


IBM i 6.1 Technical Overview

Manage complexity and risk with new Web-based management and security functions

Boost efficiency with virtualization and dynamic workload management

Realize innovation with open application design choices



Jim Cook
Ahmed Mansour
Giancarlo Omati
Ana Flavia Pavan
Mark Roy

Redbooks



International Technical Support Organization

IBM i 6.1 Technical Overview

December 2009

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

First Edition (December 2009)

This edition applies to Version 6, Release 1, Modification 0 of IBM i (5761-SS1) and related licensed program products.

© Copyright International Business Machines Corporation 2009. 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	xv
Trademarks	xvi
Preface	xix
The team who wrote this book	xix
Become a published author	xx
Comments welcome	xx
Chapter 1. IBM i 6.1 introduction and summary	1
1.1 Overview of what is new in IBM i 6.1	2
Chapter 2. Security enhancements	5
2.1 Intrusion detection and prevention	6
2.1.1 Variable dynamic throttling	6
2.1.2 Setting up intrusion detection	7
2.1.3 Remove TCP/IP Table (RMVTCPTBL)	11
2.2 Encryption of data at rest on disk	11
2.3 Encrypted backups	12
2.4 New security system values	13
2.4.1 Password rules (QPWDRULES)	13
2.4.2 Password expiration warning (QPWDEXPWRN)	16
2.4.3 Block password change (QPWDCHGBLK)	16
2.4.4 Secure Sockets Layer cipher specification list (QSSLCSL)	16
2.4.5 Secure Sockets Layer protocols (QSSLPCL)	18
2.4.6 Secure Sockets Layer cipher control (QSSLCSLCTL)	18
2.5 Changed security system values	18
2.5.1 Limit device sessions (LMTDEVSSN)	18
2.5.2 Verify object on restore (QVFYOBJRST)	19
2.6 Restricting changes to security settings	19
2.7 User profile enhancements	19
2.7.1 Block password change (PWDCHGBLK)	19
2.7.2 Limit device sessions (LMTDEVSSN)	20
2.7.3 Display User Profile (DSPUSRPRF)	20
2.8 Service Tools user profile and IBM i profile	21
2.8.1 Display Service Tools User ID (DSPSSTUSR)	21
2.8.2 Service tools user ID and IBM i user profile association	23
2.9 Object signing and signature verification	23
2.10 Digital Certificate Manager	24
2.10.1 Starting Digital Certificate Manager	24
2.11 Single sign-on	24
2.12 Enterprise Identity Mapping	24
2.13 Network Authentication Service	25
2.13.1 New Kerberos control language commands