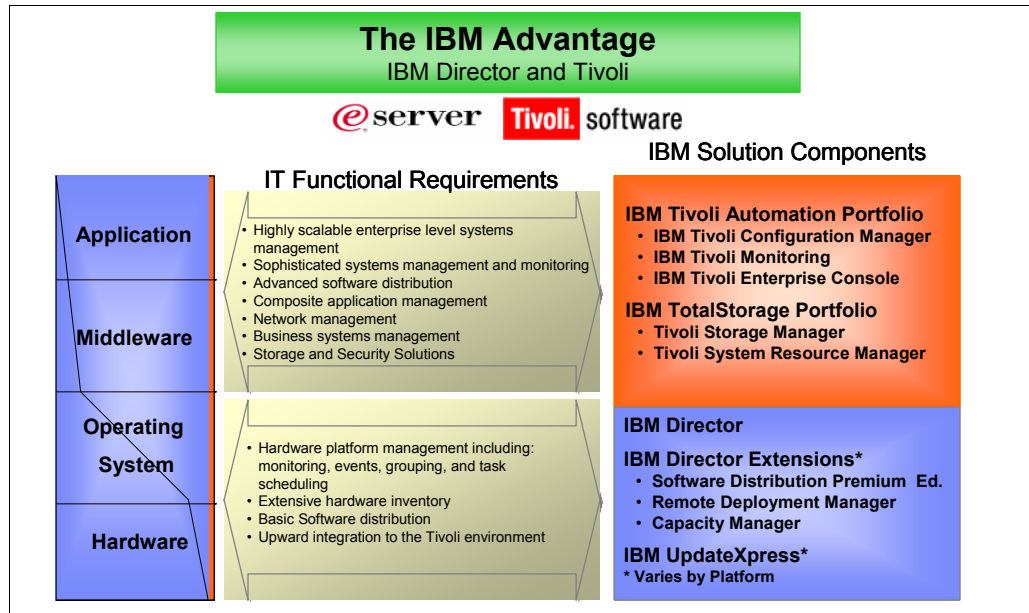


Figure 17-16 Tivoli overview

There will be adjustments and changes over time under both the Virtualization Engine suite of products and the Tivoli suite of products. From an iSeries viewpoint, the best single source of iSeries System Management capabilities and Tivoli product capabilities is included in the redpaper *IBM eServer iSeries Systems Management Handbook*, REDP-4070.

The full list of Tivoli products and on what platforms they currently execute can be found at the following Web site:

<http://www-306.ibm.com/software/tivoli/sw-atoz/index.html>

17.9 References

See the following:

- ▶ IBM Director Web site
http://www.ibm.com/servers/eserver/xseries/systems_management/ibm_director
- ▶ IBM Redbooks
<http://www.ibm.com/redbooks>
- ▶ IBM Redbook *Implementing IBM Director 5.10*, SG24-6188
<http://www.redbooks.ibm.com/abstracts/sg246188.html?open>
- ▶ Redpaper *IBM eServer iSeries Systems Management Handbook*, REDP-4070
<http://www.redbooks.ibm.com/abstracts/redp4070.html?open>



i5/OS and Lotus products

This chapter covers V5R4-related changes to the Lotus product family.

18.1 Lotus products

Lotus software makes your business more effective, improves communications, and makes it easier for people to work together. These products are the tools that provide your business with the ideal environment for effective collaboration and communications. You can perform e-mail, workflow-based computing, and data management, and you can make your server faster, more flexible, and more user friendly. You can bring the flexibility and efficiency of real-time communications to your business.

18.2 Lotus Domino for i5/OS

Lotus Domino for i5/OS is a powerful, popular, versatile, and integrated groupware product. It provides functions that include e-mail, workflow-based computing, and the integration and management of both structured and unstructured data. Domino is a server product that runs on a variety of platforms, providing easy-to-manage interoperability in a heterogeneous network.

18.2.1 Domino releases on i5/OS

Regularly scheduled updates to Domino feature releases are called Maintenance Releases (MRs). These are denoted by the use of a third digit in the release number. Fix packs (FPs) are issued as needed to fix specific problems identified between scheduled MRs. (Fix packs have previously been referred to as interim fixpacks and critical fixpacks.)

Table 18-1 shows Domino 6 releases, MRs, and compatible operating system releases.

Table 18-1 Domino 6 releases, MRs, and compatible operating system releases

Releases and critical fix packs	i5/OS V5R4	i5/OS V5R3	OS/400 V5R2
Domino 6.5.5 6.5.5 FP1	X	X	X
Domino 6.5.4 ▶ 6.5.4 FP3 ▶ 6.5.4 FP2		X	X
Domino 6.5.3 6.5.3 FP1		X	X
Domino 6.5.2 6.5.2 FP1		X	X
Domino 6.5.1 6.5.1 IF1		X	X
Domino 6.5		X	X
Domino 6.0.5		X	X
Domino 6.0.4		X	X
Domino 6.0.3		X	X
Domino 6.0.2 CF1 ^a ▶ 6.0.2 CF2			X
Domino 6.0.1 ▶ 6.0.1 CF3 ▶ 6.0.1 CF2 ▶ 6.0.1 CF1			X

Releases and critical fix packs	i5/OS V5R4	i5/OS V5R3	OS/400 V5R2
Domino 6.0			X

a. For most regions, the initial release of Domino 6.0.2 for iSeries already includes CF1.

Table 18-2 shows Domino 7 releases, MRs, and compatible operating releases.

Table 18-2 Domino 7 releases, MRs, and compatible operating system releases

Releases and critical fix packs	i5/OS V5R4	i5/OS V5R3
Domino 7.0.1	X	X
Domino 7.0		X

The Lotus developerWorks product page for Notes/Domino contains links to all of the latest Notes/Domino-related information and downloads for all platforms, including status of future MRs, the Fix List database, Notes client downloads, and more.

Domino 6.0.5/6.5.4 became available in April 2005, with Fix Pack 2 available in October 2005. 6.5.4 Fix Pack 3 became available by mid-February 2006.

Domino 6.0.5 is the last 6.0.x maintenance release (MR), and there are no plans for Domino 6.0.6 or beyond.

Domino 6.5.5 became available in December 2005, and included mostly fixes. Fix Pack 1 was released in April 2006. Two more fix packs are expected in 2006.

Current planning for Domino 6.5.6 has availability planned for 1Q 2007.

Domino 7.0 became available in August 2005, and requires i5/OS V5R3. The Lotus Extended Products also shipped concurrently include:

- ▶ QuickPlace® 7.0
- ▶ Sametime® 7.0
- ▶ LEI 7.0
- ▶ Domino Document Manager 7.0
- ▶ Workflow 7.0

Domino 7.0.1 became available on February 7, 2006, and is required for i5/OS V5R4. Domino 7.0.1 Fix Pack 1 has been released for download, and another release is planned for 2007.

In a future release, the Domino 7 product ID may be changed to match the Lotus Passport Advantage PID # (as has been done already for Sametime). For example, the Domino 7.0 PID # is 5724E62. Here is what is displayed on DSPSFWRSC:

```
5733LD7 *BASE 5050 Lotus Domino 7 (5724E62)
```

Extensive documentation exists on the Web for each release, in a variety of formats. Find the release you are interested in at the Lotus documentation page:

<http://www.lotus.com/ldd/notesua.nsf/RN?OpenView>

18.2.2 End of marketing

On November 8, 2005, IBM announced end of service and end of marketing for Lotus 6.0.x programs. The release letter can be found on the Web at:

<http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&htmlfid=897/ENUS905-243&apname=isource>

Effective on the dates listed in Table 18-3, IBM withdrew from marketing part numbers from the following product releases licensed under the IBM International Program License Agreement.

Table 18-3 Withdrawal from marketing

Program number	Program release name	Withdrawal from marketing date
5724-E89	IBM Lotus Enterprise Integrator@ V6.0.1	1/12/06
5724-E89	IBM Lotus Enterprise Integrator V6.0.2	1/12/06
5724-E70	BM Lotus Notes/Domino V6.0.3	4/11/06
5724-E70	IBM Lotus Notes/Domino V6.0.4	4/11/06
5724-E70	IBM Lotus Notes/Domino V6.0.5	1/19/07

Effective on the dates listed in Table 18-4, IBM will withdraw support from the following product releases licensed under the IBM International Program License Agreement.

Table 18-4 Withdrawal from service

Program number	Program release name	Withdrawal from support date
5724-E89	IBM Lotus Enterprise Integrator V6.0, V6.0.1, and V6.0.2	4/30/07
5724-E70	IBM Lotus Notes and Domino V6.0, V6.0.1, V6.0.2, V6.0.3, V6.0.4, and V6.0.5	4/30/07

18.3 Lotus Workplace on i5/OS

IBM Workplace and System i together provide an integrated solution. They allow organizations to deliver customized, role-based electronic workplaces to their employees and strategic partners.

The Workplace solution is a tool to help you achieve significant improvements in employee productivity and overall workforce effectiveness. For many organizations, IBM Workplace and iSeries can be a major game-changer to empower twenty-first century workers. The product is discussed in depth at:

<http://www-03.ibm.com/servers/eserver/iseries/software/workplace/index.html>

18.3.1 Workplace Services Express (WSE)

Workplace Services Express Version 2.6 for i5/OS, 5733-WSE has been released. The product comes with a default evaluation license that lets you install and use Workplace Services Express for 120 days at no charge. To continue using the product after the

evaluation license expires, you must purchase Workplace Services Express. After you purchase Workplace Services Express, you can convert the evaluation license to a full production license.

For more information about the evaluation license, see Getting Started in the Workplace Services Express Information Center at:

<http://www.ibm.com/developerworks/workplace/documentation/servicesexpress/>

You cannot use the Workplace Services Express Version 2.6 media or downloaded e-image to upgrade an existing installation based on a prior version of the product.

New in this release

IBM Workplace Services Express provides improved performance, supports a broader range of customer configurations, and includes the enhancements described below:

- ▶ Workplace Services Express supports IBM i5/OS V5R4 and includes an updated installation program to simplify deployment on i5/OS.
- ▶ Workplace Services Express includes a special license that lets you install IBM DB2 Universal Database Express Version 8.2 for use with Workplace Services Express.
- ▶ Available portlets now include the IBM Portlet Application for Microsoft Exchange 2003 and the Lotus Notes Version 6 portlet. Workplace Services Express also comes with updates to existing portlets such as news, stock quotes, and weather.
- ▶ Developers can use Workplace Designer to create components and build applications for a Workplace Services Express environment.
- ▶ This release also includes resolutions to some Japanese translation issues.

For more information see the following IBM Passport Advantage Web site, select **Trials and betas**, and search for *Workplace Services Express*:

<http://www-306.ibm.com/software/howtobuy/passportadvantage/>

See also the IBM Redbook *Deploying IBM Workplace Services Express on the IBM eServer iSeries Server*, SG24-66644.

18.3.2 Workplace Collaboration Services (WCS)

IBM Workplace Collaboration Services is an integrated family of collaborative capabilities that are based on open standards and delivered to users in a secure work environment.

Collaborative capabilities such as e-mail, calendars, training, instant messages, shared documents, and so on, are integrated into a single user experience. To simplify administration, all capabilities are activated, administered, and provisioned from the server.

Workplace Collaboration Services employs a standards-based platform that includes IBM WebSphere Application Server, WebSphere Portal, IBM WebSphere Member Manager, IBM Tivoli Directory Server (optional), and DB2 Universal Database Enterprise Edition.

The IBM Workplace browser client and the IBM Workplace Managed client are the two clients that are available for users to use to access the installed capabilities of Workplace Collaboration Services.

New features for IBM Workplace Collaboration Services

Introduced with Version 2.6, which requires i5/OS V5R4, are the enhancements listed in Table 18-5.

Table 18-5 Enhancements for both clients

Feature	Enhancements
IBM Workplace Collaborative Learning™	<ul style="list-style-type: none"> ▶ Improved management capabilities for offerings. IBM Workplace Collaborative Learning users with the appropriate permissions can copy offerings and save the copies to a folder. They can subsequently delete those copies without deleting the original offering, or they can delete the original offering and all of its copies in a single operation. ▶ Improved course results. IBM Workplace Collaborative Learning users with the appropriate permissions can modify course results for a number of students simultaneously rather than having to modify course results for one student at a time.
Team spaces	Users can create roles that give members of the team space read-only access to the team calendar.
Members portlet	Users can add, move, and delete members and groups in fewer steps.

Table 18-6 lists server enhancements.

Table 18-6 Server enhancements

Feature	Enhancements
Client provisioning	Using supplied scripts, create an installation CD that installs only the IBM productivity tools and the IBM Workplace Managed Client™ framework needed to run those tools for standalone use.
User policies	<ul style="list-style-type: none"> ▶ Change user operating system file associations to open with IBM Productivity Tools. ▶ Specify how IBM Productivity Tools should be used by Workplace Managed Client. ▶ Allow access to IBM Workplace Designer. ▶ Allow access to Activity Explorer. ▶ Disable mail when user exceeds specified maximum mail storage. ▶ User can set the out-of-office application to send automatic mail replies. ▶ Set the maximum allowed size of an e-mail message, including attachments. ▶ Allow Workplace Managed Client users to open mail attachments for an expanded list of file types: .jpg, .bmp, .gif, .tif, .txt, .rtf, .pdf, .zip, .mpp, .ppt, .xls, .prz, .123, .lwp, .wpm, .sxi, .sxc, .sxw, .odt, .ods, and .odp.
Messaging	<ul style="list-style-type: none"> ▶ Log mail routing events at the server or cell-wide level. ▶ Schedule the out-of-office task to collect the addresses of e-mail senders and send the list to the user. ▶ Automatically disables the out-of-office task when a user is scheduled to return. ▶ Specify a range of TCP/IP addresses when setting SMTP filters for trusted, suspect, or blocked connections. ▶ In addition, the user interface for selecting the junk mail filter has been simplified by eliminating ports for learning and filtering.
FIPS enablement	Enable FIPS support on the IBM Workplace Collaboration Services server to allow browser clients to communicate with the server using TLS. The FIPS 140-2 standard specifies the use and implementation of an approved set of cryptographic algorithms when encrypting or decrypting data that is at rest or in transit.

Table 18-7 lists IBM Workplace Managed Client enhancements.

Table 18-7 IBM Workplace Managed Client

Feature	Enhancements
CD Install	Using the installation CD mentioned in Table 18-6, users can install the stand-alone productivity tools without accessing the Workplace software server. Users can later (optionally) connect to a server and provision their client from that server. They can also update their stand-alone client from a local site.
Activity Explorer	Users can use this feature to more effectively track and manage multiple points of interaction (called activities) that relate to a project or process. People can share notes, files, persistent chats, and shared snapshots, associating related activities to ensure that key information is organized, readily available, and easy to find.
Data access tool	Users can create forms, grids, and reports in order to add, edit, delete, and view summaries of database records stored in an IBM Cloudscape™ database.
Instant messaging	New contacts dialog box and chat window design.
Template support	Users can create new documents from templates in their document libraries.
Productivity tools	Word processing, presentation, and spreadsheet editors now support alternative text for floating frames, graphics, and objects. Also, users of the word processing editor can specify whether to use auto breaks across pages for selected table rows.

18.3.3 IBM Workplace Designer

IBM Workplace Designer is new in this release. Workplace Designer is available from the IBM Workplace Managed Client.

Workplace Designer V2.6 is an intuitive, visual development tool that enables you to create and deploy components for Workplace Services Express applications. Script developers, Lotus Domino application developers, and the software developer community can easily create and deploy components to fit specific industry and business needs within Workplace Services Express.

The visual, intuitive user interface, powerful scripting capabilities, and simplified component deployment make Workplace Designer V2.6 the fastest and easiest way to create and deploy new components for Workplace Services Express. Developers who use Workplace Designer do not need Java development skills. Developers with these skills can make use of Workplace APIs and other Java APIs exposed through Workplace Designer V2.6.

Workplace Designer comes with sample applications that you can use to familiarize yourself with IBM Workplace application development.

During October 2006, a new update to Workplace Services Express Workplace Designer V2.6 became available.

For more information about Workplace Designer and available documentation see the topic IBM Workplace Designer in the Information Center.

18.4 References

See the following:

- ▶ Domino for iSeries
<http://www.ibm.com/eserver/series/domino>
- ▶ IBM Workplace for iSeries
<http://www.ibm.com/eserver/series/workplace>
- ▶ IBM Redbook *Implementing Domino 7 for i5/OS*, SG24-7311
- ▶ iSeries Domino Performance
<http://www-128.ibm.com/developerworks/lotus/performance>

Search through the following links to find iSeries performance-related information about Lotus products:

- Technical articles
- White papers
- IBM Redbooks
- ▶ Lotus DeveloperWorks
<http://www.ibm.com/developerworks/lotus>
- ▶ IBM Workload Estimator
<http://www.ibm.com/eserver/series/perfmgmt>
 - a. Select **Sizing resources**.
 - b. Select **IBM Systems Workload Estimator**.



A

Hardware Management Console microcode and system firmware compatibility levels

In this appendix we provide summary level information about how to download the latest code levels, as well as step-by-step guides on how to view the levels currently installed.

HMC machine code upgrade

Before you start upgrading the HMC machine code, you should be familiar with the topics in this section and understand how to obtain the latest software and how to view the level of the current software, saving HMC and configuration data and performing the actual upgrade.

For a detailed step-by-step instruction of these procedures, refer to the *Upgrading your HMC machine code* topic in the IBM Systems Hardware Information Center at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iph5/fixeshmc_upgrades.htm

Getting the upgrade

You may contact IBM service and support nearest you to order the HMC upgrade or Recovery media. Another way is to order it on your own from the HMC Web site (by selecting the Recovery media tab). From there you can order the recovery media or you may download the upgrade package using your browser.

The HMC Web site can be found at:

<https://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html>

The recovery media is a bootable DVD image that can do a clean install of the version/release of HMC code you are upgrading to. If you have downloaded the code in the ISO format, you can burn this image onto a DVD to create your own copy of the recovery media.

For instructions on how to burn ISO images onto DVDs in Linux go to:

https://www14.software.ibm.com/webapp/set2/sas/f/hmc/burn_dvd.html

Viewing the current HMC machine code level

Before upgrading the HMC machine code to a certain level, we recommend that you first check your current level of code installed.

Saving the managed system's profile data

We strongly recommend that before any upgrade is started to your HMC code that you should first save the profile of your managed system. This is a precautionary step. In case the upgrade should fail and the profile data be corrupted, you can still recover back to the state before the upgrade.

Saving critical console information

We also recommend that you save critical console information before installing a new version of HMC code so that you can restore previous levels of code should your upgrade fail.

Record HMC configuration information and remote command status

This is another precautionary measure before doing the upgrade.

Saving the upgrade data

You can save the current HMC configuration in a designated disk partition on the HMC. Do this data just before doing the upgrade to a new release.

Perform the HMC microcode upgrade and verification

After saving the upgrade data in the previous section, you may now start the actual upgrade. At this point we highly recommended that you have the needed upgrade media. It is also

advised that you verify whether upgrade is successful after the upgrade procedure completes.

Firmware history: releases and service pack levels

The chart in Figure A-1 shows the firmware history, service pack, and release levels and their supported HMC releases available up through August 2006.

HMC release level	Power Hypervisor or System Firmware (SF)				
	SF Release level	Sometimes called	Date base is first available	Base level	Service Packs and associated level
V6R1	SF240	GA7	2006 December	284	Service Pack 4
V6R1	SF240	GA7	2006 October	261	
V5R2	SF240	GA7	2006 Feb	201	202, 219 +(future SPs)
V5R1	SF235	GA6	2005 Oct	160	180, 185, 206, 209
V4R5	SF230	GA5	2005 Sep	120	126, 141, 145, 156
V4R5	SF225	GA4	2005 Feb	96	
V4R3	SF222	GA3	2004 Nov	71	75, 81
V4R2	SF220	GA2	2004 Aug	43	45, 46, 49, 51
V4R1	SF210	GA1	2004 Jun	21	22, 23, 24, 25

Find in HMC version information Find in ASMI or Service Tools

Figure A-1 HMC level and firmware level

ASMI is the Advanced System Management Interface and Service Tool is the i5/OS Service Tools functions access through the Start Service Tools (STRSST) command. We discuss these tools later in this appendix.

SF240_284 contains miscellaneous updates but is required to support feature #7119, which is a 2-port optical HSL/RIO bus expansion card for the model 595. #7119 adds two optical ports, enabling capacity for an optical loop for model 595.

Up-to-date information about the firmware and HMC level capabilities can be found at the following Web pages:

<http://www14.software.ibm.com/webapp/set2/sas/f/power5cm3/supportedtable1.html>

<http://www14.software.ibm.com/webapp/set2/sas/f/power5cm3/supportedtable2.html>

2007 Daylight saving time changes support

The Energy Policy Act of 2005 (Public Law 109-58) was passed by the United States Congress on July 29, 2005 and signed into law on August 8, 2005. This law changes the time-change dates for daylight saving time in the United States. Starting in the spring of 2007, the U.S. daylight saving time (DST) will be extended by four weeks, starting three weeks earlier on the second Sunday in March and ending one week later on the first Sunday in November.

This change could impact current applications running on any operating system using time dependent values.

i5/OS and IBM support for this change is documented in several iSeries and IBM Web sites. Consult the following IBM support Web site for information about this subject:

<http://www.ibm.com/support/alerts/daylightsavingsttimealert.html>

For i5/OS changes listed in this book see 3.2, “Daylight saving time support for March 2007” on page 78.

The following are the HMC Version Packages available and any associated new daylight saving time support fixes that must be installed:

- ▶ HMC 6.1: Support already included. There is no need to install an additional package.
- ▶ HMC 5.2.1: Support already included. There is no need to install an additional package.
- ▶ HMC 5.2: APAR MB01365/PTF MH00586.
- ▶ HMC 5.1: APAR MB000005/PTF MH00593 (replaces MH00493).
- ▶ HMC 4.5: APAR MB001276/PTF MH00494.
- ▶ HMC 3.3.7: APAR IY87070/PTF U808917.
- ▶ HMC 3.3.6: APAR IY79212/PTF U806370.

Reference the HMC Support Web site listed just before this section.

Viewing HMC firmware and system firmware levels

In this section we give examples of how to display what level of firmware you have on your HMC and your System i5 model.

Displaying HMC level

Perform the following steps to view the firmware level of your HMC:

1. At your actual HMC device or having connected to the HMC using a Web-based System Manger Remote Client Browser session, enter a valid HMC user and password.
2. Select HMC option **Licensed Internal (LIC) Code Maintenance**.
3. Select **HMC Code Update**.
4. Select **HMC Version Information**.

Figure A-2 shows an example window of version information.

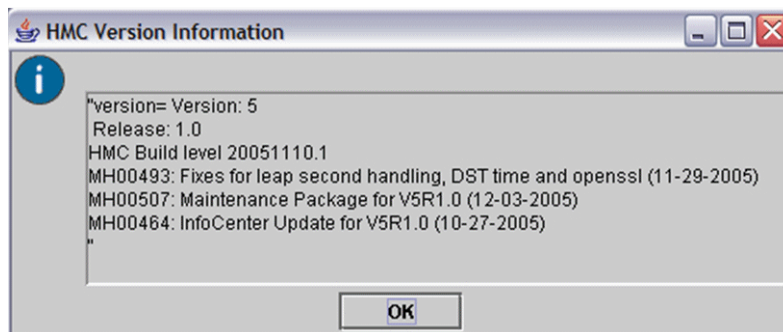


Figure A-2 Viewing the HMC level

Displaying IBM System i5 system firmware level

There are three ways to view the system firmware level. Choose the one most convenient for you, depending on your particular environment. Only an authorized customer user or IBM service representative should have access to the following functions:

- ▶ From an HMC session, use the Licensed Internal Code Maintenance option.
- ▶ From a 5250-based workstation session use the i5/OS SST command to display FLASHLEVEL information.
- ▶ From a properly connected browser, use the Advanced System Management Interface (ASMI) from a properly connected Web browser.

Display system firmware level using an HMC

At your actual HMC or having connected to the HMC using a Web-based System Manager Remote Client Browser session, enter a valid HMC user and password:

1. Select HMC option **Licensed Internal (LIC) Code Maintenance**.
2. Select **Licensed Internal (LIC) Code Update**.
3. Select the managed system you want to look at and click **OK**.

Figure A-3 shows an example window before clicking OK. In our example the HMC is managing two systems.

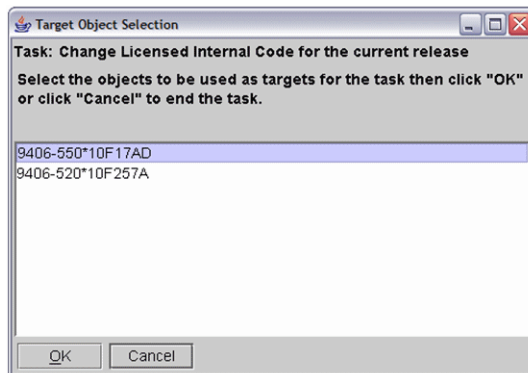


Figure A-3 HMC-based displaying system firmware level 1 of 2

4. On the next window select **View system information**. Then select **None – display current values**. The next window (Figure A-4) shows the firmware level.

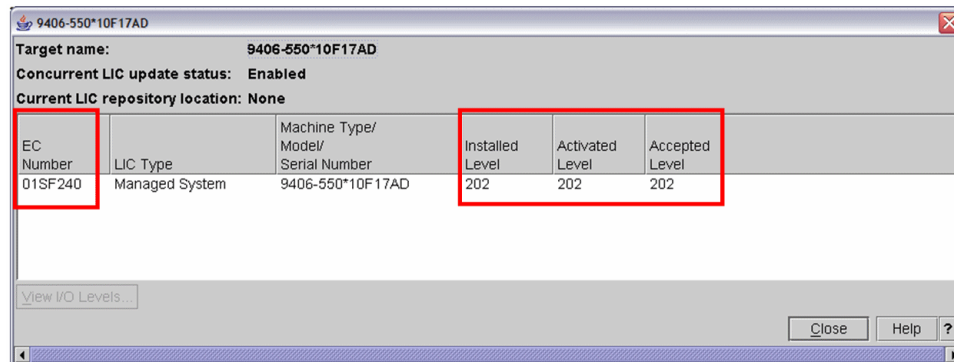


Figure A-4 HMC-based displaying system firmware level 2 of 2

Level 202 is essentially level 201 with some integrated fixes.

Display system firmware level using i5/OS Service Tools (SSTs)

Use the i5/OS System Service Tools to display the firmware level, identified on this interface as FLASHLEVEL information:

1. On a 5250 session, enter the STRSST command and sign on with a valid service tools user ID and password.
2. Select option 1 (Start a service tool).
3. Select option 4 (Display/Alter/Dump).
4. Select option 1 (Display/Alter Storage).
5. Select option 2 (LIC data).
6. Select option 14 (Advanced analysis). This option is not listed on the LIC data window.
7. Select **FLASHLEVELS**.
8. Press Enter (no need to specify any options).
9. The flash levels are displayed.

Figure A-5 shows an example of a flash level that indicates the firmware level.

Base Level:			
Side	Date/Time	MI Keyword	PTFs (V5R4M0)
-----	-----	-----	-----
Memory	20051202/1310	SF235_180	MH00506
Flash P	20051202/1310	SF235_180	MH00506
Flash T	20051202/1310	SF235_180	MH00506
Load Source A	20060130/1303	SF240_202	MH00578
Load Source B	20060505/1020	SF240_219	MH00677

Figure A-5 SST-based view of system firmware level

Display system firmware level using ASMI

The Advanced System Management Interface is a specialized IBM System i5 interface for viewing and managing low-level hardware and software. The Firmware Service Processor (FSP) component includes a secure HTTP server. This HTTP server can be accessed either from an HMC or directly from a browser by attaching to one of the FSP ports within the processor enclosure.

To view the system firmware level, see below.

The FSP has two Ethernet ports. If an HMC is not attached and the FSP port addresses have not previously been changed by an authorized administrator, the default IP addresses are:

- ▶ FSP-HMC0 defaults to IP address 192.168.2.147.
 - ▶ FSP-HMC1 defaults to IP address 192.168.3.147.
1. Attach a standard Ethernet cable from the HMC0 or the HMC1 port to a PC workstation with a browser.
 2. Assign an address to the PC in the 192.168.x.x subnet range, using normal network configuration options for the PC (typically with a Windows or Linux operating system).
 3. Start a secure HTTPS connection to the FSP using an available browser (for example, <https://192.168.2.147>).
 4. The FSP HTTP sign-on screen appears. The currently installed firmware level is displayed in the upper right-hand corner of the window.

See Figure A-6 for an example. In this case level SF253_180 is installed.

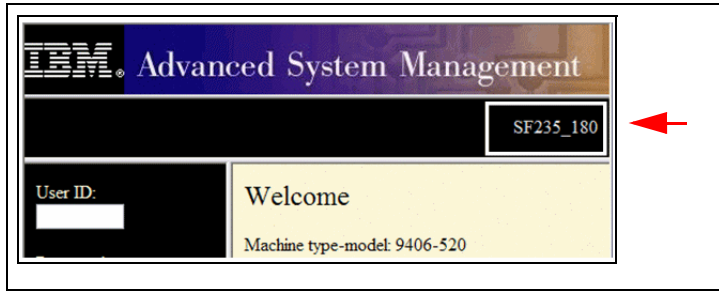


Figure A-6 ASMI-based displaying system firmware level



IBM System i5 memory details

In this appendix we provide details on available memory features and memory card placement rules for the IBM System i5 POWER5-based models.

System i5 model DDR1, DDR2 memory and plugging rules

The original IBM System i5 models were introduced with DDR1-level memory. These models and DDR1 memory remain supported and available for new orders. The IBM System i5 520, 550, and 570 models introduced during January 2006 require Double Data Rate Dual Channel (DDR2) memory. During mid-2006 the DDR2 memory became available on the 595 models.

There are memory *rules* that are common across all of the System i5 models, and some that are unique for each model. Therefore, we discuss memory rule details under model-specific topics.

System i5 520 and 550 model memory technology considerations

The DDR1 memory cards (DIMMs) are required on the 1.5 GHz and 1.65 GHz 520 and 550 models. The DDR2 memory cards are required on the 1.9 GHz 520 and 550 models.

With DDR1 technology, one memory feature number delivers four DDR1 DIMMs. With DDR2 technology, one memory feature delivers two DDR2 DIMMs.

As shipped from IBM, the model-specific DDR1 or DDR2 memory DIMMs are plugged on each available processor card appropriately for good performance under busy workloads. On the 520 and 550 models, there are up to two processors per processor card, and the 550 model comes with two processor cards.

Because additional or new memory can be ordered after the initial system delivery configuration is ordered after the original configuration was installed, we indicate, for better performance:

- ▶ Minimum of one memory feature per processor card. This is very important when more than one processor card is being used. The 550 has two processor cards. If only two processors are physically activated, these suggestions are not as important.
- ▶ Plug pairs of DIMMs so that each processor card has the same/similar total MB of memory. Optimal performance can be achieved when the DIMMs are plugged to that same/similar total MB of memory per processor card.

The better positioning of memory cards becomes more important as the workload increases in a manner such that the active threads need larger amounts of working storage, and processors on more than one processor card are used.

Figure B-1 shows a table of available DDR1 and DDR2 feature numbers and DIMM sizes.

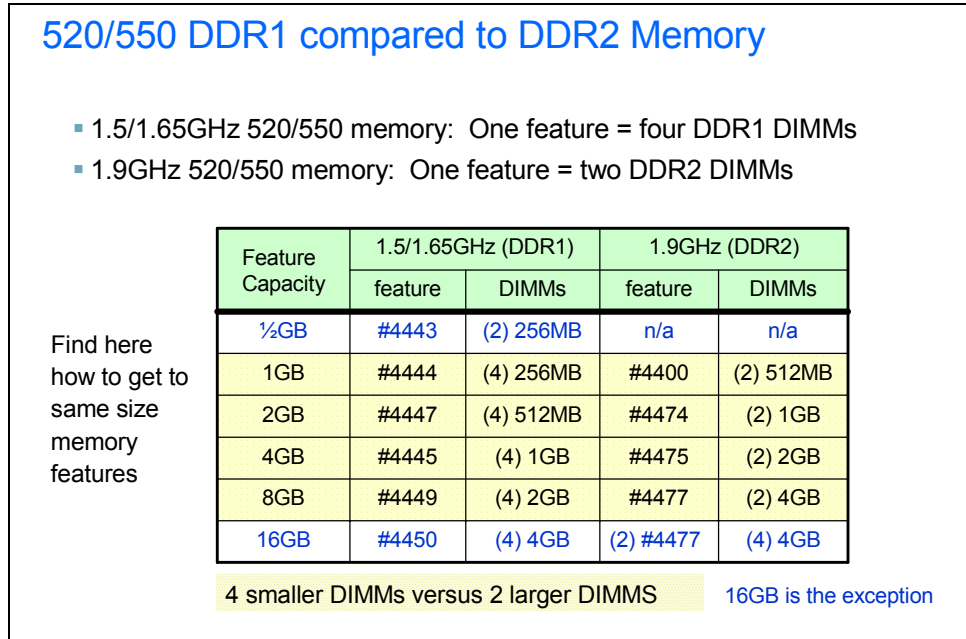


Figure B-1 Model 520 and 550 memory feature numbers

The way to read this table is to first select a total memory size increment in the first Feature Capacity column. Then for the DDR1 memory technology, find the feature number and note the quantity and size of the associated DDR1 DIMM cards. Then, for the corresponding memory size increment, note the quantity and size of the associated DDR2 DIMM cards. For example, you can see for the 8 GB memory capacity that you have feature numbers 4449 (DDR1) and 4477 (DDR2).

System i5 570 model memory technology considerations

The DDR1 memory cards are required on the 1.65 GHz 570 models (processor feature numbers 8961 and 8971). The DDR2 memory cards are required on the 2.2 GHz 570 models (processor feature number 8338).

With either DDR1 technology or DDR2 technology on the 570 models, one memory feature numbers delivers four DDR1/DDR2 DIMMs. Each processor card must contain a minimum of four DIMMs. Each processor card has eight DIMM slots.

As shipped from IBM the 570 model, DDR1 and DDR2 memory DIMMs are plugged on each available processor card appropriately for good performance under busy workloads.

Since additional memory cards may be ordered after the original configuration was installed, we indicate that better performance can be achieved when the DIMMs are plugged to that same/similar total MB of memory per processor card. On the 570 models, there are up to two processors per processor card and two processor cards per physical enclosure (one enclosure for up to four processors). This means that with up four physical enclosures, with two processor cards each yields up to 16 processors.

The better positioning of memory cards becomes more important as the workload increases in a manner such that the active threads need larger amounts of working storage, and processors on more than one processor card are used.

The DDR1 570 model memory features are:

- ▶ #4452 (2 GB) with 4 x 512 MB DIMMs
- ▶ #4490 (4 GB) with 4 x 1024 MB DIMMs
- ▶ #4454 (8 GB) with 4 x 2048 MB DIMMs
- ▶ #7890 (4/8 GB) with 4 x 2048 MB (2 GB) DIMMs (This is an on demand memory feature.)
- ▶ #4491 (16 GB) with 4 x 4096 MB DIMMs
- ▶ #4492 (32 GB) with 4 x 8192 MB DIMMs

The DDR2 570 model memory features are:

- ▶ #7892 (2 GB) with 4 x 512 MB DIMMs
- ▶ #7893 (4 GB) with 4 x 1024 MB DIMMs
- ▶ #7894 (8 GB) with 4 x 2048 MB DIMMs
- ▶ #4495 (4/8 GB) with 4 x 2048 MB (2 GB) DIMMs (This is an on demand memory feature.)
- ▶ #4497 (16 GB) with 4 x 4096 MB DIMMs
- ▶ #4499 (16 GB) with 4 x 4096 MB DIMMs
- ▶ #4496 (8/16 GB) with 4 x 4096 MB DIMMs (This is a memory on demand feature.)
- ▶ #4498 (32 GB) with 4 x 8192 MB DIMMs

Figure B-2 shows the available DDR2 memory feature numbers on the 2.2 GH model 570. It also indicates the range of memory feature numbers based upon the specific model minimal and maximum number of activated processor.

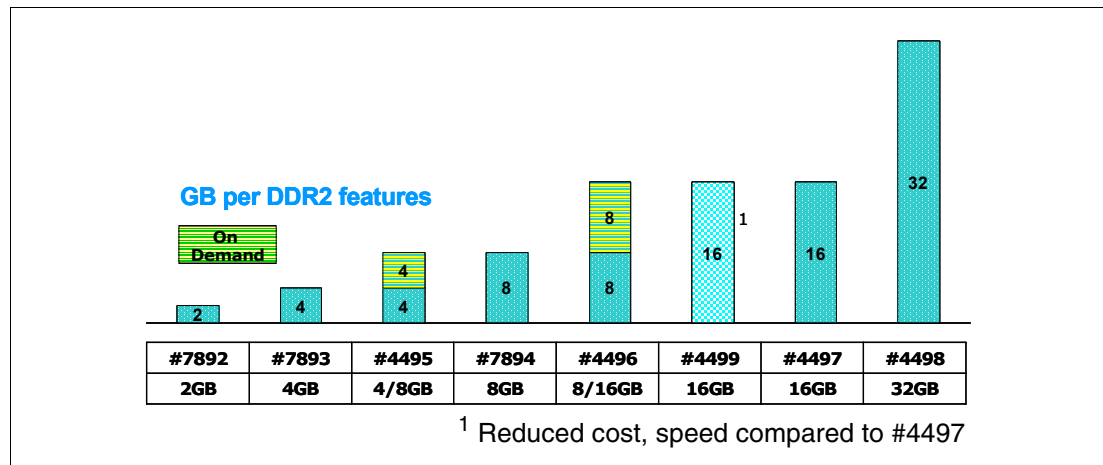


Figure B-2 570 model DDR2 available memory features

With restrictions, you can mix speeds (processor card feature numbers) in the same processor enclosure (on the same processor card or on different processor cards) and different processor enclosures (for example, two processor enclosures for a 5-8 Way 570 model). This is true except for the originally available 1/2-way 570 model, which only has one processor card.

Important: The following are important rules for mixing different memory cards and a related performance impact statement:

- ▶ Consider first the following general rules for DDR-1 and DDR-2:
 - For better performance, have the same amount of memory on each processor card.
 - Minimum of two features per processor enclosure.
 - Same memory feature maximums, governed by the 570 model maximum memory supported.
- ▶ Using DDR-1 memory:
 - Memory Feature Codes 4492, 4494, and 7049 can be mixed on the same processor card. They cannot be mixed with other memory feature codes on the same processor card.
 - Use of a memory feature code 4492, 4494, or 7049 within the same processor enclosure (CEC) reduces the memory speed to 200 MHz for all memory cards within the entire system.
- ▶ Using DDR-2 memory:
 - Memory Feature Code 4498 and 4499 can be mixed on the same processor card. Neither the 4498 nor the 4499 memory feature can be mixed with other DDR-2 memory feature codes on the same processor card.
 - Use of a memory feature code 4498 or 4499 within the same processor enclosure (CEC) reduces the memory speed to 400 MHz for all memory cards within the entire system.
- ▶ Plans for future memory upgrades should take into account the need to achieve improved performance without limiting future memory expansion.

The 16 GB feature #4499 and the 32 GB feature #4498 of the 2.2 GHz DDR2 570 memory cards run slower than the other DDR2 memory features — at up to a 400 MHz rate. The other DDR2 memory features operate at up to a 533 MHz rate. #4499 was introduced during October 2006. #4499 requires V5R3 or V5R4 and POWER5 firmware level SF240_261 or later. 261 is the latest level currently shipping as of October 2006.

Note that the price per GB of memory is lowest with the 2 GB feature and increases up to the 32 GB memory size. This cost per GB progression is similar among the DDR1 and DDR2 memory technologies. Consider this when estimating your memory requirements. Note that there is no *trade-in* value if, when adding new memory features, you want to replace the lower cost, lower memory size DIMMs.

32 GB feature memory speed considerations

More memory will easily overcome any potential performance impact of using lower MHz DIMMs, so any performance lost by lower MHz memory is overcome by using a larger amount of memory. Always consider the relative cost per GB of memory card. You must consider this in estimating your memory requirements. So as long as you have adequate memory and there are empty memory slots for future growth, there is little reason for most configurations to use the 32 GB feature.

Assuming that the system is not memory constrained, developers predicted that most System i5 customer applications would not see much of a difference in overall commercial processing environments. The speed difference could be noted in applications such as scientific applications that may manipulate huge matrixes.

570 model on demand memory feature considerations

On demand memory features for DDR2 memory work exactly like the DDR1 on demand memory features:

- ▶ You need memory activation features (one per 1 GB activation). For DDR2, this is feature 7663.
- ▶ On/off memory capability is billed in 1GB days.
- ▶ An HMC is required to activate or deactivate the 1 GB memory increments.
- ▶ Memory activations are unique to the memory technology and unique to the specific system.

Memory activations are unique to the memory technology, for example, between DDR1 and DDR2 DIMMs on a 570. In other words, memory activations purchased for DDR1 memory are not present after a 1.65 GHz model 570 is upgraded to a 2.2 GHz model 570 or a model 595.

Memory activations are also unique to the specific system, not a specific memory DIMM or memory feature. This provides the system with a total activated memory quantity, which the operating system can then select to better relate physical processors doing work and the physical memory located closest to those physical processors. It also means that if you move on demand memory to a different model 570, only base (first half capacity) of on demand memory feature will be usable by other 570 unless one of the following happens:

- ▶ New memory activations are purchased for the second system.
- ▶ There are pre-existing activation features on the second system, which were not being used.

Any original memory activations remain on the first system and can be used for any new on demand memory features (of the same technology) that are added to the machine.

System i5 595 model memory technology considerations

The 595 model DDR1 and DDR2 memory card considerations offer more flexibility, and thus rules are a little more complex than on the 570 models. Thus, we first describe 595 DDR1 memory considerations, then 595 DDR2 considerations, and then provide some positioning of DDR1 and DDR2 memory usage on 595 models.

595 DDR1 memory

Figure B-3 summarizes DDR1 memory card features and memory slot rules for the 595 model.

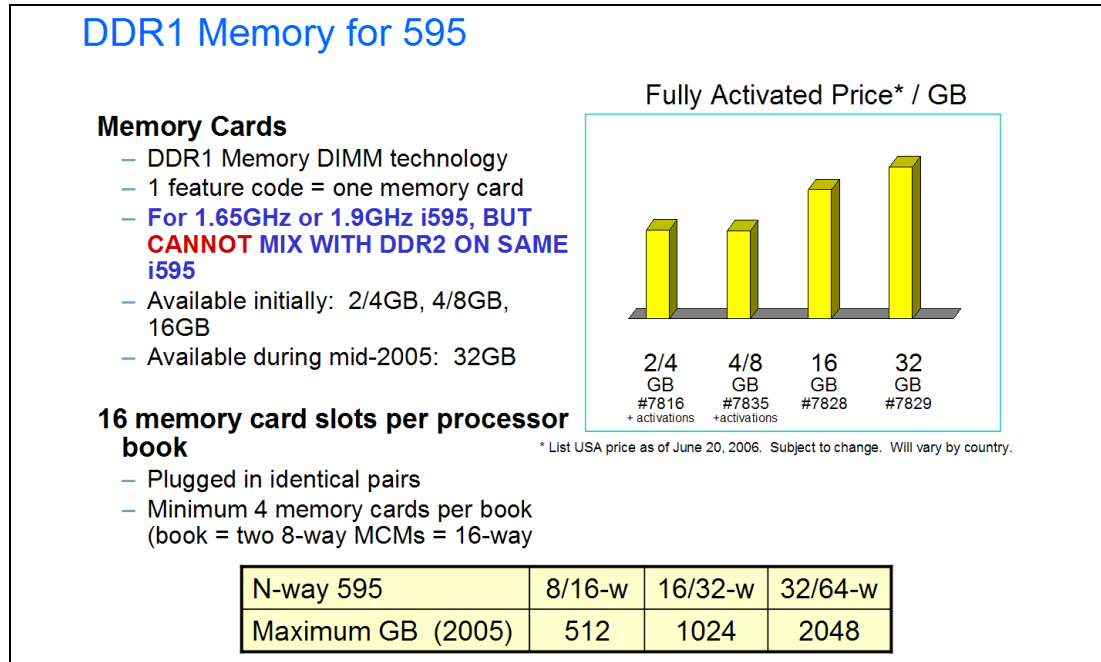


Figure B-3 595 model DDR1 memory summary

You can see for DDR1 on the 595 model:

- ▶ One feature equals one memory card.
- ▶ You cannot mix DDR1 memory cards with DDR2 memory cards.
- ▶ There are 16 memory card slots per 595 processor book (two 8 processor per Multi-Chip Modules (MCMs) per book).
- ▶ The memory cards must be plugged as identical pairs.
- ▶ An HMC is required to activate on demand memory.

This figure summarizes the 595 DDR1 memory cards and plugging rules. It is provided here to offer a contrast to the DDR2 memory made available mid 2006. The graph shows the relative USA list price (subject to change, and actual price varies by country), assuming all memory features are activated. The table shows the maximums of 595 memory. Note that the maximum changed in 2005 based on the availability of a 32 GB memory feature.

The 595 uses industrial strength DDR1 memory, using very strong Error Correction Code (ECC), redundant bit steering and other reliability technologies. Unlike the 520/550/570 where one DDR1 feature code equals four DIMMs, each 595 feature code relates to one memory card. Initially, two on demand memory features, 2/4 GB and 4/8 GB, and a 16 GB memory feature were available. During 2005 a 32 GB memory feature became available. This feature provided the maximum memory size of 2 TB.

The 16 GB and 32 GB features are treated like an on demand memory feature of 8/16 GB and 16/32 GB if moved to a different machine.

Similar to the 570 memory, the price of these memory features changes on a per GB basis. Bigger memory features are more expensive to manufacture. As a result, the 16 and 32 GB memory features are considerably more costly on a per-GB basis.

The plugging rules for the 16 memory slots per 595 processor book are based on identical pairs. The underlying rule is that you must have one pair of memory cards per MCM, which translates to a minimum of four memory cards per processor book.

The 595 DDR1 memory cards features are:

- ▶ #7816 - 2/4 GB - 2 GB on demand, 2 GB on demand if move
- ▶ #7835 - 4/8 GB - 4 GB on demand, 4 GB on demand if move
- ▶ #7828 - 16 GB - ship 16 GB active, 8 GB on demand if move
- ▶ #7829 - 32 GB - ship 32 GB active, 16 GB on demand if move (available 2Q05)

On demand memory features (1GB DDR1 memory increments):

- ▶ #7970 charge, #8460 no charge: CUoD (permanent activation)
- ▶ #7799 charge, #8467 no charge: Package of 256 DDR1 activations
- ▶ On/off Capacity on Demand
 - 1 GB memory day billing increments
 - Not included in Reserve Capacity on Demand

Similar to 570 DDRn model rules, the activations are for a specific system.

595 DDR2 memory

Figure B-4 summarizes DDR2 memory card features and memory slot rules for the 595 model.

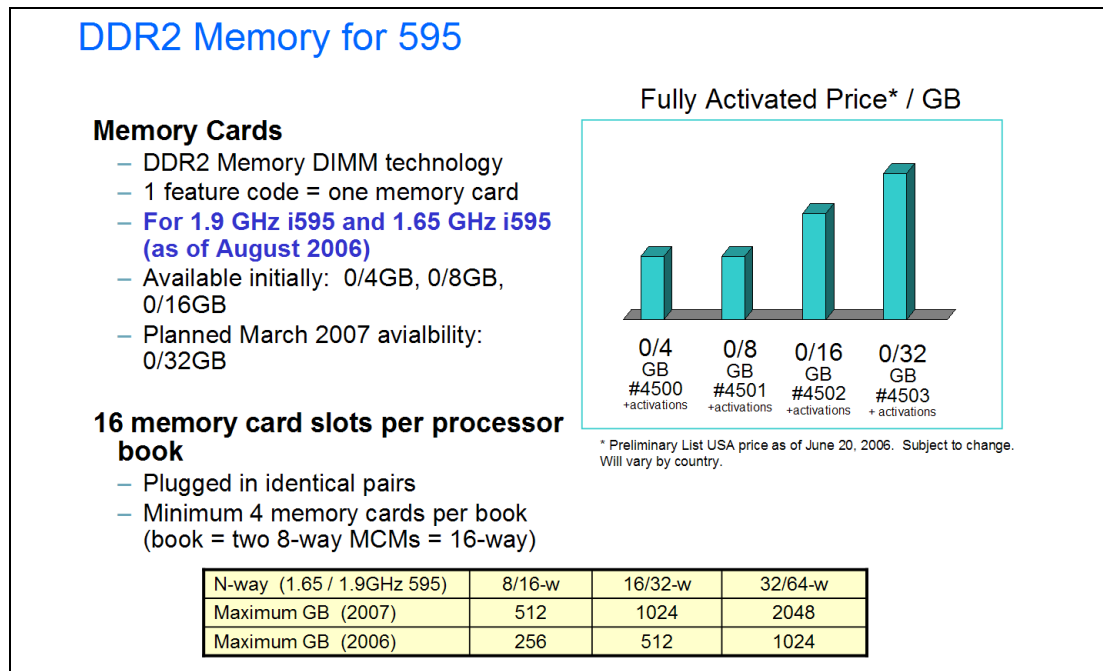


Figure B-4 595 model DDR2 memory summary

Similar to DDR1 technology, you can see for DDR2 on the 595 model:

- ▶ One feature equals one memory card.
- ▶ You cannot mix DDR1 memory cards with DDR2 memory cards.
- ▶ There are 16 memory card slots per 595 processor book (two 8-processor per Multi-Chip Modules (MCMs) per book).
- ▶ The memory cards must be plugged as identical pairs.

- ▶ An HMC is required to activate on demand memory.

This figure summarizes the 595 DDR2 memory cards and plugging rules. The graph shows the relative USA list price (subject to change, and actual price varies by country) assuming all memory features are activated. The table shows the maximums of 595 memory. Note that the maximum changes based on the availability of a 0/32 GB memory feature, which has a 2007 planned availability.

The 595 uses industrial strength DDR2 memory, using very strong ECC, redundant bit steering, and other reliability technologies.

Unlike the 520/550/570 where one feature code equals two or four DIMMs, each 595 feature code relates to one memory card. Initially, two on demand memory features, 0/4 GB and 0/8 GB, and a 0/16GB memory feature, are available. The 0/16GB feature is not really on demand when purchased from IBM manufacturing, as 16 GB of activations are required, but is fully on demand if moved to a different machine.

In 2007, a 0/32 GB memory feature will be available. This feature will provide the maximum of 2 TB.

Like the 570 memory, the prices of these memory features change on a per-GB basis. Larger memory features are more expensive to manufacture. As a result, the 0/16 and 0/32 GB memory features are considerably more costly on a per-GB basis.

The plugging rules for the 16 memory slots per 595 processor book are based on identical pairs. The underlying rule is that you must have one pair of memory cards per MCM, which translates to a minimum of four memory cards per processor book.

The 595 DDR2 memory card features are:

- ▶ #4500 - 0/4 GB: There is a minimum of 2 GB activations when purchased from IBM. No GB activations can be moved to another system.
- ▶ #4501 - 0/8 GB: There is a minimum of 4 GB activations when purchased from IBM. No GB activations can be moved to another system.
- ▶ #4502 - 0/16 GB: There is a minimum of 16 GB activations when purchased from IBM. No GB activations can be moved to another system.
- ▶ #4503 - 0/32 GB: There is a minimum of 32 GB activations when purchased from IBM. No GB activations can be moved to another system. 32 GB storage is planned to be available March 2007.

On demand memory features (1 GB DDR2 memory increments):

- ▶ #7669 charge, #8494 no charge: CUoD (permanent activation) 1 GB DDR2 memory increments
- ▶ #7280 charge, #8495 no charge: package of 256 DDR2 activations
- ▶ On/Off Capacity on Demand:
 - 1 GB memory day billing increments
 - Not included in Reserve Capacity on Demand

DDRn model rules specify that activations are for a specific system. Note that DDR1 offers some package features to reduce overall memory cost. For 595 model DDR2, you must pursue a special bids process.

595 model memory card placement performance considerations

When using either DDR1 or DDR2 memory technology, balanced memory optimizes performance:

- ▶ Acceptable performance: minimum of four memory cards per 16-way book (required).
- ▶ Good performance: same or similar total GB of memory on each processor book.
- ▶ Better performance: same or similar total GB of memory on each MCM. It is slightly better to not mix 32 GB card pairs with other size card pairs in the same processor book.
- ▶ Best performance: all memory slots filled with same size memory cards.

Note that moving from *OK* to *good* is highly recommended for performance, but moving to *better* and *best* will not provide a significant performance improvement. In other words, do not expect as much percentage improvement going from better to best as received by going from OK to good.

As stated for 550 and 570 models, optimal performance can be noticed on a system with a large number of active threads.

Also, though DDR2 offers faster (MHz) memory speeds, generally, this faster memory does not dramatically improve performance compared to similarly plugged DDR1 memory cards. Performance improvements are estimated to be less than 3% unless the application workload is performing large matrix manipulations, such as weather forecasting. In such an environment more than 3% improvement may occur.

Table B-1 shows the DDR1 and DDR2 memory speeds as of August 2006.

Table B-1 DDR1, DDR2 memory speeds

Technology/ MHz	200 MHz	266 MHz	400 MHz	533 MHz
DDR1	32 GB	2/4 GB, 4/8 GB, 16 GB		
DDR2			0/32 GB	04/GB, 0/8 GB, 0/16 GB



C

System i5 Capacity Backup edition software-based details

In this appendix we provide details on Capacity Backup editions in addition to those presented in 2.20.4, “Capacity Backup (CBU) Editions” on page 61. This appendix assumes that you have read that information before reading this appendix.

System i CBU edition software details

In this appendix we provide details on the Capacity Backup edition offerings available through August 2006. Note that the following levels of Licensed Internal Code (LIC), as indicated on the LIC media label, are required to support these CBU offerings:

- ▶ LIC V5R3M0 RS 530-K (RSK)
- ▶ LIC V5R3M5 RS 535-E (RSE)
- ▶ LIC V5R4M0 RS 540-D (RSD)

Capacity Backup edition example

Figure C-1 shows an example of a configuration with a primary IBM System i5 570 configuration with eight base processors, one i5/OS license, and one Enterprise Enablement (EE - 5250 OLTP capacity). The customer wishes to enable seven processors on the CBU model 570 2/16-way with the original base capabilities of two processors permanently activated, one base i5/OS license, and one Enterprise Enablement feature.

You must purchase six additional permanent processor activations for the CBU configuration. All prices are US prices and subject to change without notice. Prices in each country may vary.

When you register at the time the CBU edition is ordered, register that you want to enable seven CBU processors for temporary transfer of i5/OS license entitlement and Enterprise Enablement (EE) features.

In this example you see seven i5/OS licenses and three EE features being transferred.

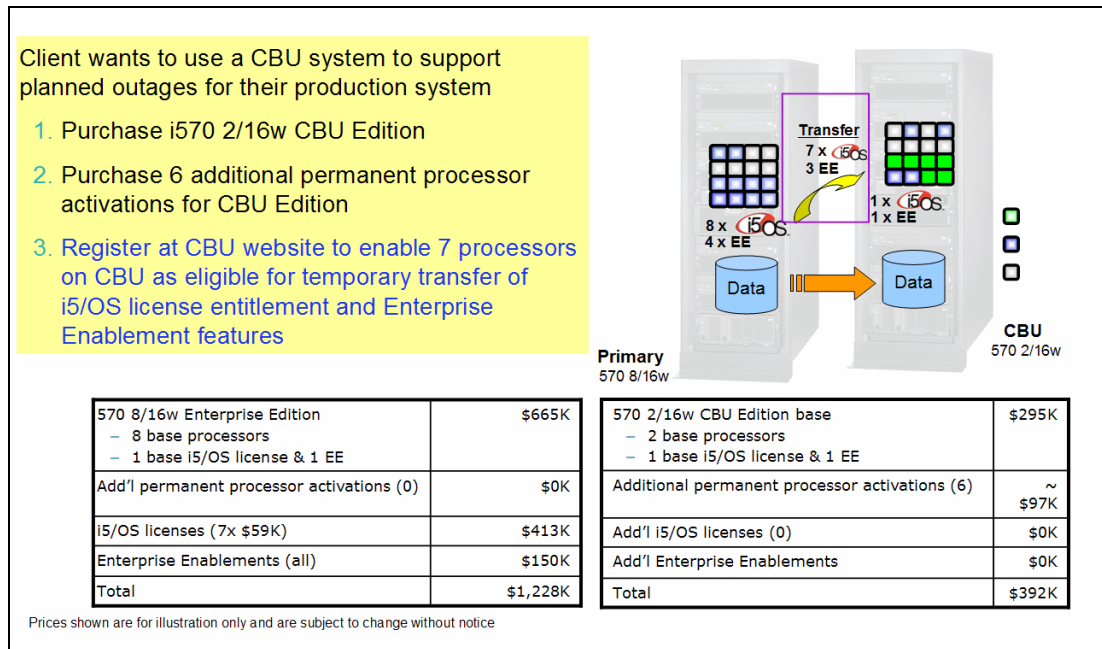


Figure C-1 Capacity Backup edition example - System i5 570 primary and CBU systems

Capacity Backup edition offerings

Figure C-2 shows the CBU Edition and CBU Standard Edition offerings available as of August 2006.

	CBU Edition		CBU Standard Edition*
	i5/OS	Enterprise Enablements	i5/OS
595 4/32-way	4	4	4
595 2/16-way*	2	2	2
570 2/16-way	1	1	1
570 1/8-way*	1	1	1
570 1/4-way*	1	1	1
550 1/4-way*	1	1	1

Figure C-2 CBU edition i5/OS and Enterprise Edition options

An asterisk (*) indicates a product preview capability. A product preview provides insight into IBM plans and directions. All statements regarding IBM plans, directions, and intent are subject to change without notice.

Permanent processor activations (beyond the base) may now be purchased on a CBU system to support additional application workloads (for example, development and business intelligence).

IBM i5/OS license entitlements, Enterprise Enablement features, and IPLA software licenses to primary system processors may be transferred (assuming proper system registration has been completed) to the CBU system while the respective primary system processors are inactive. This enables a CBU system to provide failover/role-swap for a full range of test, disaster recovery, and high-availability scenarios.

Key basic CBU Edition offerings include:

- ▶ Function
 - Support for multiple operating systems (i5/OS, AIX 5L, and Linux)
 - Support for Web modernization capabilities including enhanced WebFacing support and Host Access Transformation Services (HATS) support
 - Support for Dynamic Logical Partitioning (up to 10 i5/OS partitions per processor)
 - Enterprise Enablement processor entitlements: one for the 550, one for the 570, four for the 595
 - Support for Capacity on Demand
- ▶ Software
 - Licensing for i5/OS (550 one processor license entitlement, 570 one processor license entitlement, 595 four processor license entitlements)
 - e-business Solution Tools: none included
 - Datacenter Management Tools: none included
- ▶ Services/Education
 - Vouchers: none included for the CBU Edition

Key CBU Standard Edition offerings include:

- ▶ **Function**
 - Support for multiple operating systems (i5/OS, AIX 5L, and Linux)
 - Support for Web modernization capabilities including enhanced WebFacing support and Host Access Transformation Services (HATS) support
 - Note that for HATS to not require 5250 OLTP capability, you need i5/OS V5R4 and IBM WebFacing Deployment Tool for WebSphere Development Studio V6.0.1 with HATS Technology. See 12.10, “WebFacing and HATS” on page 307 for more information about HATS not requiring 5250 OLTP on i5/OS V5R4.
 - Support for Dynamic Logical Partitioning (up to 10 i5/OS partitions per processor)
 - Enterprise Enablement: none included
 - Support for Capacity on Demand
- ▶ **Software**
 - Licensing for i5/OS (550 one processor license entitlement, 570 one processor license entitlement, 595 four processor license entitlements)
 - ebusiness Solution Tools: none included
 - Datacenter Management Tools: none included
- ▶ **Services/education**
 - Vouchers: none included for the CBU Standard Edition

Positioning HA and CBU editions

With the August 2006 announcements, additional positioning among High Availability (HA) and Capacity Backup editions is required.

Figure C-3 shows key HA and CBU configuration considerations and attributes and how the HA and CBU editions compare in these areas.

	HA Edition	Earlier CBU Edition	Enhanced CBU Edition
Disaster Recovery	Yes	Yes	Yes
High Availability	Yes	No	Yes
Solutions/options which allow edition to be ordered:			
HA ISV Solutions	Yes	Yes	Yes
Independent Auxiliary Storage Pools	No	Yes	Yes
Cross Site Mirroring (XSM)	No	Yes	Yes
IBM DS8000	No	Yes	Yes
Permanent processor activations	Yes	No	Yes
Temporary i5/OS, AIX 5L, Linux entitlement transfers	No	No	Yes
Temporary enterprise enablement entitlement transfers	No	No	Yes
Pre-requisite of a 570 or 590 as primary system	No	No	Yes
Upgrade into edition from 825/870/890	No	No	Yes
Registration of primary system to qualify	Yes	No	Yes
No charge On/Off CoD if primary system down by disaster	No	Yes	Yes
Primary system must be same size or larger	Yes	No	Yes

Figure C-3 Multiple system HA and CBU attribute comparisons

Figure C-4 lists the HA/DR Independent Software Vendor (ISV) enhanced CBU companion offers and the URLs to find out more information about them.

HA/DR ISV	Companion Offer	Website URL
DataMirror	DataMirror iCluster for CBU	www.datamirror.com/cbu
iTera	Echo ² for CBU	www.iterainc.com/cbu
Lakeview Technology	MIMIX CBU Express	www.lakeviewtech.com/cbu
Maximum Availability	<i>Not defined August 2006</i>	www.maximumavailability.com
Trader's	Quick EDD CBU Edition 2006	www.quick-EDD.com
Vision Solutions	Orion for CBU	www.visionsolutions.CBUOffer.com

Figure C-4 HA/DR ISV enhanced CBU companion offers



IBM Support Assistant

In this appendix we provide an overview of an enhanced support tool, the IBM Support Assistant.

The IBM Support Assistant

The IBM Support Assistant (ISA) is a very useful tool for simplifying use of IBM software products, especially in WebSphere Application Server environments. This is a free Eclipse-based tool, which can be installed on AIX, HP-UX, Linux, Solaris, and Windows systems. You can download the software from the Web site:

<http://www.ibm.com/software/support/isa/>

With the ISA, you can get to information quickly. ISA provides a concurrent search tool, which spans across most IBM documentation. The returned results are categorized by source for easy review.

This, too, provides a product information feature, with key product information links. These include:

- ▶ Product support pages
- ▶ Product education roadmaps and the IBM Education Assistant
- ▶ Product home pages
- ▶ Product recommended updates
- ▶ Product troubleshooting guides
- ▶ Product newsgroups and forums

Included in ISA is a Tool workbench. This is the same problem determination toolset used by IBM Support. Future enhancement plans will include more tools of this nature.

Installation and setup

This product is easy to use and set up:

1. Download the product from the Web site and run the installer. The first time you do this, you get the window shown in Figure D-1, which includes the notice about not currently having any products installed and configured.

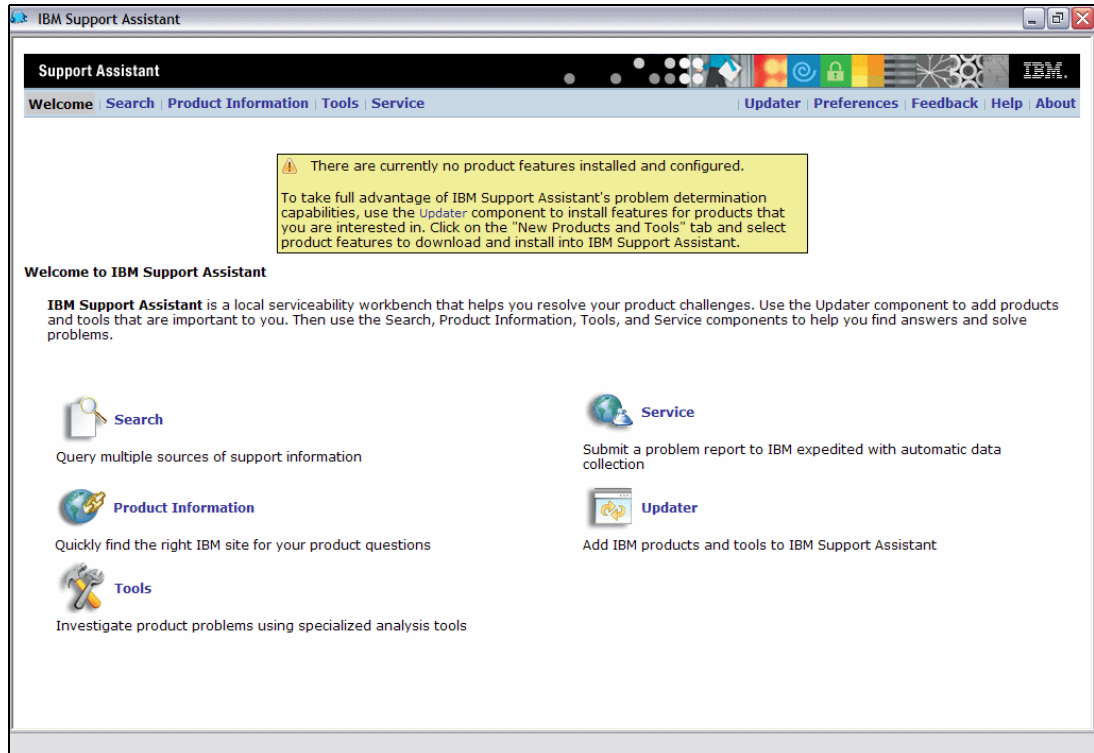


Figure D-1 ISA, initial screen

You can see from Figure D-2 that there are two areas of navigation. Along the top of the screen are drop-down menus, while the body has navigational icons.

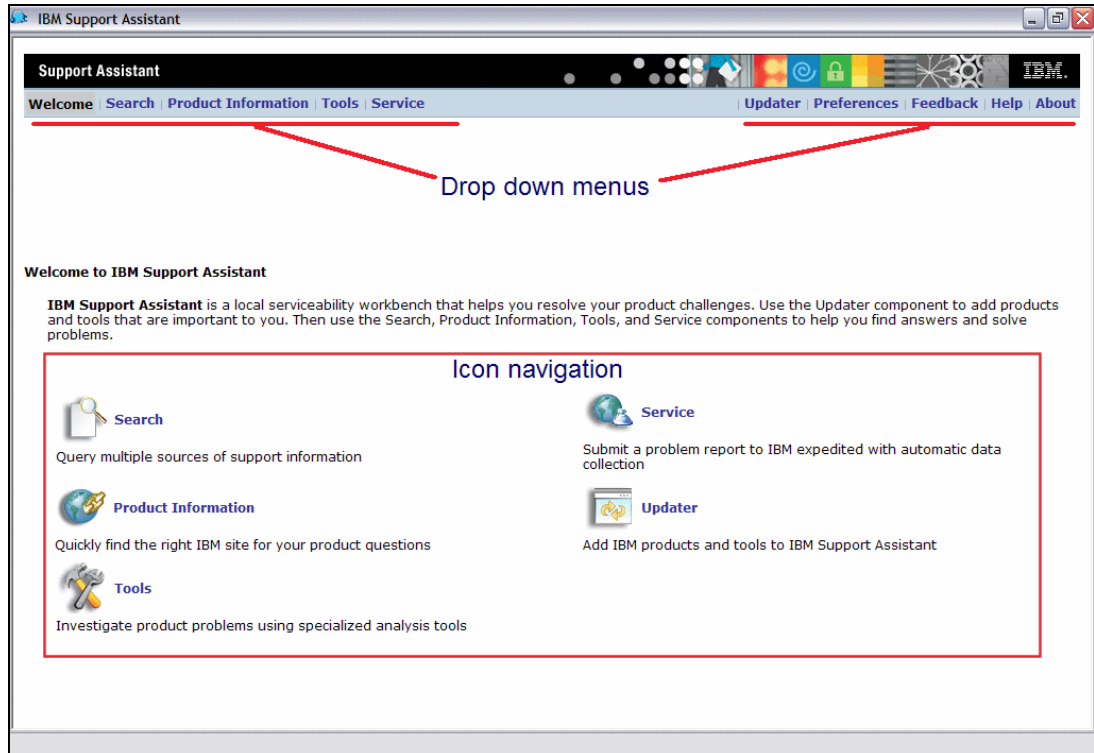


Figure D-2 Navigation options

2. The next step is to select the Updater. This opens the window shown in Figure D-3. From the left-hand column, you can select as many or as few options as you wish. Populating this screen can take some time, and the process bar in the lower right of the screen is rather dim.

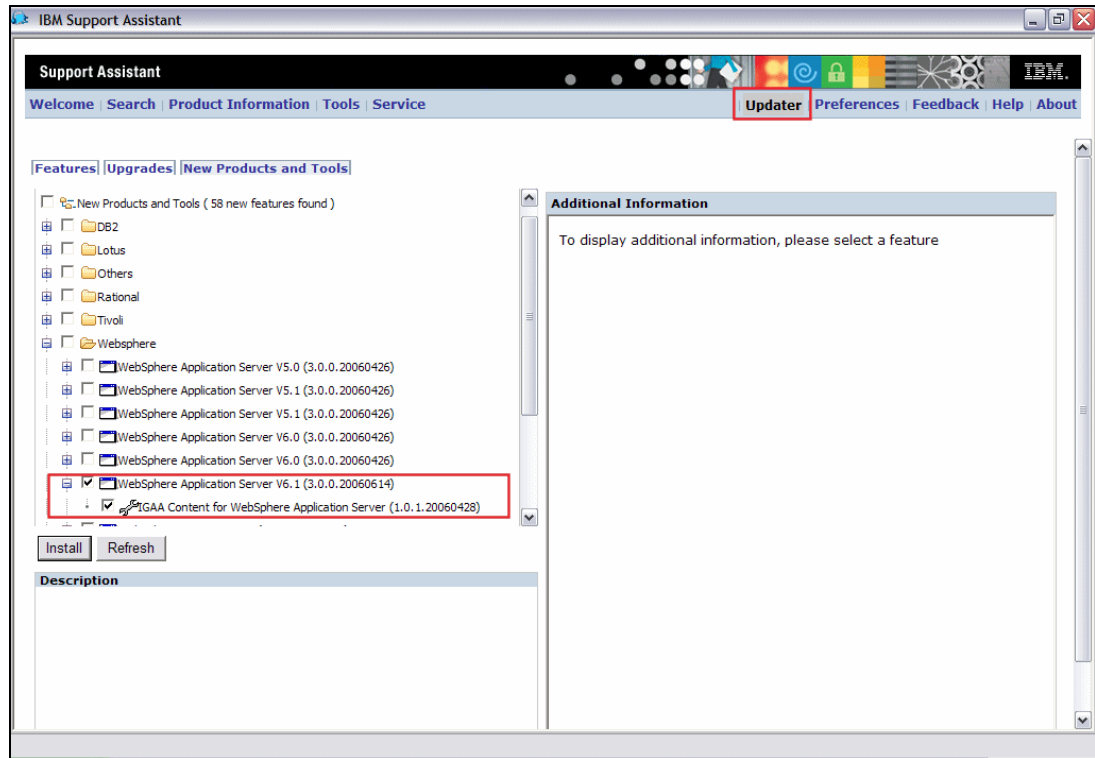


Figure D-3 Updater screen

After you have downloaded and installed your selections, which is a single process, restart the application. After the restart, the other selection links become populated with data.

Search

The first significant tool we discuss is the Search tab. You can see in Figure D-4 on page 520 that there are two sections to the search selector area. There is a default set:

- ▶ IBM Software Support Documents:
 - IBM downloads
 - IBM APARs
 - IBM books and articles
 - IBM Technotes
- ▶ IBM developerWorks
- ▶ IBM newsgroups and forums
- ▶ Google Web search

Then there is the section that has products. The inset in Figure D-4 shows an addition of Rational Application Developer 6.0 and WebSphere Application Server 6.0. This inset was captured after the update processed in Figure D-3 on page 519. We went back and added more products, which resulted in more potential search selections.

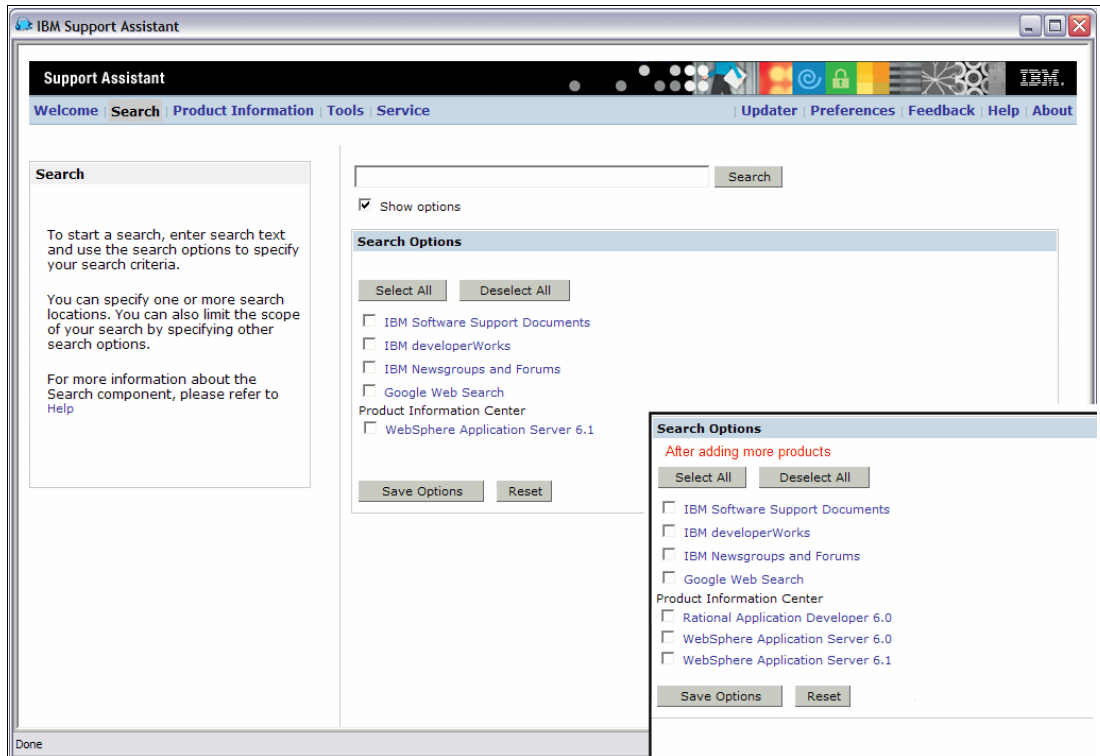


Figure D-4 Search

Product information

The Product Information screen, shown in Figure D-5, provides direct hot links to the selection. This helps you get to needed documentation quickly and with much less effort.

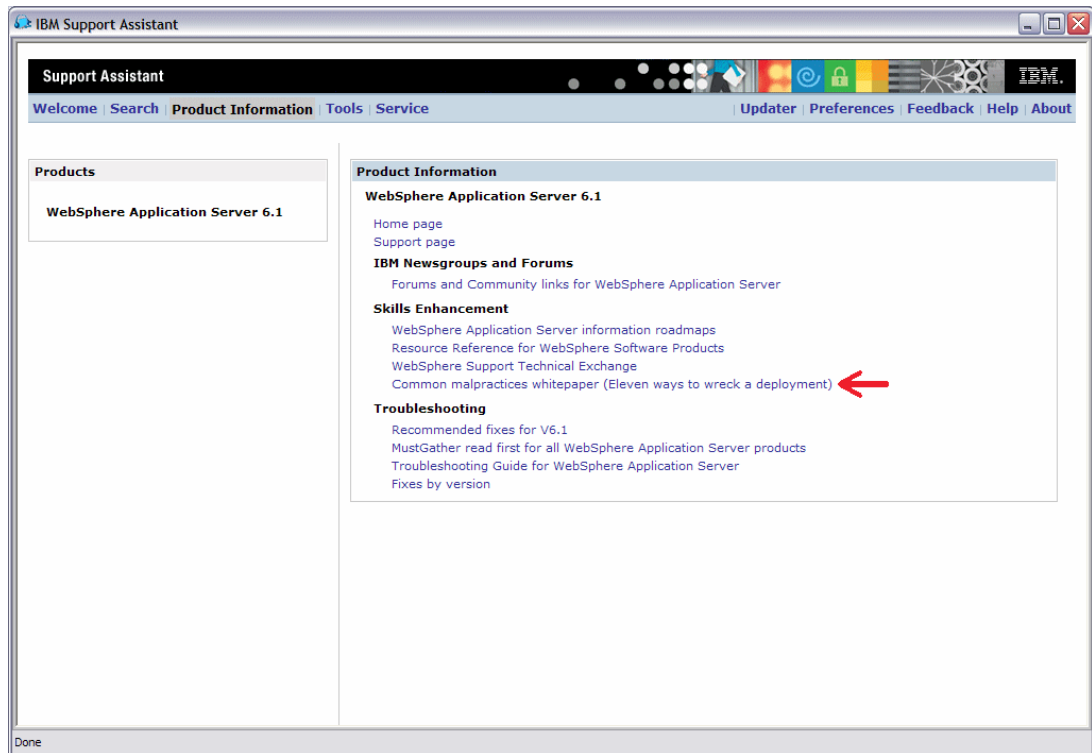


Figure D-5 Product information

For instance, there is a link to the “Common malpractices” white paper. This document was compiled by the WebSphere Serviceability Team from actual customer involvements. The paper explains the practices and patterns from years of work and evaluation. It is designed to give you an understanding of why and how trouble occurs, with possible solutions.

With the ISA, when you click that link, a window will pop up directly to the white paper at:

<http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg27007543>

The traditional approach is seen in Figure D-6. From the WAS home page, we selected **Support**, then selected **WebSphere Application Server**. In the search, we entered best practices. This returned 287 items, displayed ten per page. The white paper that we were after was on the sixth page, in position 64.

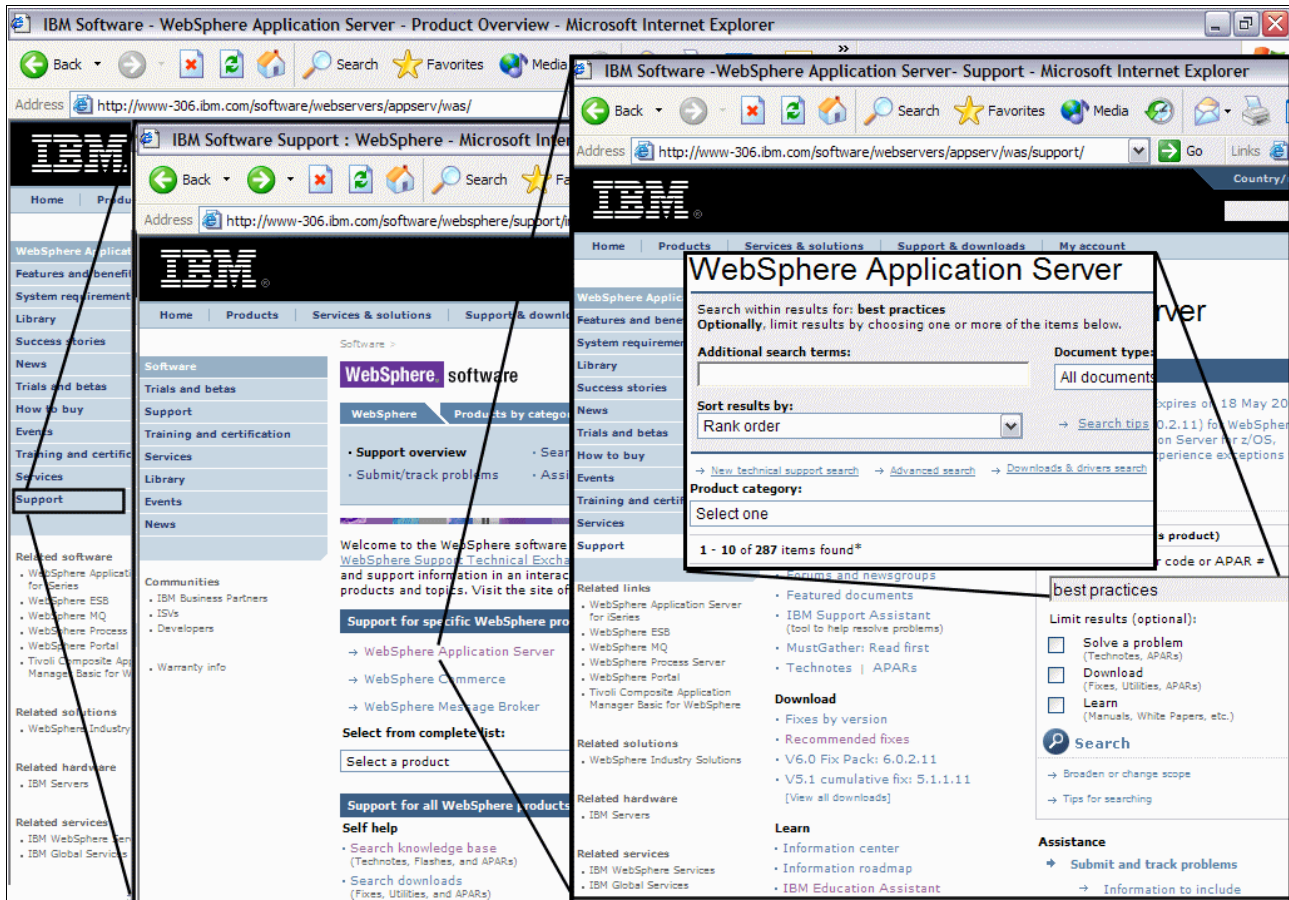


Figure D-6 Traditional search

Tools

The tools page (Figure D-7) populates to your installed products. Only products that have applicable tools will be selectable. When you select a link in the Products column, the Tools column populates.

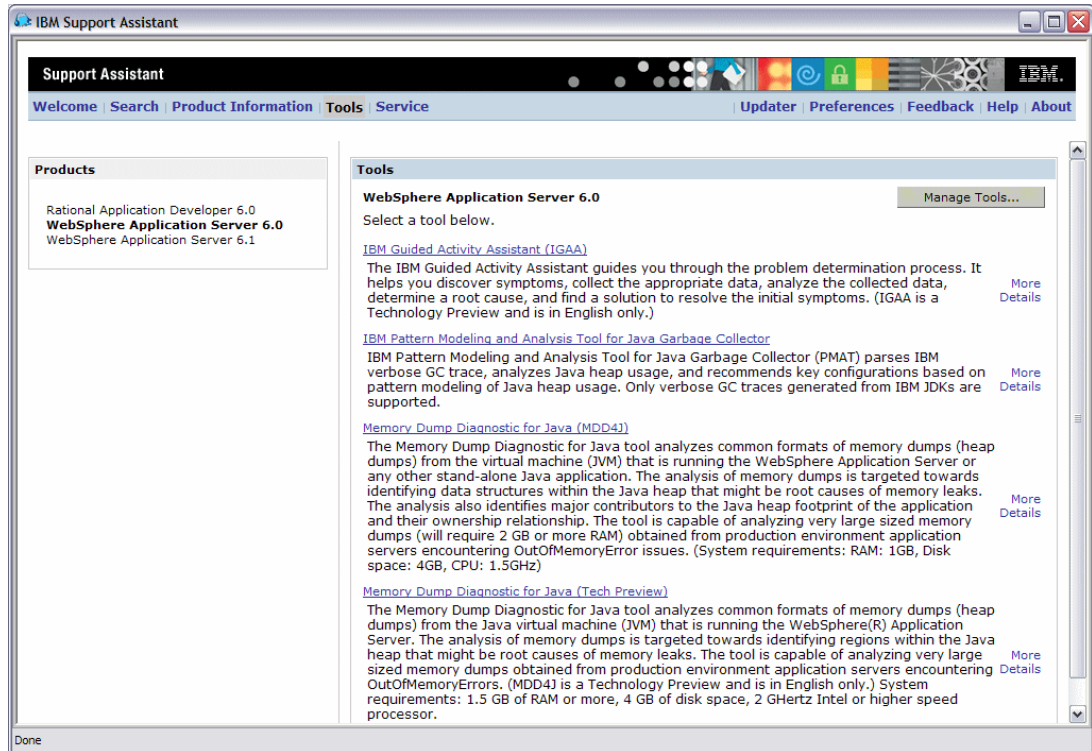


Figure D-7 Tools

If you select the IBM Guided Activity Assistant (IGAA), you get the tool in a separate window, as shown in Figure D-8.

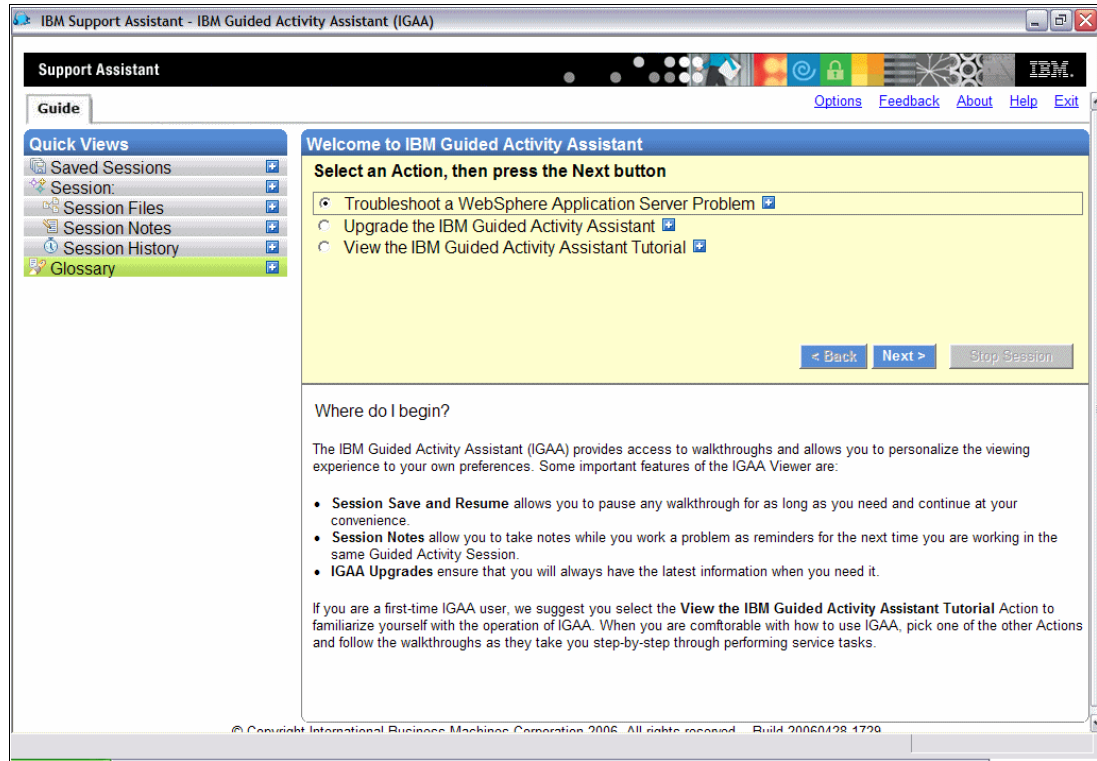


Figure D-8 The IBM Guided Activity Assistant

Support

Another very useful feature is the ability to submit a problem report directly from this tool. As you can see in Figure D-9, you can create and submit a new problem report, collect your diagnostic data, view and update an existing problems report, and send files for problem determination.

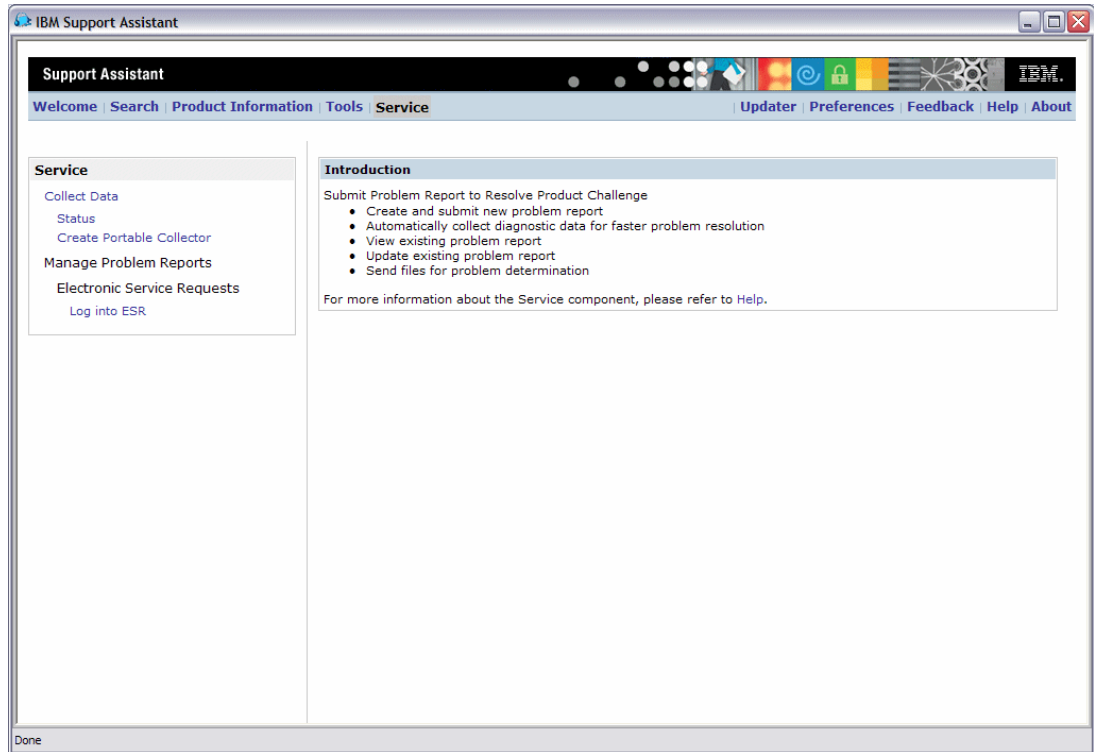


Figure D-9 Support Assistant

Help

There is a very useful set of online help text files for every aspect of the tool. Figure D-10 shows an example.

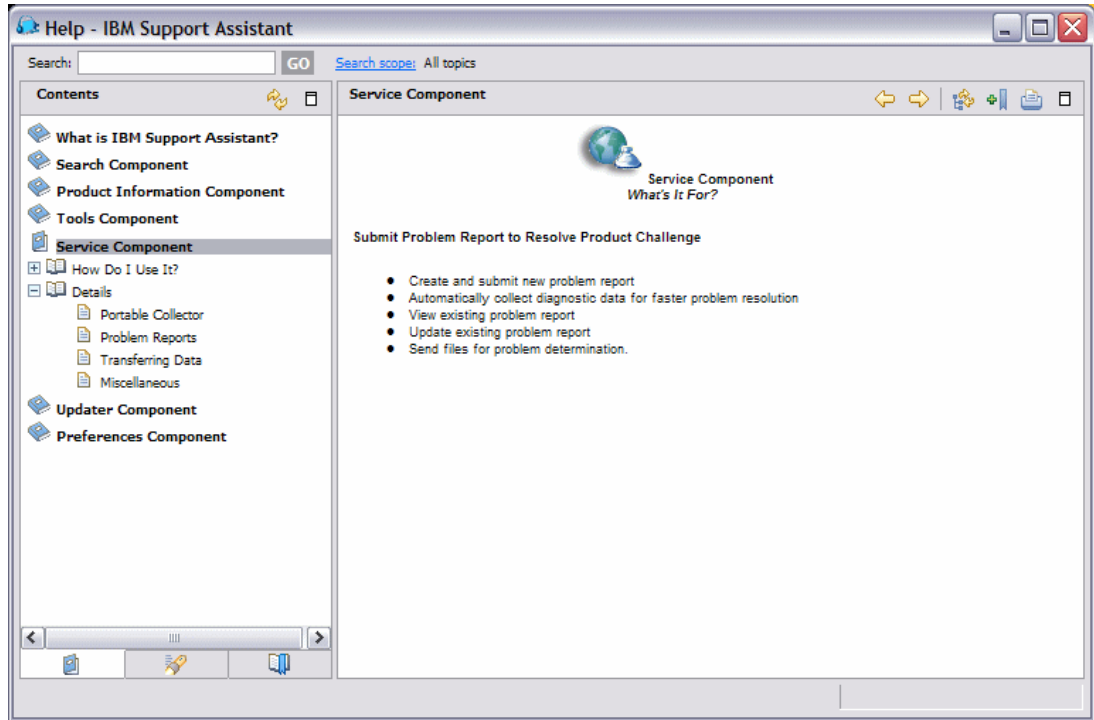


Figure D-10 Help screens

We expect you to find this tool very useful for a wide range of products.

IBM Education Assistant

A related tool is the IBM Education Assistant (Figure D-11). This tool provides narrated presentations, flash simulations, tutorials, and resource links. This tool is also based in the Eclipse Help System framework.

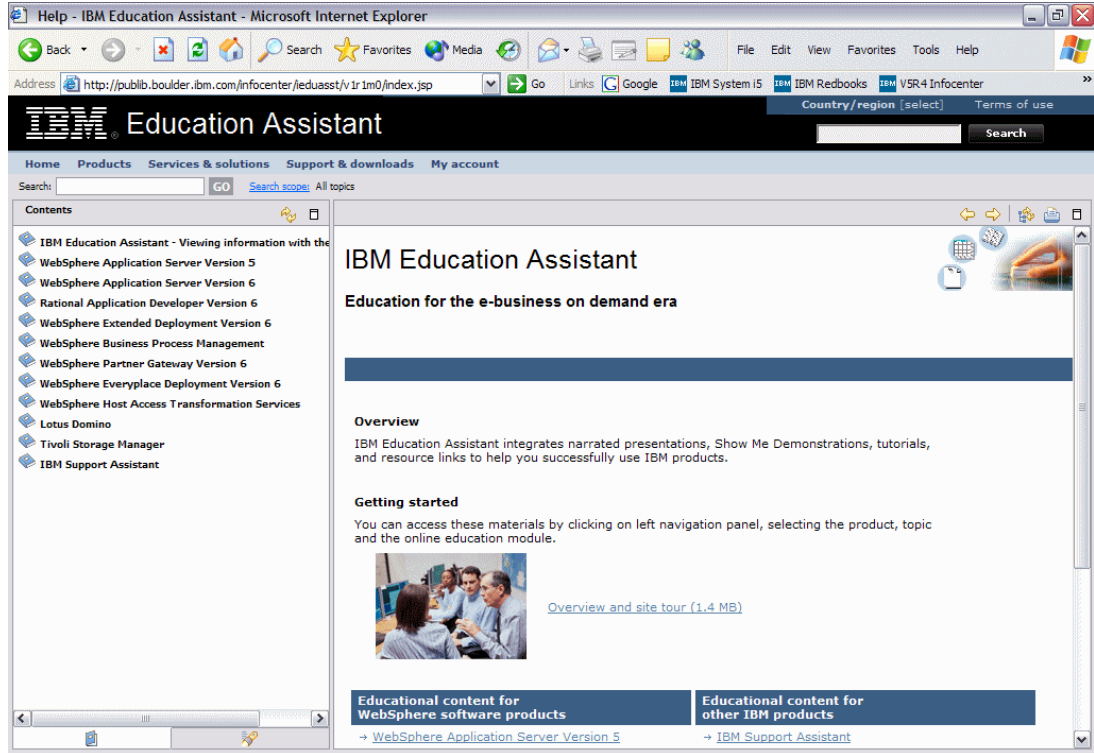


Figure D-11 IBM Education Assistant

Access IBM Education Assistant content at:

<http://www.ibm.com/software/info/education/assistant/>

Covered products include Lotus, Rational, Tivoli, and WebSphere software products, as well as systems and servers.

Using these tools can significantly expedite gaining knowledge on products available and how to use them productively.



System i IP Telephony powered by 3Com

In this appendix we provide a December 2006 level summary of IP Telephony support (also referred to as Voice over IP) products available from 3Com that run in a Linux partition on an IBM System i5 configuration.

For more complete information refer to:

- ▶ IBM Redbook *IBM System i IP Telephony*, SG24-7382, available at:

<http://www.ibm.com/redbooks>

- ▶ 3Com Web site at:

<http://www.3com.com>

Current System i information is at:

http://www.3com.com/index_jump/ibm_ipitel.html

- ▶ System i Telephony information at:

<http://www.ibm.com/systems/i/solutions/iptelephony/>

- ▶ Information available using various *Telephony* search words on the Google Web site

<http://www.google.com>

We have included excerpts and summarized information available from these sources in this appendix.

System i IP Telephony, Voice over IP, and 3Com

Voice over IP (VoIP) is the packetizing and transporting of voice traffic and signaling over a LAN or WAN using Internet Protocol (IP).

Use of VoIP Telephony technology is an important capability for businesses primarily in two areas — cost savings and business advantages:

- ▶ Cost savings can be realized by reducing infrastructure and operating costs because IP Telephony systems have a single infrastructure across multiple offices for voice and data, and these systems are typically easier to manage and change according to the needs of the business. Companies may also experience lower long-distance calling fees by eliminating toll charges.
- ▶ Business advantages are driven by increased collaboration for employees, customers, and partners. Advantages can be gained through:
 - Advanced voice applications such as:
 - Contact centers, where organizations can drive customer satisfaction and business revenues
 - Unified messaging, which includes voice mail/e-mail integration and find me/follow-me capabilities
 - Advanced collaborative applications such as:
 - Instant messaging
 - Presence awareness
 - Highly scalable audio conferencing capabilities

IP Telephony solutions are built on open architectures and industry standards.

There are several different enterprise-class protocols, or methods, of implementing IP Telephony technology that have been developed and implemented. These include H.323, Megaco, and Session Initiation Protocol (SIP). IP Telephony product vendors can implement their product using any of these protocols. Currently, SIP has emerged as the industry leader due to its relative simplicity, multi-media capability, and strong interoperability between devices and applications.

Telephony has real-time demands, which, relative to IP networks, require use of two relatively new IP network improvements:

- ▶ Quality of service packet prioritization services. For example, this can prevent local print jobs from interrupting conversations.
- ▶ Power over Ethernet (to eliminate power bricks at each IP phone). Though certainly not required in many Telephony configurations, power over the Ethernet may become more common in the future. Power over Ethernet technology allows many IP devices (such as wireless LAN access point, VoIP phone, Bluetooth Access point, network camera, and other network appliances) to receive power as well as data over existing LAN cabling, without needing to modify the existing Ethernet infrastructure.

This capability is based upon an international standard, called IEEE802.3af, originally available in 2003, as an extension to the existing Ethernet standards. Power Over Ethernet-capable products are available wherever needed. You can view this configuration as having your network appliance connected to the Power Over Ethernet-capable box, which in turn is connected to an Ethernet switch, through which data flow is managed.

This enables much flexibility that includes:

- Only one set of wires to connect to each appliance
- Easy movement of the appliance to another location
- An optional UPS connected to the Power Over Ethernet-capable box, to guarantee power to the appliance during local power failure.
- Use of SNMP network management infrastructure to monitor and control the appliance
- Network transmitted shutdown, reset, and power on control commands

Note: The information about power control over the Ethernet IP network is primarily based on the information found at the following Web site:

http://www.poweroverethernet.com/articles.php?article_id=52

During mid-2006, IBM and 3Com announced a joint project to port 3Com's Session Initiation Protocol (SIP)-based VCX platform to run in a Linux partition on a System i5 configuration. The 3Com suite of Telephony application product characteristics include:

- ▶ Open system - Standards Based.
- ▶ Linux and SIP.
- ▶ Audio, video, and data conferencing - supports multimedia communications.
- ▶ Instant messaging and Web collaboration.
- ▶ All Telephony features and functions in software.
- ▶ Phones receive calls and power over Ethernet, including *find and follow*."This allows you to get your phone calls on your same number regardless of where you are in the world.
- ▶ Enables simple integration with existing applications.

You can have interaction with the Public Switched Telephone Network (PSTN) all the way to the bare-metal detail of UDP traffic over firewalls. If you have connectivity to your IP network, you can get phone calls. Unified messaging allows you to send voice, video, and e-mail using the same connection. Multicast can easily broadcast information to multiple people, while conferencing lets you to pull multiple people into a single session. Because all of this is digital, it means that you can do it through your workstation: Drag and drop a person onto a conference call, and point and click to send a voicemail. You also have the ability for someone to point and click to dial an information or support line directly from your Web site.

Both 3Com IP Telephony and System i5 technology deliver simplicity, integration, and support for open standards with a mid-market focus. 3Com is making a significant investment to deliver a solution that runs on the System i platform that will provide the framework for future standards-based integration with collaboration and business processing applications running in i5/OS.

Voice over IP solutions and messaging platforms are included in the offerings running on a System i configuration. The i5/OS based-configuration characteristics include:

- ▶ An i5/OS partition serving as the virtual I/O server to one or more Linux partitions running a 3Com application. This enables use of partition-to-partition movement of processor and memory resources, virtual LAN connectivity, standard i5/OS network description objects, and network storage spaces. With the i5/OS partition providing virtual I/O support you can take advantage of an i5/OS-based backup of these objects using standard i5/OS save and restore commands or Backup Recovery and Media Services for iSeries (BRMS), 5722-BR1.

Note: Individual 3Com application file backup requires running 3Com-provided scripts to produce the resultant file. These files can be backed up using FTP to the i5/OS partition and placed into an i5/OS Integrated File System directory. Placed in such a directory, that file can be individually backed up using i5/OS commands or BRMS.

- ▶ The option to have a Linux partition run the 3Com application, and that partition manages its own I/O devices (direct I/O) and backup procedures.

Minimal configuration is one V5R3 or V5R4 i5/OS partition and one POWER Linux partition running the Telephony application. You can run this partitioned configuration using either:

- ▶ A System i5 with no HMC that is running Virtual Partition Manager (VPM), which can support one i5/OS partition and up to four POWER Linux partitions.
- ▶ A System i5 managed by an HMC. With this configuration, multiple i5/OS, POWER Linux, and AIX 5L V5.3 or later partitions can be configured and active provided that sufficient processor, main memory, and disk configurations are being used.

3Com offerings on System i configurations

There are two sets of 3COM Telephony support that will run within a Linux partition on a System i configuration:

- ▶ 3Com IP Telephony Suite for System i
 - This is the 3Com convergence application suite for IBM System i.
- ▶ 520 Telephony Express packages
 - Pre-sized and competitively priced for IP Telephony
 - Packaged System i hardware, i5/OS, and 3Com IP Telephony software and licenses
 - 100, 250, 500, 1000 Handsets
 - Flexible options for high availability

These are specific orderable IBM System i5 hardware configurations that come with i5/OS and the 3Com application suite already installed.

3Com IP Telephony Suite for System i

This suite of products runs only on IBM System i5 models 520, 550, 570, and 595 running i5/OS V5R3 or V5R4 and includes multiple functional 3Com software modules. The IBM PID number is 5639-3CM. 3Com offers software maintenance for this suite of products as IBM PID number 5771-3CM.

Based on end-to-end SIP signaling, the 3Com IP Telephony Suite for System i platform is comprised of the following software modules:

- ▶ IP Telephony
- ▶ IP Messaging
- ▶ IP Conferencing
- ▶ IP Presence

3Com products provided include the following with various numbers of licenses per specific System i configuration ordered:

- ▶ 3Com IP Telephony Server License
- ▶ 3Com IP Messaging Server License
- ▶ 3Com IP Messaging nnn Seat License

- ▶ 3Com IP Standard nnn Phone License
- ▶ 3Com Analog FXS Media Gateway 1 Port Licenses
- ▶ 3Com IP Telephony, Messaging, and Conferencing

Figure E-1 depicts System i Telephony with 3Com products installed.

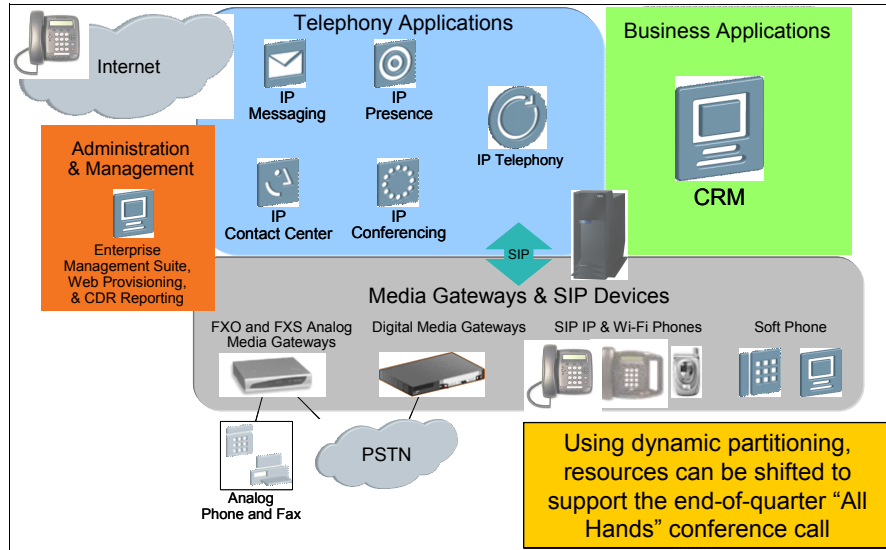


Figure E-1 System i IP Telephony Suite

This figure demonstrates two key points:

- ▶ System i IP Telephony uses built-in virtualization capabilities to run a full IP Telephony suite on a single system (versus each element of a suite taking one or more servers from the competition) and the business processing for collaboration workloads on a single system.
- ▶ With dynamic resource allocation, resources on the System i configuration can be shifted between applications in partitions to accommodate demand and increase resource utilization. In this example a large end-of-quarter conference call requires additional processing power, which is temporarily sourced from the CRM and contract center applications.

Figure E-2 depicts a basic configuration on i5/OS host partition and Linux Telephony application partitions with messaging and conference modules. i5/OS is the hosting partition for the Telephony servers. i5/OS provides the disk device space and Ethernet connectivity for the Telephony server partitions.

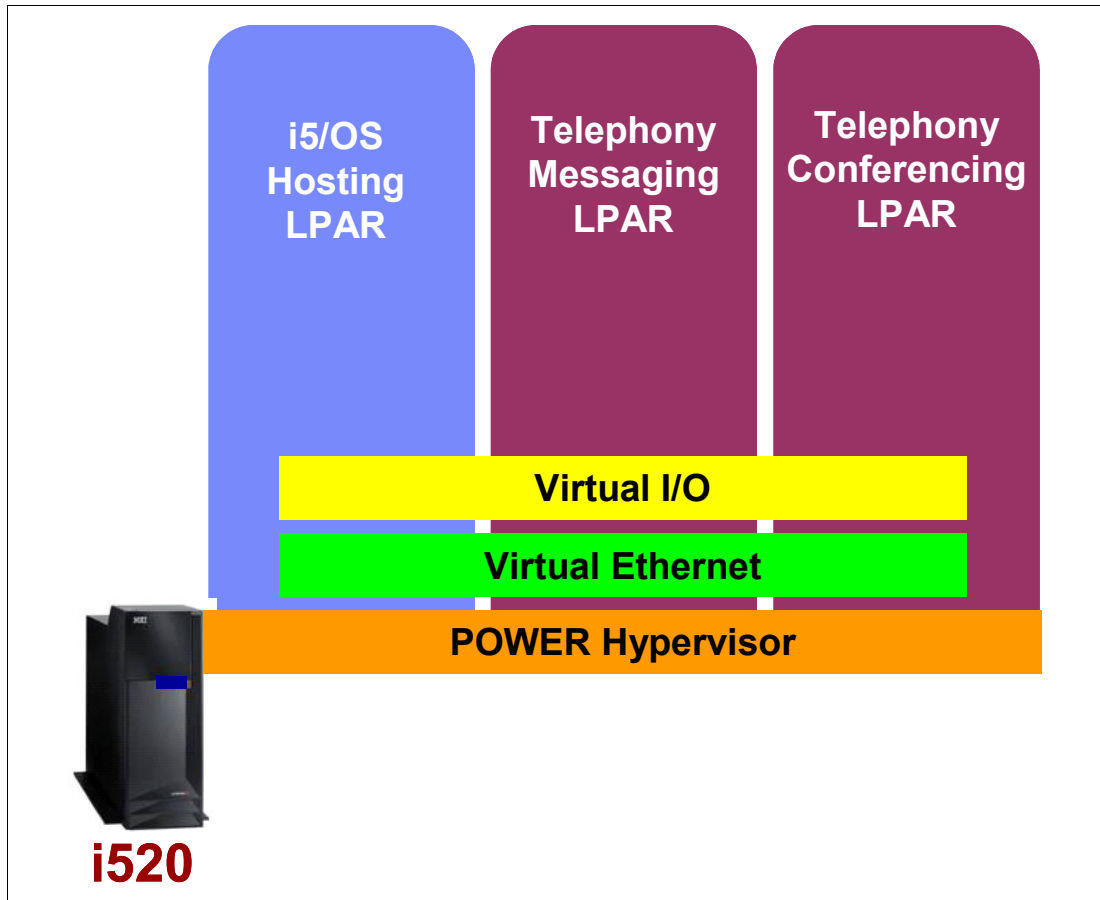


Figure E-2 Hosted Telephony with messaging and conference modules example

The Telephony servers use the same technology and i5/OS support as the hosted AIX and Linux partition on any System i5 model. No separate Linux license is required.

Because the Telephony environment must be up 24 hours, 7 days a week, a high availability capability configuration is available. The 3COM Telephony servers have built-in support for replication between a primary and a secondary server.

A High Availability (HA) configuration is strongly recommended, with primary and backup systems meeting your high availability requirements. In this environment, consider:

- ▶ Backup does not need to be 1:1 in multi-site environments.
- ▶ The backup machine may be stand-alone or shared with existing applications.

Figure E-3 depicts a simple two-system high availability configuration. Either server may be set to primary to allow maintenance on the other server.

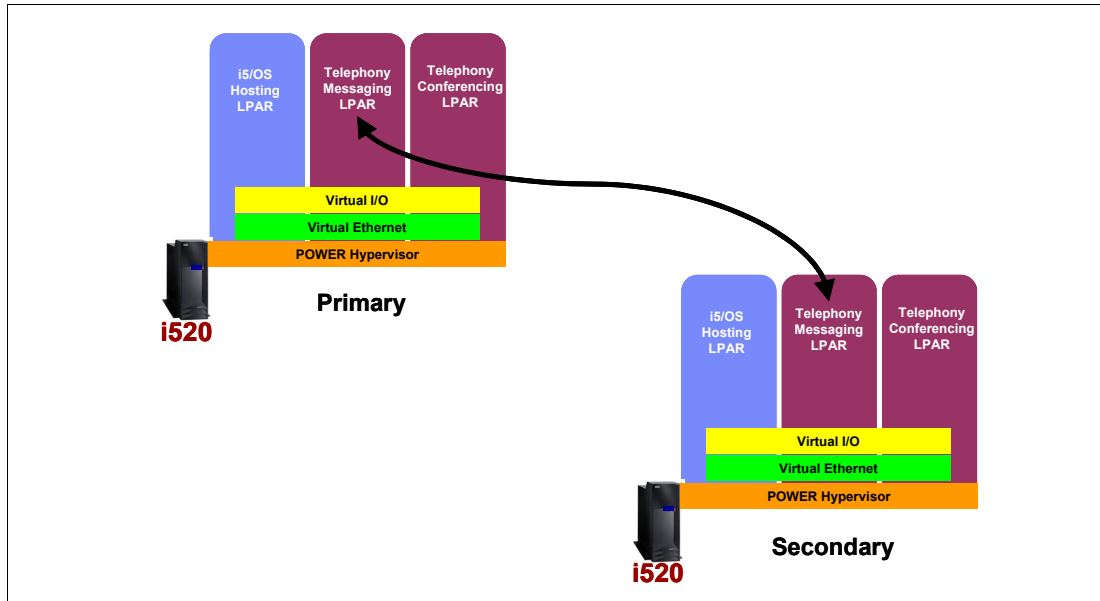


Figure E-3 HA Telephony using two-systems example

Other high availability configurations are possible, including a single system with multiple partitions configured to provide high availability capabilities on a single system.

IBM System i5 520 model Telephony Express packages

These are combined 520 model hardware configurations, i5/OS, and already installed 3Com Convergence Application Suite for IBM System i. These offer:

- ▶ Pre-sized and competitively priced for IP Telephony
- ▶ Packaged System i and 3Com IP Telephony software and licenses
- ▶ 100, 250, 500, 1000 Handsets
- ▶ Flexible options for high availability

Figure E-4 shows a summary level of the model 520 Telephony Express packages available for IBM System i5.

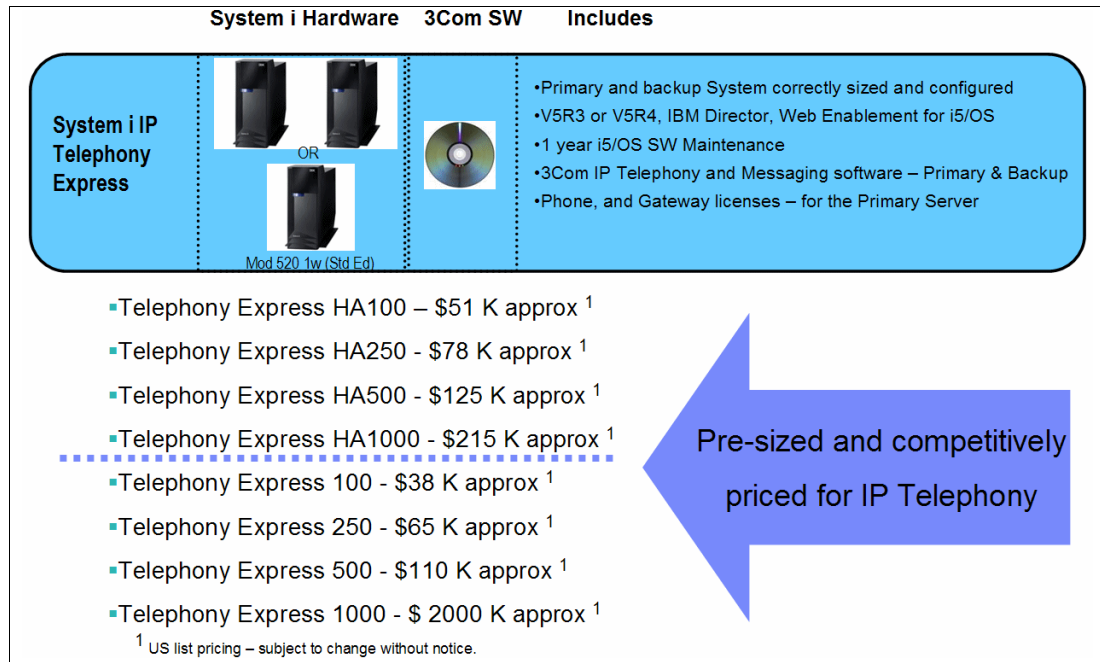


Figure E-4 IBM System i5 520 model Express packages as of December 2006

Note that the following products are included for no charge with an i5/OS V5R3 and V5R4, and you should be aware of these when planning a VOIP partitioned configuration on a System i model:

- ▶ IBM Director V5.10 Server for i5/OS, 5722DR1. Any i5/OS V5R3 or V5R4 order, if selected, includes i5/OS Server and Agent functions and agents for AIX 5L V5.3, Linux partitions, agents for Linux and Windows operating systems running on a System i5 integrated xSeries server, and the IBM Director Console for AIX, Linux, and Windows operating systems.

For more information about IBM Director refer to Chapter 17, “i5/OS and IBM Director” on page 455.

- ▶ Web Enablement for i5/OS. Any i5/OS V5R3 or V5R4 order, if selected, includes this software package, which includes:
 - 5722-WE2: This can include, if ordered, WebSphere Application Servers for i5/OS Express for:
 - Version 5.1.x
 - Version 6.0.x
 - Version 6.1.x
 - IBM Express Runtime Web Environments for i5/OS (5733-SO1). This product:
 - Installs, if not already installed, WebSphere Application Server - Express V6.0 for OS/400, 5733-W60. This includes fix pack level 6.0.2.9.
 - Installs, if not already installed, V5R4 IBM HTTP Server, 5722-DG1. This includes installing PTFs.
 - Configuring (using a wizard) a WebSphere Application Server profile (corresponds to an instance in pre-V6.n versions of WebSphere Application Server) and an associated HTTP server instance, generates the WebSphere Application Server plug-in, and starts both servers.

- Installs, if not already installed, V5R4 iSeries Access for Web, 5722-XH2. This includes installing PTFs.
- Installs iSeries Access for Web to run in the wizard. This is created at the WebSphere Application Server profile.
- Installs a 5250-based flight400 application in HATS, WebFaced, and WebSphere Services versions into the wizard. This is created at the WebSphere Application Server profile.

For further information about this support refer to Chapter 14, “i5/OS-based Web enablement enhancements” on page 337.

The following are short summary lists of the IBM System i5 hardware configurations available for each Telephony Express order feature number. See the IBM System i Telephony Web site for software and associated number of user license feature number requirements.

Single system Telephony Express configurations

For these:

- ▶ #7381 Telephony Express 100 Users hardware:
 - One System i model 520 1.9 GHz 1 Way (3800 CPW)
 - One #9613 Base 70 GB 15K rpm disk drive
 - One #9614 Base 1 GB DDR2 main storage
 - One #9540 Base DVD-ROM
 - One partition - #0140 Logical Partitioning Specify
 - One partition - #0142 Linux Partition Specify
 - One #7381 Telephony Express 100
- ▶ #7382 Telephony Express 250 Users hardware:
 - One System i model 520 1.9 GHz 1 Way (3800 CPW)
 - One #9613 Base 70 GB 15K rpm disk drive
 - One #9614 Base 1 GB DDR2 main storage
 - One #9540 Base DVD-ROM
 - One partition - #0140 Logical Partitioning Specify
 - One partition - #0142 Linux Partition Specify
 - One #7381 Telephony Express 250
- ▶ #7383 Telephony Express 500 Users hardware:
 - One System i model 520 1.9 GHz 1 Way (3800 CPW)
 - One #9613 Base 70 GB 15K rpm disk drive
 - Two #9614 Base 1 GB DDR2 main storage (2 GB total)
 - One #9540 Base DVD-ROM
 - One partition - #0140 Logical Partitioning Specify
 - One partition - #0142 Linux Partition Specify
 - One #7383 Telephony Express 500
- ▶ #7384 Telephony Express 1000 Users hardware
 - One System i model 520 1.9 GHz 1 Way (3800 - 7100 CPW)
 - One 8410 Base Processor Activation
 - Two #9613 Base 70 GB 15K rpm disk drives (140 GB total)
 - Two #9614 Base 1 GB DDR2 main storage (2 GB total)
 - One #9540 Base DVD-ROM
 - One partition - #0140 Logical Partitioning Specify
 - One partition - #0142 Linux Partition Specify
 - One #7384 Telephony Express 1000

High Availability primary, secondary system Telephony Express configurations

See Table E-1.

Table E-1 Telephony Express, High Availability hardware configurations

#0486 Telephony Express HA100	
<p>Primary system</p> <ul style="list-style-type: none"> ▶ One System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ One #9614 Base 1 GB DDR2 main storage ▶ One #9540 Base DVD-ROM ▶ One partition - #0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #0486 Telephony Express 100 	<p>Secondary system</p> <ul style="list-style-type: none"> ▶ 1 x System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ Two #9614 Base 1 GB DDR2 main storage (2 GB) ▶ One #9540 Base DVD-ROM ▶ One 0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #9734 Telephony Standard Edition
#0487 Telephony Express HA 250	
<p>Primary system</p> <ul style="list-style-type: none"> ▶ One System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ Two #9614 Base 1 GB DDR2 main storage (2 GB) ▶ One #9540 Base DVD-ROM ▶ One #0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #0487 Telephony Express 250 	<p>Secondary system</p> <ul style="list-style-type: none"> ▶ 1 x System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ Two #9614 Base 1 GB DDR2 main storage (2 GB) ▶ One #9540 Base DVD-ROM ▶ One 0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #9734 Telephony Standard Edition
#0488 Telephony Express HA 500	
<p>Primary system</p> <ul style="list-style-type: none"> ▶ One System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ Two #9614 Base 1 GB DDR2 main storage (2 GB) ▶ One #9540 Base DVD-ROM ▶ One #0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #0488 Telephony Express 500 	<p>Secondary system</p> <ul style="list-style-type: none"> ▶ 1 x System i model 520 1.9 GHz 1 Way (3800 CPW) ▶ One #9613 Base 70 GB 15K rpm disk drive ▶ Two #9614 Base 1 GB DDR2 main storage (2 GB) ▶ One #9540 Base DVD-ROM ▶ One 0140 Logical Partitioning Specify ▶ One #0142 Linux Partition Specify ▶ One #9734 Telephony Standard Edition
#0489 Telephony Express HA 1000	
<p>Primary system</p> <ul style="list-style-type: none"> ▶ One System i model 520 1.9 GHz 1 Way (3800–7100 CPW) ▶ One 8410 Base Processor Activation ▶ Two #9613 Base 70 GB 15K rpm disk drives (140 GB total) ▶ Two #9614 Base 1 GB DDR2 main storage (2GB total) ▶ One #9540 Base DVD-ROM ▶ One partition - #0140 Logical Partitioning Specify ▶ One partition - #0142 Linux Partition Specify ▶ One #0489 Telephony Express 1000 	<p>Secondary system</p> <ul style="list-style-type: none"> ▶ One System i model 520 1.9 GHz 1 Way (3800–7100 CPW) ▶ One 8410 Base Processor Activation ▶ Two #9613 Base 70 GB 15K rpm disk drives (140 GB total) ▶ Two #9614 Base 1 GB DDR2 main storage (2GB total) ▶ One #9540 Base DVD-ROM ▶ One partition - #0140 Logical Partitioning Specify ▶ One partition - #0142 Linux Partition Specify ▶ One #9734 Telephony Standard Edition

Further description of 3Com support on a System i5 configuration is beyond the scope of this book. The following are places to go for more information:

- ▶ IBM Redbook *IBM System i IP Telephony*, SG24-7382, available at:
<http://www.ibm.com/redbooks>
- ▶ 3Com Web site at:
<http://www.3com.com>
Current System i information at:
http://www.3com.com/index_jump/ibm_ipTel.html
- ▶ System i Telephony information at:
<http://www.ibm.com/systems/i/solutions/ipTelephony/>
- ▶ IBM System Workload Estimator (WLE) for sizing 3COM applications
A fast path link to the Workload Estimator IP Telephony solution sizer is:
<http://www.developer.ibm.com/graphics/estimator/HTML/IP3Com.html>

Figure E-5 depicts one of the early WLE windows during a sizing of the 3Com applications available under WLE.

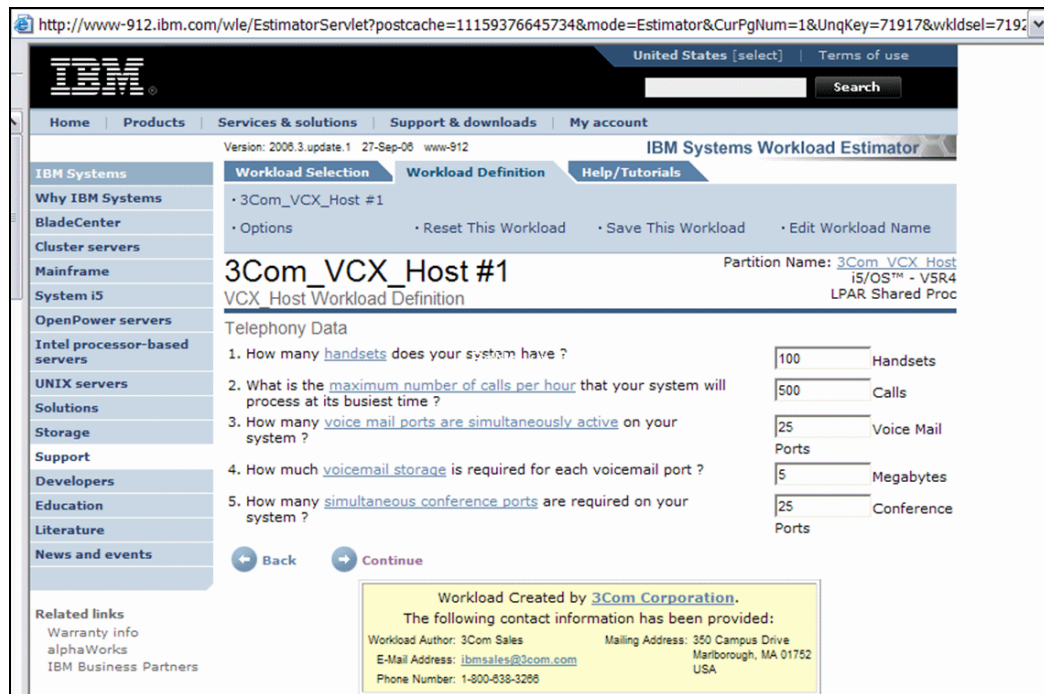


Figure E-5 WLE example: sizing a 3COM application

Note: You can also get the to the IP Telephony solution WLE sizing windows starting at the standard iSeries Performance Management Web site, selecting **sizing tools**. You get to the first Workload Estimator page by going to:

<http://www.ibm.com/eserver/series/perfmgmt/sizing.html>

Select **IBM Systems Workload Estimator**.

From there, add partitions and application workloads.

However, for sizing a 3Com application workload, we recommend starting with the fast path direct link URL listed before Figure E-5. This is the most efficient way to size the 3Com application suite of functions. This is because, depending on the 3Com options you select, you may end up with multiple Linux partitions to run the 3Com workload on the recommended IBM System i configuration. In most cases this set of 3Com application workloads and partitions is all you want to run on that configuration, and you are done.

If you want to run other applications in other partitions on the same system, use the direct link to size the 3Com workload first. Then use the standard entry to WLE to add and size additional workloads and partitions.



IBM-Zend PHP on i5/OS

This chapter give a brief update on PHP support on the IBM System i platform.

IBM-Zend PHP Web deployment environment

This appendix summarizes the Zend PHP Web deployment environment.

Because many IBM System i5 clients strive to make their businesses more innovative and look for an easy, cost-effective, and rapid way to deploy Web applications, adding the capability to run develop Zend PHP Web applications helps meet these expectations. IBM and Zend Technologies, Inc., provide an integrated and open Web application deployment environment using PHP on i5/OS. PHP is a powerful, open, yet easy-to-use Web application environment that has the support of a large community with more than 10,000 applications and components to share. PHP is a natural fit for i5/OS clients and developers, offering improved productivity and rapid deployment.

With PHP software on i5/OS, you can easily integrate PHP applications with existing investments in DB2 Universal Database (UDB) and RPG, COBOL, Java, and other business applications on i5/OS. With millions of PHP developers, you have access to a larger pool of skilled resources for Web enablement projects.

Zend offers selected PHP products for i5/OS, including:

- ▶ Zend Core for IBM i5/OS V1.5
- ▶ Zend Platform for IBM i5/OS V2.1.2
- ▶ Zend Studio for IBM i5/OS V5.2

Hardware and OS requirements

Zend Core for i5/OS will run on System i5 server models 520, 550, 570, and 595 and iSeries models 270, 800, 810, 820, 825, 830, 840, 870, 890, and SB2 with i5/OS V5R4. Additional requirements are:

- ▶ Portable App Solutions Environment - 5722-SS1 (Option 33)
- ▶ Qshell - 5722-SS1 (Option 30)
- ▶ System Openness Includes - 5722-SS1 (Option 13)
- ▶ Digital Certificate Manager - 5722-SS1 (Option 34)
- ▶ CCA Cryptographic Service Provider - 5722-SS1 (Option 35)
- ▶ IBM Portable Utilities For i5/OS - 5733-SC1 (*base)
- ▶ OpenSSH, OpenSST, Zlib - 5733-SC1 (Option 1)

For more information about the Zend PHP support running on i5/OS see IBM Redbook *PHP: Zend for i5/OS*, SG24-7327.



i5/OS V5R4 IPv6 restrictions

This chapter summarizes the status of i5/OS TCP/IP function and application use of IPv6.

IPv6 support restrictions through V5R4

There is no IPv6 support for the following:

- ▶ CL commands to configure IPv6
- ▶ IPSec/VPNs over IPv6
- ▶ Packet filtering, QOS, and NAT support
- ▶ Virtual Private Addresses with IPv6
- ▶ Local host table support for IPv6 host names
- ▶ Ability to communicate with DNS over an IPv6 network
- ▶ Ability to use iSeries Navigator over an IPv6 network

The limited number of IBM-provided applications that have been updated to support IPv6 include:

- ▶ V5R2: ping, traceroute, netstat, comm trace, xsockets
- ▶ V5R3: LDAP, Extended Dynamic Remote SQL Server

V5R2 or V5R3 based IPv6 configurations cannot be migrated to V5R4. You must reconfigure any IPv6 on V5R4 because of the significant enhancements made in V5R4.

Tunnels are no longer supported in V5R4 because:

- ▶ The requirement for tunnels has been reduced now that IPv6 and IPv4 can be used on the same Ethernet adapter.
- ▶ Network routers can be used if IPv6 packets need to be sent over an IPv4 network.

The V5R2 and V5R3 IPv6 Configuration Wizard in iSeries Navigator is not available on V5R4:

- ▶ With V5R4 there is a new interface to configure Stateless Address Auto configuration.
- ▶ With V5R4 there is a new wizard to create IPv6 interfaces.



IBM future planning statements

In this appendix we provide a February 2007 summary of planning statements made for IBM System i models and i5/OS-based software. The items included in this appendix are either simple *as is* statements or overviews of new or changed support announced during 2006 up through February 2007. The February 2007 announcement letter documents include:

- ▶ Withdrawal Announcement letter 907-021
- ▶ IBM i5/OS V5R3 and selected IBM System i programs and features Announcement letter 907-022

Also see the information in this appendix for references to additional information.

Remember that statements about IBM plans and statements of direction are made to assist customers in planning for continued or future use of system hardware configuration capabilities and software product capabilities. Plans and statements of direction can be made, changed, or withdrawn any time and without prior notice.

Planning statements

In this appendix we provide a summary of important hardware and software marketing and technical support statements. These cover new support of planned or planned removal of marketing as of February 2007.

For the latest information in this area, go to the Web site:

<http://www.ibm.com/servers/eserver/support/series/planning/migrationupgrade.html>

Or:

<http://www.ibm.com/systems/support/i/planning/upgrade/index.html>

Remember that IBM planning statements are made to provide you with insight into IBM plans for future support or removal of support. This information is subject to change or withdrawal without notice.

Marketing withdrawals - hardware

In this section we discuss hardware marketing withdrawals:

- ▶ Twinax controller planning statement withdrawn:
Twinax controller will *not* be withdrawn from marketing on June 1. #4746 controllers will continue to be available.
- ▶ The #0588/#5088 will be withdrawn from marketing July 1, 2007:
This enclosure is not ROHS compliant. For countries that do not require ROHS compliance, IBM still has a modest supply of #0588 for a couple more months. Order them using the withdrawn product process. IBM must have firm orders by February 2007.
Alternatives you can consider include:
 - #5790 I/O drawer (for 5xx models). This assumes that your configuration is not I/O drawer/tower constrained.
 - #5094/5294 I/O tower.
 - As of February 2007 option: #5096/5296.
- ▶ Withdrawals effective May 8, 2007 :
 - Integrated xSeries Adapter (1519-200)
 - VXA-2 and VXA-320 tape drives
 - 50 USB keyboards worldwide superseded by other keyboardsSee announcement letter 907-021 for more information.
- ▶ Withdrawal from marketing effective June 1, 2006:
 - New box shipments of 1.5/1.65 GHz POWER5 systems.
 - 5560/61/62/63 mirroring packages (#5095/0595+12 disk+2 #2757).
 - Replace many 520 upgrades to 1.5/1.65GHz model 520 with 1.9GHz.
 - 8xx unique features (for example, 800/810/825/870/890 memory features, 8xx CEC HSL adapters, 8xx CEC disk expansions).
On demand features such as processor activations are not withdrawn.
 - 870 5/8-w to 870 8/16-w and 890 16/24-w to 890 24/32-w upgrades.
 - #0588/5088 expansion unit (14 PCI slots, no disk slots).
 - #2849 10/100 Ethernet adapter (SNA capable adapter).

- ▶ IBM plans to withdraw from marketing the sale of new 35 GB 15k rpm disk drives effective May 1, 2007.

Suggested actions and considerations:

- Start moving to 70 GB or large disk capacities as makes sense for your configuration. Unless utilization is low, use the larger 757 MB write cache disk controller to enhance performance. If using RAID, use an auxiliary write cache IOA with the controller.
 - Consider assigning 17 GB or smaller disk drives still in your configuration to separate partitions or ASPs. Mixing large capacity drives with small capacity drives in the same ASP (for example, mixing 17 GB with 70 GB) can degrade performance.
 - Note that the older #5074/5079 I/O towers do not support 15k drives.
 - IBM hardware maintenance and operating system support of the 35 GB disk drives continues until further notice.
- ▶ Specific I/O-related hardware will not be supported on the next generation of POWER technology that follows POWER5 technology on System i models.

Watch for these announcements during 2007. For the latest information in this area, refer to the System i planning and upgrade Web site listed at the beginning of this appendix.

Announced during February 2007, the models 520/550/570/595 are the last to support the following hardware:

- #5074/5079 I/O towers
- Optical HSL
- 8 GB 10k rpm and 17 GB 10k rpm disk drives
- 11 older PCI cards including:
 - #2763, #2782, #4748, #4778 Disk Controllers
 - #2765, #2766 Fibre Channel Controllers
- Several older tape drives/media including
 - 9348 Tape Drive (½-inch reels)
 - 3570 and 3575 Tape Drives
 - 3490 Tape Drives when attached via #2749 HVD SCSI Tape Controller
 - 358x LTO-1 Tape Drives when attached via #2749
 - 4 GB, 16 GB, 25 GB QIC tape drives
 - All VXA tape drives

Watch for more announcements during 2007. For the latest information in this area, refer to the System i planning and upgrade Web site listed at the beginning of this appendix.

Software withdrawals

Effective January 4, 2008, IBM will withdraw from marketing i5/OS V5R3 and selected programs and features licensed under the IBM International Program License Agreement. For more information see announcement letter 907-022.

General software planning statements

General software planning statements follow:

- ▶ HATS Limited Edition, 5722-XH2, withdrawn from marketing effective May 1, 2006.
- ▶ i5/OS support of selected iSeries Models.

IBM plans for i5/OS V5R4 to be the final release supported on iSeries models 270, 820, 830, 840, SB2, and SB3. The next release is planned to be supported on models 520, 550, 570, 595, 800, 810, 825, 870, and 890.
- ▶ iSeries Access for Windows Data Transfer for Excel.

IBM plans for iSeries Access for Windows V5R4 to be the final release to support Excel 95 and Excel 97 with the Data Transfer add-in support. Only Excel 2000 and later will be supported for data transfers in the release after V5R4.

Suggested alternatives: Upgrade to Excel 2000 or newer.

▶ Apache Tomcat, January 2006.

IBM plans for i5/OS V5R4 to be the final release to ship Apache Tomcat (jsp/servlet) support in IBM HTTP Server, which ships with i5/OS.

Suggested alternatives:

- Use WebSphere Application Server Express, which is shipped with i5/OS.
- Use an open source version of Tomcat, which can be obtained from:
<http://www.jakarta.apache.org>

▶ Enhanced Netware Integration, January 2006.

IBM plans for i5/OS V5R4 to be the final release to support Enhanced Netware Integration (5722-SS1 Option 25).

▶ Tivoli Management Agent, January 2006.

IBM plans for i5/OS V5R4 to be the final release to include Tivoli Management Agent with i5/OS shipments.

Suggested alternatives: Obtain the Tivoli Management Agent for i5/OS through Tivoli channels.


▶ S/36 and S/38 environments.

IBM plans to continue supporting the S/36 and S/38 environments within i5/OS in the next release of i5/OS.

▶ RPG and COBOL System/36™ and System/38™ compatible compilers, July 2006.

IBM plans to continue to ship and support the following S/36 and S/38 compiler options as part of WebSphere Development Studio for iSeries in the release after V5R4. IBM is withdrawing the previous planning statement that said V5R4 will be the final release to include these compilers in the Software Maintenance or Software Subscription contract.

- 5722-WDS Option 32 - System/36 Compatible RPG II
- 5722-WDS Option 33 - System/38 Compatible RPG III
- 5722-WDS Option 42 - System/36 Compatible COBOL
- 5722-WDS Option 43 - System/38 Compatible COBOL



IBM System iSociety

This appendix summarizes what is available within the iSociety online community being developed collaboratively by members of the System i community, including individual clients, the COMMON user group, students, software providers, and publishers.

iSociety overview and facilities

iSociety is an online community being developed collaboratively by members of the System i community, including individual clients, the COMMON user group, students, software providers, and publishers. Its intent is to foster awareness of IBM System i capabilities and enable assistance among members. There are many capabilities running under i5/OS that many i5/OS customers are unaware of. This community helps these customers use i5/OS and solutions under i5/OS to their maximum advantage.

To register go to:

<http://www.isociety.org>

Once registered, you can examine all of the offerings, which include, for example:

- ▶ Business and social networking
 - Chat and collaboration
 - Job and skills board
 - Training
 - Peer-acclaim system
- ▶ System i information portal
 - Blogs, podcasts, and wikis
 - News and RSS feeds
- ▶ System i publications
- ▶ Discussion groups and forums

You can link to many other related Web sites, including:

- ▶ Truth site: This is designed to help iSociety reach out to IT professionals who are unfamiliar with System i and are seeking alternatives to complex, Windows-based systems. The Truth site enables System i professionals to share real experiences and stories that demonstrate the unique strengths of System i.
<http://ibm.com/iwantcontrol/truth>
- ▶ i want an i site: At this site you can find powerful tools that enable the entire System i community (business partners, ISVs, and clients)- to clearly and consistently communicate the unique business value of System i.
<http://www.ibm.com/systems/i/iwantani/>
- ▶ Marketing in a Box site: This site is designed to help IBM System i sales and marketing teams quickly develop and execute client and partner events, breakfast briefings, and other demand generation activities.
<http://www.ibm.com/systems/i/marketinginabox/>

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these redbook and redpaper publications, see “How to get IBM Redbooks” on page 553. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *Implementing Integrated Windows Server through iSCSI to System i5*, SG24-7230
- ▶ *LPAR Simplification Tools Handbook*, SG24-7231
- ▶ *Deploying IBM Workplace Services Express on the IBM eServer iSeries Server*, SG24-6664
- ▶ *Sizing IBM i5/OS Work on IBM System i5 Partitions*, SG24-6656
- ▶ *A Systems Management Guide to Performance Management for i5 and p5 Systems*, SG24-7122
- ▶ *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155
- ▶ *IBM System i5 Handbook IBM i5/OS Version 5 Release 4 January 2006*, SG24-7486
- ▶ *PHP: Zend for i5/OS*, SG24-7327
- ▶ *IBM eServer iSeries Systems Management Handbook*, REDP-4070
- ▶ *Planning for IBM eServer i5 Data Protection with Auxiliary Write Cache Solutions*, REDP-4003
- ▶ *PCI and PCI-X Placement Rules for IBM System i5, eServer i5, and iSeries servers with i5/OS V5R4 and V5R3*, REDP-4011
- ▶ *IBM i5/OS Intrusion Detection System*, REDP-4226
- ▶ A direct link to ITSO Technotes is:
<http://www.redbooks.ibm.com/redbooks.nsf/tips/>
Search using *i5/OS*.

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ The IBM Systems - Servers
<http://www.ibm.com/servers/index.html>
- ▶ The IBM System i
<http://www.ibm.com/systems/i/>
- ▶ The IBM Systems Information Center
<http://publib.boulder.ibm.com/eserver>

- ▶ The POWER5-based IBM System Hardware Information Center
<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>
 This contains much information for System i5 and p5 models, such as has how to install and activate things like the POWER5 system, making your choice of console setup, setting up and administering HMC-based LPAR, setting up and using Capacity on Demand offerings, and more.
- ▶ The IBM iSeries (System i) Information Center
<http://publib.boulder.ibm.com/iseriess/>
- ▶ IBM System i planning, including upgrades information
<http://www.ibm.com/systems/support/i/planning/migrationupgrade.html>
- ▶ IBM Director details at the IBM Systems Software Information Center
<http://publib.boulder.ibm.com/eserver/swic.html>
- ▶ Direct link to a list of the PDFs for IBM Director at the IBM Support
<http://www.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?brandind=5000016&indocid=MIGR-61788>
- ▶ i5/OS-based documentation for WebSphere Application Server V6.0
<http://www.ibm.com/servers/eserver/iseriess/software/websphere/wsappserver/docs/docws60.html>
- ▶ i5/OS-based documentation for WebSphere Application Server V6.1
<http://www.ibm.com/servers/eserver/iseriess/software/websphere/wsappserver/docs/docws61.html>
- ▶ The latest WebSphere Application Server for i5/OS PTFs
<http://www.ibm.com/servers/eserver/iseriess/software/websphere/wsappserver/services/service.html>
- ▶ IBM Electronic Software Delivery
<http://www.ibm.com/servers/eserver/ess>
- ▶ IBM Support Assistant
<http://www.ibm.com/software/support/isa/>
- ▶ Workload Estimator
<http://www.ibm.com/systems/support/tools/estimator>
- ▶ All IBM System i and System p white papers
<http://www-03.ibm.com/servers/enable/site/education/ibo/view.html?wp>
- ▶ Information about power control over the Ethernet IP network
http://www.poweroverethernet.com/articles.php?article_id=52
- ▶ ITSO Technical Overview set of presentations (V5R2, V5R3, V5R4) under the Technical Library link at the technical support Web site shown below. This book contains more up-to-date information about V5R4 capabilities, as the Web content was last updated in March, 2006.
<http://www.ibm.com/servers/eserver/support/iseriess/index.html>
 Select **Technical Library** → **System i Technical Overviews**.

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

#0308 54
#0454 56
#0455 56
#5706/5707 247
#5737, #5776, #0648 cache sizes 30
#5738, #5777, #5582, #5583 cache sizes 30
#5739, #5778, #5781, #5782, #5799, 5800 cache sizes 30
%ADDRESS or %ADDR 277
%HANDLER 281
%XML (RPG) 281
(BRMS) 168
(DSPPFRDTA 438
(SLIC 41
(SQL_ATTR_MAX_ROWS 91
*FLDBDY 96
*JOBEND 70
*JOBLOGSVR 70
*LINK Omit (BRMS) 158
*MODULE 422
*NULLIND 279
*PGM 422
*PND for job log output 70
*SQLPKG 422
*SRVPGM 422
*SYSTEM certificate store 374
*SYSVAL for job log output 70
.NET 86
.NET Data Provider 87
.NET Framework 87
.NET plug-ins 90
.NET provider 86
.NET support 243
.Net support 243
.sysplan 260
.sysplan file 258, 260
"HATS-ed" version of Flight400 311

Numerics

0269 55
0647 26
0647 controller 25
0647 IOP not required PCI-X controller 25
10 GB Ethernet LAN IOA 52
128 byte column name support 289
16xx printers 400
2 MB SQL statement 86
2 MB SQL statement support 289
2 port Ethernet LAN 247
2002/95/EC Directive 20
2007 xvii
2757 cache sizes 30
2765 30

2766 30, 49
2780 cache sizes 30
2787 30, 49
2787 performance 31
2793 32
2794 32
2844 IOP 49
2847 IOP 49
2D bar code types
 Data Matrix 391
 MaxiCode 391
2D barcodes 391
2EE Deployment, 318
32 bit or 64 bit (Java) 286
32-bit versus 64-bit architecture (JVM) 284
32K keys 86
32R1923 126
3Com 531
3Com Convergence Application Suite for IBM System i 532
3Com IP Telephony Suite for IBM System i 532
3Com IP Telephony Suite for System i 532
3Com offerings on System i 532
3Com Web site 529
5/OS integration of IBM System x-based servers 117
520 and 550 model memory technology 500
520 model Express packages 535
520 model Telephony Express packages 535
520 Telephony Express packages 532
5372-IS5 55
5553 feature code 247
5639-3CM 532
570 model memory technology 501
570 redundant CPU power regulator 39
5700 32
5701 32
5702 23
5704 30
5705 23
5709 RAID Enabler Card (controller) 22
5712 23
5715 23
5722-AC3 141, 325
5722-CR1 (future plans) 141
5722-DG1 340
5722-IP1 389
5722-LSV 120
5722-PT1 428
5722-SS1 option 34 325
5722-WE2 536
5722-WSV 120
5722-XE1 240
5722-XH2 240
5722-XL1 240
5722-XP1 240

- 5722-XW1 240
- 5726 22
- 5727 22
- 5728 22
- 5733-ID1 390
- 5733-SO1 536
- 5733-SO1 requirements 342
- 5733-WE2 339
- 5733-WSE 486
- 5736 24
- 5736 IOP-based PCI-X disk, tape attachments 24
- 5737 26
- 5760 31, 49
- 5761 31
- 5766 IOP-based PCI-X tape controller attachment 24
- 5771-3CM 532
- 5775 24
- 5775 and 0647 controllers
 - tape devices supported 25
- 5776 26
- 5783 - Host Bus Adapter (HBA) 124
- 5784 - Host Bus Adapter (HBA) 124
- 595 DDR1 memory 505
- 595 DDR2 memory 506
- 595 model memory card placement performance considerations 508
- 595 model memory technology 504
- 64-bit ODBC driver 244
- 6574 4-Disk Slot Expansion Base Controller 23
- 6592 4-Disk Slot Expansion Base Controller 23
- 6800 32
- 6801 32
- 6803 32
- 6804 32
- 7316-TF2 249
- 7463 (processor feature) 360
- 7722-XH2 340
- 8677xxx 126
- 8852xxx 126
- 9493 32
- 9494 32
- 9510 22
- 979 32
- 9793 32

A

- A microcode level 41
- Accelerator feature 10
- Accelerator on System i5-520 408
- access plan 95
- access ports 127
- Active Server Page (ASP) - Microsoft 272
- ActivitySession Service 317
- adapter placement rules 248
- Add Environment Variable 286
- ADDENVVAR 286
- Additional Settings (advisor) 370
- Additional Settings tab (advisor) 368
- Additional web application server settings 366
- Address Resolution Protocol 229

- Administrative Console 319
- administrative console port number 329
- ADO.NET 90
- Advanced POWER Virtualization 256
- Advanced Power Virtualization 255
- Advanced System Management Interface 249, 496
 - ASMI 493
- advanced system management interface 251
- Advanced System Management Interface (ASMI) 10
- Advise window for Threads to process 365
- Advisor 107
- AFP - Advanced Function Presentation™ 304
- Agents
 - Level 0 458
- Agents (IBM Director) 458
- AH 235
- AIX 33, 174
- AIX 5L partition 255, 260
- AIX 5L partitions 445
- Alias Name (IP address) 228
- All IBM System i and System p whitepapers URL 552
- Allow reorganization to be suspended 413
- Alter Procedure 95
- ALWCANCEL(*YES) 413
- ALWCPYDTA(*NO) 409
- American National Standards Institute 86
- analyze monitor enhancement 110
- ANSI 86
- ANSI/ISO SQL Standard and ISO formats 290
- antivirus scanning 144
- apacity on Demand 249
- API 382, 386
- APIs 379, 382
- application developer roadmap 270
- application development enhancements (WebSphere) 331
- application development languages and tools 269
- application flexibility and portability 84
- Application Profiling 317
- Application Response Measurement 289, 352
- Application server resources 366
- APPN Connection Network support
 - Connection Network support 224
- APV 255
- Archive (BRMS) 168
- Archive options (BRMS) 155
- ARM 289, 352
- ARP 229
- ASMI 249, 251, 496
 - displaying system firmware level 496
- ASP performance information 433
- AST enhancements 331
- Asynchronous Bean 317
- Aurema, Inc. Web site 474
- authority setup (stored procedures) 95
- authorized machine interface instruction 137
- Autodesk Inventor 399
- Autodesk Streamline 399
- automatic creation of index 98
- auto-negotiation 32

ava 2 Enterprise Edition 305
AX-RPC - APIs for representing WSDL-based services as
RPCs in Java 316

B

B microcode level 41
BACKMGN parameter 396
Backup and Recovery Media Services 153, 379
Backup and Recovery Media Services (BRMS) 147
Backup enhancements (BRMS) 159
bar code exception reporting 392
Barcode/400 399
Baseboard Management Controller (BMC) 129
best practice 111
BidiConversionProperties 289
Blade servers 126
BladeCenter 126
BladeCenter (iSCSI initiator) installation requirement 126
BladeCenter chassis 126
BladeCenter chassis IBM Director 464
BladeCenter components for System i5 integration 126
BladeCenter management modules 466
BladeCenter server performance 425
BladeCenter support 121
blades 464
Blanks are allowed in hex literals (SQL) 302
BM Infoprint 16xx Series Printer 395
BMC® PATROL® for iSeries - Predict 441
BRMS 153, 379
 *HSM (Hierarchal Storage manager) option 154
 *NETWORK option 154
 BASE option 154
BRMS - no auto create 173
BRMS Archive Policy Report 155
BRMS interfaces 158
BRMS iSeries Navigator 155
BRMS plug-in 159
BRMS product overview 154
Bulk power processor 42
Bypass 5250 session sign on 210

C

C6004 508 reference code) 254
cache sizes (disk controllers) 30
Call Level Interface 91
call stack display 74
call stack with Java entries 75
CallStackEntry 289
CALLSUBR command 275
Capacity Backup edition example 510
Capacity BackUp Edition Offerings 511
Capacity Backup editions 509
Cartridge Type under BRMS 164
Catalog security 186
CBU Standard Edition offerings 512
CCSID UTF8 237
ceate the virtual tape image catalog 175
cell 318
certificate store 374

CGI 272
CGU - Character Generator Utility 304
Change Journal (CHGJRN) command 96
Change Printer File 386
Change Query Attributes 412
Change Spooled File Attributes 386
Change TCP/IP IPv4 Interface 229
Changes in Web Services with V6.0 316
Channel Framework 321
CHGIPLA spool parameters 398
CHGPRTF 386
CHGQRYA 102
CHGSPLFA 386
CIDR 226
CIFS 238
CIM 32, 463
CIMOM 463
CL 275
Classless Inter-Domain Routing 226
cleaning up your existing spool files 387
CLI 91
CLI enhancements 91
Client Encryption (5722-CE3) 141
CLR 87
COD 249
CODE and VisualAge RPG V6.0 305
CODE128 bar codes 391
coexistence - IXS, IXA, iSCSI 122
Coginord 378
collect inventory 469
Collection Services 370, 447
Column-wise blocked insert binding 91
COMMIT 92
Common Cryptographic Architecture Cryptographic Ser-
vice Provider (CCA CSP) 141
Common Gateway Interface 272
Common Information Model 463
Common Information Model Object Manager 463
Common Information Module 463
Common Internet File System 238
Common Language Runtime 87
communications 221
compare monitors 111
Compatibility differences (BRMS) 159
Component Report changes 435
Concurrency (Java) 288
connection objects 92
Connection Security object 130
Console for Express Runtime 335
console not found 254
consoles 245
constraints 98
Containers support (BRMS) 156
Content Manage 114
Control group attributes (BRMS) 159
Control Language 275
Control specifications (RPG) 278
Coordinator 321
copper connections (iSCSI) 127
Copy Spooled File 380

- Core Flagger 85
- Core Group 321
- Core level of the SQL 2003 Standard 85
- CPU power redundancy 39
- CPU Used and CPU Available field
 - CONFOT 437
- CPU utilization for all processors 434
- CPW ratings 404
- CPYSPLF 380
- CQE 409
- create a virtual tape image 177
- Create Image Catalog 184
- Create Printer File 386
- create spooled job logs when needed 388
- cross-reference files 96
- CRTDEVNWSH command 132
- CRTNWSCFG TYPE(*CNNSEC) 132
- CRTNWSCFG TYPE(*RMTSYS) command) 132
- CRTNWSCFG TYPE(*SRVPRC)command 132
- CRTPRTF 386
- Cryptographic Access Provider 325
- Cryptographic Services APIs 141
- cryptographic support 141
- Cryptographic Support for AS/400 (5722-CR1) 141
- CSP 56
- cumulative PTF level of C6045530 8
- CURRENT DEGREE 100
- Current SQL for a job 113
- current user 77
- current user for a job 76
- current user profile 76
- Cursor Sensitivity 91
- Custom Data Protection Specify 55
- Custom Package (IBM Director) 476
- Customer Specified Placement 56
- Customizable string processing 88
- Cycle clause (SQL) 297

D

- daemon process 272
- DASD protection schemes diagram with #5738, #5777, #5739, #5778 IOAs 44
- data area journaling 96
- Data Description Specifications 272
- data queue journaling 96
- Data Replication Service 321
- Data replication service 321
- Data Retention products 165
- Data Source 357
- data source 369
- Data transfer 243
- database and availability 96
- database and On Demand 96
- Database application development enhancements 86
- database connection 214
- database cross reference files 414
- Database function enhancements (iSeries Access for Web) 211
- Database host server trace support 289
- database limits and sizes 113

- DataBase Management Systems 85
- database performance monitor 109
- database performance update 408
- Datamax 399
- Datasource 212
- Date Reset Days Used Count 105
- Daylight Savings Time 78
- Daylight Savings Time (HMC) 493
- Days Used Count 105
- DB Group PTF 94
- DB Group PTF, SF99503 109
- DB2 - additional information 116
- DB2 Add-Ins for Visual Studio 90
- DB2 Connect Unlimited Edition for iSeries 90, 93
- DB2 Connect Unlimited edition website 91
- DB2 Content Manager Standard Edition 115
- DB2 Expression Evaluator 95
- DB2 family compatibility 86
- DB2 for i5/OS 83
- DB2 Multisystem 409
- DB2 Multisystem feature 98
- DB2 SMP parallelism 100
- DB2 UDB C code generation engine 95
- DB2 Universal Database (DB2 UDB) for iSeries 83
- DB2 Universal Database Express Edition 334
- DB2 Universal Database for i5/OS 83
- DB2 Universal JDBC driver 93
- DBC Resources tab (advisor) 368
- DBMS 85
- DCMs 441
- DDR1 memory 15
- DDR1, DDR2 memory and plugging rules 500
- DDR2 12, 15
- DDR2 memory 10
- DDS 272
- DDS printer file enhancements 395
- Dedicated Service Tools 254
- DefaultCoreGroup 321
- Delete Expired Spooled files 386
- Denial of service attacks 139
- Dependent catalog 181
- dependent catalog definition 183
- dependent catalog overview 181
- Development Studio subscription 304
- devices/media management enhancements (BRMS) 162
- DFTRDBCOL 95
- DFU - Data File Utility 304
- Digital Certificate Manager 325
- DIMMs 16
- Direct Execution (Java) 286
- disk hardware performance 417
- Disk Magic 442, 444
- disk sizing algorithms - Workload Estimator 443
- DISKPART command line utility 121
- Display Call Stack 74
- Display Licensed Program 412
- Display Performance Data 438
- Display Performance Data IP address, remote port number 430

- Display System Value (DSPSYSVAL) QPRCFEAT 360
- Distributed Management Task Force 463
- Distributed Partitioned Views 90
- DistributedMap 317
- DLTEXPSPLF 386
- D-mode IPL 254
- DMTF 463
- Domino for iSeries reference documentation 490
- Domino releases on i5/OS 484
- Domino Server Statistics 437
- Double Data Rate Dual Channel (DDR2) 15
- DPVs 90
- DR450 165
- DR550 165
- Driver class 212
- DS6000 48
- DS8000 48–49
- DSPLICPGM 412
- DSPOPT 380
- DSPSAVF 380
- DSPTAP 380
- DST 254, 493
- Dual Core Modules 441
- Dual mode IOA 19, 52
- dual mode IOA 8, 19–22, 54, 417
- duplicate media 165
- DUPMEDBRM 165
- DUPMEDBRM (BRMS) 165
- DUPTAP 181
- DVD ROM 22
- DVD-RAM 22
- dynamic memory changes 268
- Dynamic Query 317
- DYNDFTCOL 95

E

- ebSphere Enterprise Workload Manager 353
- Eclipse 306, 332
- Eclipse Update Manager 331
- e-configurator 26
- edition software details 510
- editions 9, 12–13
- EDRS 244
- EDTRBDAP 107
- Education Assistant 527
- Ended Sequence Number 235
- EGL COBOL Generation 305
- EIM 135
- EJB 2.1 332
- EJB QL 315
- Electronic Service Agent 473
- Electronic Software Delivery Web site 552
- E-mail policy (BRMS) 156
- Embedded LAN ports support 246
- embedded SQL rules (RPG) 282
- Encapsulating Security Payload 235
- encoded vector indexes 96
- End Job 70
- ENDJOB 70
- ENDPAGE keyword 397

- Enhanced Hardware Storage Protection 136
- Enterprise Edition (J2EE) 314
- Enterprise Enablement (EE) feature 510
- Enterprise Extender
 - Connection Network support 224
 - SNA based remote workstation support 224
- Enterprise Extender restrictions 224
- Enterprise Extenders 18, 221
- Enterprise Identity Mapping 135
- Enterprise Identity Mapping (EIM) 142
- Enterprise Java Beans (EJB) 305
- Enterprise Java Beans (EJB) 2.1 315
- Enterprise JavaBeans Query Language 315
- Enterprise Workload Manager 91, 256
- ESN 235
- ESP 235
- Ethernet adapters 32
- Ethernet LANs and SNA 33
- Ethernet switch requirements (iSCSI) 127
- Ethernet UTP 33
- EVAL-CORR (RPG) 279
- event action 466
- Event Action Plans 466
- event filter 466
- EVI 96
- EWLM 91, 353
- eWLM 256
- eWLM Correlator support 289
- Existing adapters without input/output processor 32
- Exit program 160
- EXP24 disk enclosures 28
- EXP24 example attachment 29
- export or import the performance profile (Web Performance Advisor) 359
- Express Runtime 334, 350
- Express Runtime Developer 335
- Express Runtime middleware 334
- Express Runtime V2.1.1 350
- Express Runtime Web Environments for i5/OS 311
- Extend iSeries data to mobile devices
 - DB2 Mobility on Demand 90
- Extended Architecture 91
- Extended Dynamic Remote SQL 244
- Extended Dynamic Remote SQL Server 544
- Extended JTA Support 317
- Extensible Markup Language 273
- eXtensible Markup Language 272
- Extension Task functions. 471
- Extract Server Data 217
- Extract server data 217

F

- Faces Client Tutorial 323
- Failover using a preferred interface list (IP) 230
- fastest disk controllers (adapters) 419
- Federal Express 391
- federate repositories 330
- Federated Data access 90
- federated database 90
- fiber optic connections (iSCSI) 127

- fibres channel adapters 30
- Field Boundary 96
- File function (iSeries Access for Web) 217
- File Serving workload (WLE) 443
- File Serving workload for AIX and Linux 442
- File transfer package (IBM Director) 476
- Fine-grained Application Update 318
- FIPS enablement 488
- Firefox 1.5 343
- firewalls 344
- Firmware 41
- Firmware Level 496
- firmware level 496
- firmware service pack levels 493
- Firmware Service Processor 496
- FirstSteps 345
- FirstSteps Web page 345
- Fix Central 79
- flash memory 42
- Flashcopy 49
- FLASHCOPY support 161
- FLASHLEVEL 496
- Flatpanel Monitor 127
- Flexibility 84
- flight400 application 340
- Force full backup days (BRMS) 160
- Four Gbps fibres channel adapters 30
- Free Format embedded SQL in ILE RPG 84
- free format for SQL (RPG) 282
- free-format RPG and SQL 84
- FRONTMGN parameter 396
- FSP 496
- FSP-HMC0 496
- FSP-HMC1 496
- Full select in a subquery 291
- functions (SQL) 86

G

- gaining control over your output queue 387
- Garbage Collection 422
- general connectivity 221
- general LPAR updates 264
- GENERATE_UNIQUE (SQL) 301
- Global Copy 51
- Governing Service 321
- graph history
 - Print 428
 - save 428
- Greenhouse by WebSphere application 323
- Group 108
- group PTF 119
- group registry definition (EIM) 142
- Guided Activity Assistant 524

H

- H.323 telephony protocol 530
- HA and CBU attribute comparisons 512
- HA Telephony using Two Systems example 535
- HA/DR ISV enhanced CBU Companion Offers 513

- handles limit 91
- Hardware (POWER5 or later technologies) Information Center xvii
- hardware cryptographic device support 330
- Hardware Management Console 36, 40, 258, 362, 491
- Hardware Storage Protection 136
- Hardware Storage Protection object conversion 423
- HATS 270, 307
- HATS information center 311
- HATS LE 208
- HATS Limited Edition 240
- HATS Toolkit 306
- HBA 35
- Health Center 113
- Health Check example 114
- heavy journal traffic 97
- high availability in an iSCSI, AIX 5L environment 119
- High Availability primary, secondary system Telephony Express configurations 538
- High Voltage Differential 25
- Hits/Second 436
- HMC 10, 36, 40, 258, 362, 459, 492
 - displaying system firmware level 495
 - record HMC configuration information and remote command status 492
 - release level for system plan input 258
 - saving critical console information 492
 - saving the managed system's profile data 492
 - saving the upgrade data 492
 - viewing the current HMC machine code level 492
- HMC FSP IP address 496
- HMC HTTP server 496
- HMC level 492, 494
- HMC machine code upgrade 492
- HMC Manager Tools 473
- HMC models 249
- HMC Version PTFs 494
- HMC Web site 492
- Host Access Transformation Server 307
- Host Access Transformation Server Limited Edition 208
- Host Access Transformation Services 270
- Host Access Transformation Services (HATS) 341
- Host Access Transformation Services Toolkit 305
- Host Bus Adapter (HBA) 122
- Host Bus Adapter (HBA) for System i5 124
- Host on Demand 270
- Host Print Transform 391–392
- hot spare 122
- HP-GL 379
- HPT 392
- HS20 126
- HS21 126
- HSP 136
- HTML 115, 271
- HTTP Server 340
- HTTP server 345, 353, 356
- HTTP Server 6.0.2 335
- HTTP Server Activity 436
- HTTP server enhancements 237
- HTTP support 221

HVD 25
Hypervisor 41

I

i5/OS console 128
i5/OS integration of IBM System x servers 117
i5/OS partition 260
i5/OS Tasks 351
i5/OS V5R4 and XA support enhancements 93
i5/OS-based documentation for WebSphere Application Server V6.0 552
i5/OS-based documentation for WebSphere Application Server V6.1 552
IASP 172
iASP considerations 383
iASPs and spool files 385
IBM Configurator for e-business 53
IBM DB2 UDB for iSeries .NET Provider Technical Reference 87
IBM Developer Kit for Java 286
IBM Developer Kit for Java DJBC driver 369
IBM Direct 305
IBM Director
 BladeCenter Management 472
 Groups 462
 Independent Hardware Vendor 460
 Independent Software Vendor 460
 Managed objects 462
 starting and stopping 463
IBM Director
 Agent Level 0 458
 Agent Level 1 458
 Agent Level 2 458
 all events filter 466
 Application Workload Manager 474
 base extensions 472
 base functions 471
 classes 459
 components 457
 Console 458
 create package for delayed delivery 476
 critical events 466
 define a new scheduled job 477
 discovery 464
 disk usage 471
 duplication event 467
 Event Action Plan 466
 Event Action Plan Builder wizard 467
 exclusion events 467
 execution history 477
 extensions Website 471
 for iSCSI 458
 Groups 456
 Inventory 465
 Job Scheduler 468, 477
 job scheduler alerts 469
 logging options 466
 managed objects 456
 memory usage 471
 network topology example 459

Premium Edition 474
QCPMGTSVR job 463
Rack Manager 472
Real Time Diagnostics 473
remote deployment manager 474
Resource Monitors task 471
SNMP 464
 supported agent software 456
 system availability 473
 System or Group thresholds 471
 tasks 456
 threshold event 467
 UpdateXpress 473
 z/VM Center 473
IBM Director 5.10 456
IBM Director 5.10 Console customization 462
IBM Director and i5/OS Collection Services 479
IBM Director and LDAP server 480
IBM Director and Tivoli summary 481
IBM Director Console 448
IBM Director details at the IBM Systems Software Information Center 552
IBM Director File Transfer 475
IBM Director for i5/OS Server 125
IBM Director for i5/OS Version 5.20 479
IBM Director Multiplatform 456
IBM Director overview 456
IBM Director Server 125
IBM Director topology 457
IBM Director V5.20 installation process 479
IBM Education Assistant 516, 527
IBM Express Runtime Web Environments for i5/OS 337, 536
IBM Express Runtime Web Environments for i5/OS with your applications 350
IBM i5/OS license entitlements 511
IBM Infoprint Color 150xx Series printer 395
IBM Infoprint Color 160xx Series Printer 395
IBM Portlet Application for Microsoft Exchange 2003 487
IBM Rational® Web Developer 335
IBM Support Assistant 3, 516, 552
IBM System iSociety 549
IBM System Planning Tool 258
IBM System Storage™ TS3400 Tape Library 38
IBM Systems Workload Estimator 441
IBM Technology for Java Virtual Machine
 32-bit JVM 422
IBM technology for Java Virtual Machine 283
IBM Toolbox for Java 288
IBM Web Enablement for i5/OS 339
IBM WebFacing 270
IBM Workplace Designer 489
IBM.Data.DB2 provider 90
IBM.Data.DB2.iSeries 87
IBM.Data.DB2.iSeries provider 90
IBM-Zend PHP on i5/OS 541
IDS 334
IETF 225
IFSFileReader 289
IFSFileWriter 289

- IFSSystemView 289
- IGP 399
- IKE packet 234
- IL 87
- image catalog 173
- Import Query 214
- Important 570 model memory pluggng rules 503
- IN predicates (SQL) 108
- Independent Auxiliary Storage Pool 172
- Independent Software Vendor (ISV) workloads (performance Web sites) 404
- Index Advisor 107–108
- Index Evaluator 104
- Index Evaluator function 99
- index management 102
- Index Rebuilds 107
- Industry and Technical Advisory Councils, 275
- Infoprint 15xx and 16xx 400
- Infoprint 6700 398–399
- Infoprint Designer for iSeries 390
- Infoprint Server for iSeries 389
- Information Center 119
- information center - hardware 10
- Information Technology (IT) 2
- Informix Dynamic Server Express 334
- Initiative for Innovation program 274
- Installation Factory 332
- instance (WebSphere Application Server) 318
- Instead of Triggers 103
- integrated application server 338
- integrated BladeCenter models supported 124
- Integrated Development Environment (IDE) 306
- Integrated File System 172, 212
- integrated file system (IFS) 144
- Integrated Server Administration 133
- integrated server performance
 - see Performance Capabilities Reference manual 122
- Integrated Server Support 120
- Integrated Server Support (option 29 of i5/OS V5R4) 121
- integrated Small Computer System Interface 458
- integrated System x models supported 124
- integrated xSeries models supported 124
- integrated xSeries server performance 424
- integrated xSeries systems 472
- Integrating DB2 Universal Database for iSeries with Microsoft ADO .NET, SG24-6440, 90
- Integrating PM for System i5 and the IBM System Workload Estimator 451
- Integration for Linux, 5722-LSV 120
- intellimagic 442
- IntelliSense® support 88
- Intermec 399
- International Standards Organization 85
- Internationalization Service 317
- Internet Control Message Protocol (ICMP) 139
- Internet Engineering Task Force 225
- Internet Protocol 221
- Internet protocol (IP) fragment 139
- internet Small Computer System Interface controller 23
- interpreted mode (Java) 286
- Interval Planning LPAR simplification 260
- intrusion detection 139
- Intrusion Detection System redpaper 141
- IOA-level mirroring 54
- IOP 12
- IOP controlled IOA 19
- IOP less only IOA 19
- IOP-less adapter 25
- IOP-less strategy 18
- IOTs 302
- IP address and port number - Display Performance Data DSPPFRDTA 438
- IP Conferencing 532
- IP Messaging 532
- IP Presence 532
- IP Security Authentication Header 235
- IP Telephony 532
- IPL 96
- IPL clearing of spooled files 398
- IPSec 234, 544
- IPv6 and Asia, Pacific 227
- IPv6 details 226
- IPv6 support 225
- IPv6 support restrictions 544
- ISA 516
 - installation and setup 517
 - product information 521
 - Support 525
 - Tools 523
- iSCSI 118, 458
 - BladeCenter components 126
 - BladeCenter installation requirements 126
 - Ethernet switch minimum installation requirements 127
 - i5/OS V5R4 requirements 125
 - IBM Director 125
 - iSeries Navigator 132
 - Service processor configuration 132
 - Virtual Ethernet objects 130
 - xSeries installation requirements 126
 - xSeries or BladeCenter installation 127
- iSCSI controller 23
- iSCSI HBA 122
- iSCSI Host Bus Adapter 35
- iSCSI initiator 124
- iSCSI installation diagram 129
- iSCSI installation requirements 124
- iSCSI introduction 123
- iSCSI performance considerations 119, 122, 124
- iSCSI support 35
- iSCSI target 124
- iSeries 90, 338
- iSeries (System i) Information Center xvii
- iSeries .NET Provider 86
- iSeries Access and Virtual Basic 242
- iSeries Access family 240
- iSeries Access family of products 221
- iSeries Access for Linux 243
- iSeries Access for Web 189, 311, 338, 340, 349
 - 5250 emulation new functions 210

- bypass sign-on 210
- customize function 211
- Database function 211
- Find Record 214
- PDF output setting preference 215
- portlet version enhancements 215
- printer output filtering 215
- servlet version enhancements 209
- single sign on 210
- single sign-on 210
- supported WebSphere Application Server versions 207
- WebSphere data sources 212
- iSeries Access for Windows 87, 240
- iSeries developer community 90
- iSeries Integration for Windows Server (LPP 5722-WSV) 121
- iSeries Navigator 132
- iSeries Navigator and EIM 142
- iSeries Navigator DB2 performance analysis 101
- iSeries Navigator on the Web 338
- iSeries Navigator System Monitor 427
- iSeries Navigator Tasks on the Web 189, 356
- iSeries support Web site for integrated xSeries systems 472
- iSeriesNetServer 290
- ISO 85
- ISO timestamp format 290
- iSociety 3, 549–550
- ITSO Technical Overview set of presentations (V5R2, V5R3, V5R4) 552

J

- J2CA 315
- J2EE 305
- J2EE 1.4 332
- J2EE 1.4 overview 314
- J2EE Connection Architecture 315
- J2EE Deployment 318
- J2EE Management, 318
- JACC 322
- JACL scripts 323
- Java 272
- Java 2 Platform 314
- Java Adventure Builder 323
- Java API libraries 288
- Java Authorization Contract with Containers 322
- Java Cryptography Extension 287
- Java Database Connectivity 289
- Java Development Kit (JDK) level 314
- Java Management Extensions 288, 318
- Java Message Service 315
- Java performance 421
- Java Platform Debugger Architecture (JPDA) 287
- Java Platform Profiling Architecture (JVMTI) 288
- Java Platform, Enterprise Edition (Java EE) 284
- Java Rowset 288
- Java Server Faces (JSF) Widget Library 331
- Java Server Page 272, 335
- Java Server Pages 271

- Java Specification Request 318
- Java Specification Requirement 288
- Java toolbox driver interface 357
- Java tools 287
- Java Transaction APIs 92
- Java version 366
- Java Virtual Machine 87
- Java Virtual Machine Debug Interface 288
- Java Virtual Machine Profiler Interface 288
- Java Virtual Machine Tool Interface 287
- JAX-P 1.2 - New properties for XML parsers 316
- JAX-R - XML registry API 316
- JCE 287
- JDBC 94, 289
- JDBC (Java Database Connectivity) parameters 366
- JDBC Provider 357
- JDBC provider 369
- JDBC URL 212
- JDK 5.0 support (WebSphere) 328
- JIT 286, 421
- JIT performance 422
- jitc 286
- JMS 315
- JMX 288, 318
- JMX remote API,V1.0 288
- job interrupt exit program 76
- job log management 70
- Job Log Output 70
- Job log pending status cleanup 71
- job termination 70
- Job Trace Analysis Summary report 431
- Job Trace Information report 431
- Job Watcher 439
 - graph and legend support 439
 - import interface 440
 - user defined queries and graphs 440
- Join 108
- Joining of DB2 tables in different iSeries partitions 90
- Journal Caching 97
- journal maximum object limit 96
- journal minimal data 96
- Journal Object Limit 411
- journal receiver size 411
- Journal Recovery Count 97
- journal tips for i5/OS 97
- Journaling 96
- journaling data areas 96
- journaling data queues 96
- journaling enhancements 147
- Journaling for SQL tables 303
- journaling performance 410
- JRNOBJLMT 411
- JSF 1.1 331
- JSF Widget Library 331
- JSP 272, 335
- JSR 288, 318
- JSR 077 318
- JSR 109 - Web services programming and deployment model 316
- JSR 116 331

- JSR 168 331
- JSR088 318
- JTA 92
- Just In Time 286
- Just-in-Time compiler 421
- JVM 87, 283
- JVM (Java Virtual Machine) parameters 366
- JVM 1.5 runtime 422
- JVM initial heap size 366
- JVM maximum heap size 366
- JVM settings Advise examples 367
- JVM TI 288
- JVMDI 288
- JVMPI 288
- JWL 331

K

- KEEP EXCLUSIVE LOCKS (SQL) 300
- Kerberos 142
- Kerberos server 141
- Kerberos tokens 330
- KEYFILE 414
- Konqueror 238

L

- L2TP 236
- LABEL ON INDEX (SQL) 302
- LAM router 232
- LAN line description 33
- large spool environment considerations 388
- Last Participant Support 317
- Last Query Statistic Use 105
- Last Query Use 105
- Last Used Date 105
- latest WebSphere Application Server for i5/OS PTFs 552
- Layer 2 Tunnel Protocol 236
- LCA password 374
- LDAP 480
- Level 1 458
- Level 2 458
- library flight400 346
- library list considerations for WebFacing 346
- Library lists in object lists (BRMS) 160
- Licensed Internal (LIC) Code Maintenance 495
- Licensed Internal Code
 - LIC 41
- Licensed internal code (LIC) 136
- Lifecycle Management Contract 315
- LIKE 98
- LIKE performance 421
- Like/Substring predicates 409
- Linux 33, 174, 473
- Linux and SIP 531
- Linux kernel version 238
- Linux on iSCSI 34
- Linux partition 255
- load source iop 49
- LOB columns 98, 409
- Local Certificate Authority (LCA) 374

- Local Console Manager and Monitor 127
- LOCK parameter on RGZPFM 413
- LOG(4 0 *NOLIST 70
- logical partitioning 245
- LOGOUTPUT 70
- longer column names 94
- Lotus Domino for i5/OS 484
- Lotus products 483
- Lotus Workplace on i5/OS 486
- Low Voltage Differential 23, 25
- LPAR 41
- LPAR OS Preload 56
- LPAR Partitioning Initialization 56
- LPAR performance considerations 423
- LPAR simplification 257
- LPAR simplification tools 257
- LPAR Simplification Tools Handbook, 259
- LPAR Validation Tool 56, 257
- LS20 126
- LS21 126
- LS41 126
- LVD 23, 25
- LVD interface. 37
- LVT 257
- LWCANCEL(*YES)
 - considerations 414

M

- mage catalog 181
- Maintained Temporary Index 98
- Maintained Temporary Indexes 98
- Malformed packets 139
- malformed packets 136
- manage application server and HTTP server attributes (Web Performance Advisor) 359
- Manage attributes 365
- manage system attributes (Web Performance Advisor) 359
- Management Information Base 464
- Management Module 129
- Management Modules 127
- manageProfiles 330
- Manufacturer Type and Model 392
- Mapping Suite 399
- Mapping Suite (5639-AAC, AAD, AAE) 390
- Mapping Suite product 378
- MarkMagic 399
- materialized query table 102
- Materialized Query Table Evaluator 106
- Materialized Query Table support 86
- materialized query tables 408
- Max rows attribute 91
- MAXCPU 437
- maximums for a table 98
- MBeans 288
- MCU ratings 404
- Mechanical Desktop 399
- Media policy 168
- MEDPCY 168
- Megaco telephony protocol 530

- memory 15, 360
- Message authentication 238
- Message Inflow Contract 315
- Metro Mirror 51, 119
- Metro Mirror (synchronous PPRC) 51
- MFP (multi-function printer) 379
- MFP capable devices 399
- MFP devices table 399
- MFRTYPMDL 392
- MI 137
- MIB 464
- Microsoft Intermediary Language 87
- Microsoft System.Data.Odbc provider 90
- Microsoft System.Data.OleDb provider 90
- Microsoft Visual Studio 90
- Microsoft Visual Studio Integrated Development Environment. 88
- Midrange Performance Group (MPG) 441
- migrating Linux servers 120
- MINCPU 437
- minimized journal entries 96
- Minimum and Maximum Individual CPU Time Used
FIELDS 437
- missed object policy (BRMS) 160
- MOM 460
- MQ-Explorer type management 320
- MQSeries 329
- MQT 102
- MQTs 408
- MS FW 2.0 compatibility 88
- MS Operation Manager 460
- MSIL 87
- MTI 98
- Multi-column predicates(SQL) 300
- Multi-function Printers, MFP devices 399
- Multiple active result sets on a connection 88
- Multiple Proxy ARP 231
- multiple threads 73
- multiple virtual disk paths 122
- multi-threaded support 238
- multithreading (NetServer) 238

N

- NAT 226, 544
- NAT firewalls 234
- NAT traversal support 234
- Native JDBC driver 369
- Native spool file save/restore 160
- Neoware, Inc. 253
- NetServer 221
- Netserver performance 421
- NetServer scalability and performance enhancements 239
- NETSTAT 225
- Network Address Translation 226
- Network Authentication Enablement (5722-NAE) 141
- Network connectivity enhancements 33
- Network enhancements (BRMS) 168
- Network Server Description 130
- Network Server Host Adapter 130

- NWSH 130
- network sniffing 139
- Network Transfer Format (Java) 288
- New and enhanced classes (Java) 289
- new printers in HPT 392
- NLSS/CCSID translation between columns 409
- Node Groups 318
- Non-IBM sizing tools 451
- NSTEAD OF Triggers 302
- number of disk units 360
- NWSCFG 130
- NWSD 130
- NWSH 130

O

- OASIS 273
- object conversion 422
- Object Database Connectivity 92
- Object Pools 317
- ODBC 87, 92, 94, 241
- ODBC driver 241
- ODBC/JDBC jobs 77
- Odyssey 331
- OLAP expressions (Java) 297
- OLE DB provider 241
- On Demand Index capability 86
- On demand memory and cross-partition movement
(LPAR simplification) 266
- Open DataBase Connectivity 87
- Open Service Gateway Initiative 328
- Openness 84
- OpenSSL 145
- operation codes (RPG) 279
- Operations Console 248
- Operations console adapter placement 248
- Operations console support 246
- optical type reference catalog image files 182
- Option 29 of 5722SS1 120
- Order 108
- OSGi 328
- OSGi Alliance 327
- OUTBIN attribute 395
- output queue example 382
- OUTQ and spooled files in ASPs 383
- Override with Printer File 386
- OVRPRTF 386

P

- P/N 30R5201 - Copper 126
- P/N 30R5501 - Fiber 126
- Packet filtering 544
- parallel library save enhancement 153
- parallel rollback unlock 100
- Parallel save, restore enhancements 161
- Partition Load Manager 255
- Partition Preload Specifies 55
- partition restart 96
- partitioned table 98
- Partitioned Table optimization 98

- PASE 74
- Passport Advantage Web site, 487
- Pass-thru Module 126
- PATROL for iSeries - Predict 452
- Patrol Predict 452
- PCL 379
- PCL outbin control 395
- PDF 390
- PDF printing 400
- PDF printing via Infoprint Server 379
- PDF417 391
- PDFs for IBM Director at the IBM Support Web site 552
- PDM - Program Development Manager 304
- PDU 127
- Peer-to-Peer Remote Copy 51
- Pegasus 464
- performance
 - integrated xSeries disk I/Os per second rules of thumb 427
 - integrated xSeries servers 424
 - LPAR 423
- Performance Capabilities Reference manual 119
- performance data graphs
 - print 430
- performance enhancements 97
- Performance Management for IBM eServer iSeries and pSeries
 - A Systems Management Guide, SG24-7122 448
 - Performance Management for System i5 444
- Performance Management offering (PM for System i or p) 264
- performance management Website 404
- Performance Monitoring Infrastructure 370, 443
- Performance Navigator 441, 451
- Performance Setting Web page 364
- Performance Tools for i5/OS 428
 - Job Trace 428
- performance tools plug-in 430
- performance tools updates 427
- permanent processor activations 511
- PHP Web deployment environment 542
- PHP Zend for i5/OS, SG24-7327, redbook 542
- PHYP 42
- PING 225
- PJL 379
- plan cache snapshot 111
- planning statements 546, 550
- Plants by WebSphere application 323
- PLM 255
- PM for System i5 445
- PM for System i5 data integration with the IBM Work Load Estimator 451
- PM iSeries 444
- PM pSeries 442, 444
- PM redbook 264
- PMI 443
- PMiSeries IS tool 449
- PMiSeriesIS 449
- Point-to-Point Protocol 236
- policy-based output queue management 386

- Port scans 139
- Portable Application Solution Environment 74
- Portable Utilities for i5/OS 144
- Portal Tool Kit 305
- portlet 273
- Portlet programming model 331
- POSITION DDS keyword (printers) 397
- Position keyword (DDS) 396
- positioning HA and CBU editions 512
- post filtering 110
- Postscript 390
- Power Distribution Units (PDU) 127
- POWER Hypervisor 41
- POWER Linux partition 445, 532
- POWER5 Hypervisor 41
- POWER5-based systems and i5/OS console support 35
- PPRC Extended Distance 51
- PPRC with iASP 52
- PPRC-XD 51
- preferred interface list (IP) 229
- pre-filtering 109
- prepare to commit 91
- prepare to rollback 91
- print fidelity 385
- Print function enhancements 215
- Print output formatting and routing 389
- Print Services Facility 379, 396
- Printer Command Language 379
- printer output enhancements 378
- printers 400
- Process Management 470
- processor 471
- processor (SP) 125
- processor feature 360
- Product education roadmaps 516
- Product home pages 516
- Product newsgroups and forums 516
- Product recommended updates 516
- Product support pages 516
- Product troubleshooting guides 516
- profile 318
- profiles (WAS) 326
- program QSYS/QDBFIMI 423
- Programmer's Toolkit 87
- PRTJOBTRC command 429
- PSF 379
- PSTN 531
- PTF group level 7 340
- PTF groups recommendation 351
- PTF MF34406 23
- PTF SI24255 338
- PTF tabs - Web Performance Advisor 361
- PTFs
 - MF39611, MF39933 26
 - MF39614, MF39926 26
- PTFs for integrated server support 120
- Public Switched Telephone Network 531

Q

- QAITMON file 437

QAPMWASAPP file 371
 QAPMWASCFG file 371
 QAPMWASEJB file 371
 QAPMWASRSC file 371
 QAPMWASSVR file 371
 QAQQINI options 99
 QASYPOJ audit file 384
 QAUDLVL 384
 QAUDLVL2 384
 QDFTJRN data area 96
 QDFTJRN support extension 96
 QEZDEBUG 387
 QEZJOBLOG queue 387
 QLOGOUTPUT system value 70
 QOS 544
 QPEZJOBLOG 387
 QR Code 391
 QSERVER 238
 QSHELL 422
 QSPGETSP 384
 QSPLFACN system value 380
 QSQRVR job 357, 369
 QSRRSTO API 382
 QRSAVO 152
 QRSAVO and QSRRSTO APIs 382
 QSYS2 96
 QTIMZON 78
 QTOCC4IF 229
 QTPPPL2SSN 236
 QTPPSSN 236
 Quality of Service (QoS) 139
 Query Statistics Use 105
 Query Storage Limit parameter 102
 Query Use Count 105
 Quick View 219
 QWAS6 326
 QWAS61 331
 QWPIPCOLOR object 395
 QWPLEXMARKC object 395
 QYCMCIMOM job 463
 QYMAP2MIB job 464
 QZDASOINIT 95
 QZDASOINIT job 77, 357, 369
 QZLSFILET 238

R

Rack and Power Distribution Units 127
 rack features 39
 radio frequency identification 398
 radio frequency identification printersw 377
 RAID-5 442
 RAID5 adapters with write cache 22
 RAID5 and RAID6 comparison 48
 RAID-5 with #5738/#5777 27
 RAID-6 442
 RAID6 26
 RAID-6 with #5738/#5777 27
 RANK and DENSE_RANK (SQL) 299
 RAT 385
 Rational Developer family of products 332
 Rational Product Updater (RPU) 331
 Rational Web Developer 331
 Raz-lee's Security's iSecurity Suite 144
 Reclaim Database Cross-Reference (RCLDBXREF) command 96
 recursive common table expressions 291
 recursive expression (SQL) 295
 Recursive SQL 86
 recursive SQL considerations 296
 recursive views 291
 Redbooks Web site 553
 Contact us xxi
 refacing tool 271
 refacing tools 270
 refacing with IBM business partner tools 271
 reference image catalog 181
 References to DDS logical files 409
 Referential Integrity (RI) 413
 Relative position keyword (DDS) 396
 Remote - IBM Toolbox Java JDBC driver 369
 Remote Journal 97
 remote server discovery 125
 Remote Supervisor Adapter 129
 Remote System Explorer 307
 Remote System object 130
 reorganize physical file member 411
 reorganize table 411
 Requests Received 435
 Reserve Capacity on Demand 245, 265
 reserve CoD processor day 265
 Resource Allocation Table 385
 Resource Managers 92
 resource managers 92
 Resources tab (advisor) 368
 Responses Sent 435
 Restore Objects and Libraries with BRM 162
 Restore Objects List API 382
 Restricted IP options and protocol 139
 Restriction of Hazardous Substances 20
 ResultSet (Java) 288
 Resume key (RSMKEY) 165
 Retrieve Call Stack (QWVRCSTK) API 74
 RFC4304 235
 RFID 379, 398
 Infoprint 6700 398
 RFID printers 377
 RGZPFM 411
 RI (binary radix index) 413
 Ricoh Aficio SP9100DN Printer Series 395
 RLU - Report Layout Utility 304
 RMs 92
 RoHS 20
 ROLLBACK 92
 ROW_NUMBER (SQL) 298
 Rowset (Java) 288
 rowset functions 288
 RPG 280
 RPG and COBOL System/36™ and System/38™ compatible compilers, 548
 RPQ 163

- RS-232 32
- RSA II 129
- RSAll Slimline 126
- RSTLICPGM 325
- RTAS 41
- RUN SQL - new destination option 212
- Run time architecture
 - JAVA 32-bit, JAVA 64-bit 284
- Running WebSphere Application Server V6.0 prerequisites 325
- Run-Time Abstraction Services 41
- Runtime i5/OS Subsystem 326
- Runtime i5/OS subsystem 331
- RWD 331
- RWD (Rational® Web Developer) V6.0 305

S

- S/36 and S/38 Environments 548
- S/38 Environments 548
- SAAJ 1.1 - SOAP Attachments API for Java 316
- Samba version 238
- sample Flight400 application
 - HATS URL 348
 - WebFacing URL 347
- SamplesGallery 322
- SamplesGallery application 323
- SAN 31
- SASL 288
- SATF 224
- SAV 152
- SAVCFG 152
- SAVCHGOBJ 152
- SAVDLO 152
- save data queue 152
- Save joblogs (BRMS) 157
- Save System Information 160
- Save/Restore of spooled files 381
- save/restore performance consideration (520) 9
- save/restore performance consideration (550) 11
- save/restore performance consideration (570) 13
- Saved Objects (BMRS) 168
- SaveFile 290
- SAVESECDTA 152
- saving and restoring spool file examples 381
- SAVLIB 152
- SAVOBJ 152
- SAVSYS consideration for virtual tape 171
- SAVSYSINF 160
- SAVSYSINF backup strategy 152
- scalability offered by iSCSI 122
- Scalar Fullselect 84–85
- scalar-fullselect 291
- SCANFSCTL system value 144
- Scheduler Service 317
- SDA - Screen Design Aid 304
- SDK 422, 460
- SDO 323
- Secure SHell 144
- security 135
- SELECT (SQL) 300

- Sensitive Cursor 409
- Sensitive Cursors 98
- SEQUENCE object 96
- Series Access for Web 271
- Series Navigator Tasks on the Web 338
- ServeRAID Manager 473
- Service Data Objects 323
- Service integration technologies in Version 6 320
- Service Processor 42
- service processor connections (iSCSI) 127
- Service Processor object
 - NWSCFG 130
- service support level console function changes 254
- Service Tools 30
- servlet 272
- Session Initiation Protocol 331, 531
- Session Initiation Protocol telephony protocol 530
- SESSION_USER special register 94
- SESSION_USER special register 291
- SET SESSION AUTHORIZATION statement 290
- SET SESSION Authorization statement 94
- Setting up the HTTP Administration server to use SSL 373
- SEU - Source Entry Utility 304
- SF 41
- SF235_160 250
- SF99348 group PTF 119
- SF99503 94
- Shadow catalog views 185
- shared access transport facility 224
- Show Materialized Query Tables 106
- SI24716 78
- SI24717 78
- SI24906 78
- SI24908 78
- SI25364 395
- SI25365 395
- SI25430 395
- SI25431 395
- SI2545 395
- SI25450 395
- SignonHandler (interface) 290
- Simple and Protected GSS-API Negotiation Mechanism 330
- Simple Authentication and Security Layer 288
- Simple Network Management Protocol 464
- Simple Object Access Protocol 273
- simplified set up for HTTP Admin server 372
- simplified use of SSL 372
- Single system Telephony Express configurations 537
- SIP 331, 530–531
- Small and Medium Business 163
- Small Computer System Interface controller 23
- Small Computer System Interface controller for disks with 90 MB cache 25
- SMAPP 96
- smart IOA 8, 19, 22, 52, 54
- SMB 163
- smbclient 238
- smbmount 238

SMP 412
 SNA 19, 221
 SNA support with Enterprise Extenders 223
 SNMP 464
 SNMP Devices 465
 SNMP devices 459
 SNMP devices, 459
 SOAP 273
 soft commit 100
 Software Development Kit 422
 Software Information Center xvii
 Solution Visualizer (LPAR simplification) 260
 Spanning Tree Portfast 127
 Special edition offering upgrade path options 67
 special register (SQL) 290
 SPLFACN JOB attribute 380
 SPLFACN(*DETACH) 71
 SPLFRY IPL attribute 71, 398
 SPNEGO 330
 spool file expiration restriction 387
 spooled file attributes not preserved 385
 spooled file expiration 386
 spooled file restore considerations 385
 spooled files save considerations 384
 SPT 258–259
 SPT and WLE 264
 SPT education 264
 SPT output 260
 SPT Web site 259
 SQE 86, 97, 112
 SQL and CL integration 303
 SQL and table limits 94
 SQL Descriptor Area 93
 SQL enhancements) 290
 SQL optimizer 95
 SQL performance
 compare SQL monitor plan cache snapshot 111
 SQL performance monitor 109
 SQL performance monitors 111
 SQL Plan Cache 112
 SQL Plan Cache viewer 109
 SQL Procedures 86
 SQL procedures 95
 SQL program object conversion 422
 SQL Query Engine 86
 SQL Query Engine (SQE) enhancements 97
 SQL standards compliance 85
 SQL standards support 85
 SQL statistics 112
 SQLColAttributes 91
 SQLDA 93
 SQLFetchScroll block-fetch 91
 SQLGetInfo 91
 SSH 144
 SSH protocol suite 144
 SSL 325, 372
 SSL Inbound Connections 435
 SST 254
 Standard Flagger 84
 StandGuard Antivirus for iSeries from Byteware 144
 Start IP Interface command 228
 Start Job Trace maximum storage 428
 start journaling 96
 Start Recovery using BRMS 161
 Start Service Tools command 496
 Start Trace 238
 Start Trace command 428
 Starting to use Web Enablement Environment 344
 Startup Beans 317
 statement objects 92
 Statistics Requests 113
 StereoGL 399
 STGL 399
 storage pool rules of thumb for integrated xSeries performance 426
 STORAGE_LIMIT QAQQINI 102
 Stored Procedures 95
 STP Fast Start 127
 STRDBMON filter 101
 STRJOBTRC 429
 STRJOBTRC thread ID 429
 STRSST 496
 STRTC 428
 STRTCPIFC 228
 STRTRC 238
 Structured Query Language enhancements 290
 Style sheet support (iSeries Access for Web) 210
 Subquery 85
 SUBR parameter 275
 Subroutine support 275
 SUBSTR 98
 Substring 409
 Subsystem (Java interface) 290
 subsystem enhancements 73
 Support Assistant 515–516
 Supported IXS/IXA Linux versions 34
 SUSE Linux Enterprise Server 9 126
 Symmetric Multiprocessing feature 412
 Synchronous PPRC
 Peer-to-Peer Remote Copy 119
 SYSCSTCOL 96
 SYSIBM 96
 SYSINST 195
 SYSTABLES 96
 system CPW 360
 System Firmware 41
 system firmware 491
 System firmware level 493
 System i Editions
 Buying and Selling Guide, 57
 System i integration with BladeCenter and System x PTFs 119
 System i IP telephony powered by 3Com 529
 System i planning and upgrade Web site xviii
 System i Planning Web site 1
 System i strengths 84
 System i telephony information 529, 539
 System i5 520 hardware changes 8
 System i5 550 hardware changes 10
 System i5 570 hardware changes 12

- System i5 595 hardware changes 14
- System i5 code administration and terminology 40
- System i5 editions 56
- System i5 Innovation 275
- System i5 memory features 15
- System i5 System Firmware level 495
- system integrity 136
- System Licensed Internal Code 41
- System Managed Access Path Protection 96
- system model 360
- System Monitor
 - coordinate graphs 428
 - visualize colors 428
- System Monitor average line utilization 427
- System Naming and Library List support 88
- system plan 260
- System Plan Tool 56
- System Planning Tool 258
- system pointer 137
- System report
 - ASP performance information 433
 - disk utilization section 433
 - hits per second 435
 - HTTP Server summary 435
 - storage pool utilization 432
 - workload section 434
- System report changes 432
- System Resources 362
- System Service Tools 254
- System Storage DS6000 119
- SYSTEM USER special register 94
- System x integration 118
- System x support via iSCSI 121
- Systems Network Architecture 19
- SYSTRIGCOL 96

T

- table
 - maximum rows 98
 - maximum size 98
- Tables with select/omit logical files 409
- TAM 322
- tape devices 22
- tape devices supported 25
- Tasks 462
- TCP/IP 33
- TCP/IP configuration enhancements 228
- TCP/IP Offload Engine 33
- TEC 399, 460
- Technology Samples 323
- telephony 530
- Telephony edition models 443
- temporary storage governor 102
- text file format 399
- TFC 234
- thermal data stream design and production 399
- Thin Console 253
- Thin console setup 253
- tickets (Kerberos) 142
- Tivoli Authorization Module 322

- Tivoli Performance Viewer 320
- Tivoli products 481
- Tivoli Storage Manager 164
- Tivoli Storage Manager Application Programming Interface 165
- Tivoli Upward Integration 460
- TM 92
- TOE 33
- tool workbench 516
- TotalStorage ESS and DS6000 and DS8000 441
- Trace Analysis I/O Summary report 432
- TRACEROUTE 225
- Traffic Flow Confidentiality 234
- transaction counts 356
- Transaction Manager 92
- Transaction Manager enhancements 100
- transparent subnetting (virtual Ethernet) 231
- triggers 98
- triggers (database) 86
- Triple DES encryption 86
- TrueType fonts 395
- trunk ports 127
- Truth Web site 550
- TS1120 Tape Drive 37
- TS3400 Tape Library 37
- TS7510 163
- TS7510 and BRMS 163
- TSM 164
- TSM performance report 165
- Tunnel Mode 234
- two-phase commit 91
- Type 2 files 172

U

- ubquery and Scalar Fullselect (SQL) 300
- UDDI 273
- Unicode (ICU) sort sequence tables 96
- Unicode updates (printer support) 395
- Unified clustering 321
- Unique indexes 413
- United States Government Accounting Office 227
- Universal Description, Discovery and Integration 273
- UOM parameter on the CRTPRTF command 397–398
- Update Manager task (IBM director) 480
- User accounts defined on an LDAP server (IBMDirector). 480
- User Datagram Protocol 233
- user exit program (antivirus) 144
- user id management (IBM Director) 480
- User Media Library (MLB) support 157
- User profile credentials (iSeries Access for Web) 217
- Using a Dependent catalog 181
- using dependent catalogs (virtual tape) 182

V

- V5R3 24
- V5R3M0 LIC 37
- V5R3M5 Licensed Internal Code 8
- V5R4 5722-XE1 iSeries Access for Windows 141

- V5R4 i5/OS security overview 136
- V5R4 iSCSI support configuration objects 131
 - view 2 of 2 131
- variables (CL) 276
- vault credentials 217
- VCX platform 531
- VIPA 229
- VIPA proxy ARP agent 229
- Virtual Basic 242
- Virtual Ethernet 229
- Virtual Ethernet configuration 231
- virtual I/O performance 421
- Virtual I/O restructuring 421
- Virtual I/O server (AIX) 255
- Virtual IP 229
- virtual IP address 229
- Virtual IP address on a different subnet 232
- virtual IP address proxy 229
- Virtual media 173
- Virtual Partition Manager 532
- Virtual Partition Manager 256
- Virtual Private Networking 233
- virtual tape devices with guest partitions 174
- Virtual tape support of the Virtualization Engine TS7510 163
- virtual tape usage and disk storage consumption 187
- virtual tape volume storage allocation 178
- Virtual volume support (BRMS) 157
- Virtual volumes as IFS stream files 172
- Virtualization Engine 460
- Virtualization Engine for Tape Console 163
- Virtualization Engine TS7510 153
- virus 144
- Visual Basic 242
- Visual Explain 113
- VMTM Tool Interface 288
- Voice over IP 530
- Voice over IP VoIP 331
- Voice-over-IP 529
- VoIP 530
- VOLATILE table support 100
- Volatile tables 100
- volume auto creation 173
- Volume Shadow Copy 120
- VPM 256, 532
- VXA 37
- VXA-320 Internal Tape Drive 37

W

- WAN adapters 32
- WAS security model 210
- WAS V6
 - saving the product, configuration, and data 326
- WATK 332
- WBEM 463
- WCS 487
- WDHT 310
- WDS V6.0.1 packaging 306
- WDS Lite 311
- WDSC Lite technology 306

- WDSC V6.0 305
- WDSC V6.0.1 306
- WEAPPSVR port 344
- Web Application Server 345
- Web authentication 330
- Web authorization 330
- Web Based Enterprise Management initiative 463
- Web Enablement for i5/OS 536
- Web Environment area (advisor) 366
- Web Environment JVM settings 366
- Web Environment of Web Performance Advisor 364
- Web Performance Advisor 350, 359
 - performance attributes 358
- Web performance advisor 358
- Web performance management tools 350
- Web Performance Monitor 350, 352
- Web Performance tools 451
- Web Services 273, 341
- Web Services by WebSphere release 316
- Web Services Description Language 273
- Web Services Gateway Filter Programming Model 317
- Web Services Samples 323
- Web System Manager 251
- Web-based access 473
- Webfaced version of Flight400 311
- WebFacing 341
- WebFacing server 345
- WebFacing Tool 306, 308
- WebSM 251
- WebSphere Administrative Console 357
- WebSphere Application Server 271
 - multiple edition support 328
- WebSphere Application Server administrator 329
- WebSphere Application Server for i5/OS V6.0 additional documentation 326
- WebSphere Application Server for i5/OS V6.1 additional documentation 333
- WebSphere Application Server installation changes 328
- WebSphere Application Server performance data 370
- WebSphere Application Server performance expectation 422
- WebSphere Application Server profile 213
- WebSphere Application Server security enhancements 330
- WebSphere Application Server V6.0 for i5/OS 314
- WebSphere Application Server V6.0 instance 318
- WebSphere Application Server V6.0 on i5/OS packaging 324
- WebSphere Application Server V6.0 product packaging 325
- WebSphere Application Server V6.0 profile 318
- WebSphere Application Server V6.0 profiles 326
- WebSphere Application Server V6.0 security 322
- WebSphere Application Server V6.0 supplied applications 322
- WebSphere Application Server V6.0 system management features 318
- WebSphere Application Server Version 5.1.1 322
- WebSphere Application Server Version 6.1 on i5/OS 327
- WebSphere Application Support 313

- WebSphere Bank application 323
- WebSphere Datasource 212
- WebSphere Development Studio 304
- WebSphere Development Studio and Client 304
- WebSphere Development Studio Client for iSeries information center 311
- WebSphere Development Studio Client for iSeries support 311
- WebSphere Development Studio Client Studio Client Advanced 305
- WebSphere Development Studio for iSeries Client 349
- WebSphere function comparison 327, 334
- WebSphere Information Center 327, 334
- WebSphere installation 332
- WebSphere MQ 320, 329
- WebSphere performance 421
- WebSphere Studio Application Developer 305
- WebSphere Version 6 HA terminology 321
- WEHTTPSVR port 344
- WIN32 services and device drivers 470
- Windows 2003 Volume Shadow Copy 120
- Windows Internet Explorer 343
- Windows Network Neighborhood (iSeries NetServer 221, 237
- Windows Server 2003 with SP1 126
- Windows Server 2003 with SP1 or R2 (iSCSI) 126
- Windows Server release level 126
- WLE 258, 264, 445
- WLE Website 264
- Work Control Block Table 398
- Work Management Contract 315
- Work management enhancements 70
- Work with Image Catalog Entries 183
- Work with Job Log 388
- Work with Job Logs 72
- Work with Journal (WRKJRN) command 96
- Work with Saved Objects 168
- Work with Shared Pools 268
- Work With Spooled Files 388
- Work with System Activity 437
- WorkArea Service 317
- workgroup printers
 - 15xx printers 400
- Workload Developer component (WLE) 444
- Workload Estimator 258, 264, 342, 441, 552
 - TotalStorage ESS and DS6000 and DS8000 441
 - Virtual I/O workload 442
- workload Estimator 442
- Workload Estimator (WLE) 539
- Workload management high availability 321
- WorkManager 317
- Workplace Collaboration Services 487
 - new features 488
- Workplace Designer 489
- Workplace Services Express 486–487
- Workplace Services Express Information Center 487
- WorkStation Customization Object 392
- Worm 163
- write cache 22–23
- Write Once Read Many 163

- WRKACTJOB - current user 76
- WRKIMGCLGE 183
- WRKJOBLOG 388
- WRKJOBLOG) 72
- WRKSHRPOOL 268
- WRKSPLF 388–389
- WRKSYSACT 437
- WRKTCPPTP 236
- WSCST 392
- WSDL 273
- WSE 486
- WS-Security 1.0 WebServices Specification 322

X

- X/Open Distributed Transaction Processing Model 92
- XA 91
- XA and JTA 91
- XA over DRDA 91
- XGL 399
- XML 115, 271–273
- XML GENERATE (COBOL) 283
- XML processing 99
- XML Toolkit 303
- XML. 318
- XML-INTO (RPG) 280
- XML-SAX (RPG) 280
- xSeries 464
- xSeries integration
 - expand a disk drive 121
- xSeries integration - Licensed Program Product changes 120
- xSeries model support withdrawn 120

Z

- Zend
 - hardware and OS requirements 542
- Zend Core for IBM i5/OS 542
- Zend PHP Web deployment environment. 542
- Zend Platform for IBM i5/OS 542
- Zend Studio for IBM i5/OS 542
- Zip function 217
- zlib 145
- ZPL (Zebra) 399



Redbooks

IBM System i5 V5R4 Technical Overview



IBM System i5 V5R4 Technical Overview



More power to the System i family and i5/OS

This IBM Redbook provides a technical overview of IBM System i family hardware and i5/OS-based software at Version 5 Release 4 level. It is intended for those already familiar with System i and i5/OS who need a broad perspective of what is new with V5R4 either to consider upgrading to V5R4 or to take full advantage of V5R4 capabilities.

Expanding innovative solutions

Several V5R4-based licensed program capabilities are also included. These include Lotus products, WebSphere Application Server versions, WebSphere Development Studio for iSeries capabilities, and IBM Director for i5/OS capabilities.

More management tools for your IT solutions and business

Selected i5/OS V5R3 level content is included where the developers of this book determined that the function should be included. Support introduced with V5R3 is clearly noted.

This information is generally of an intermediate level of detail. Providing this information in one place is an ideal way to obtain a good understanding of capabilities. This book refers the reader to sources for more detailed information about specific topics.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7271-00

ISBN 073848962X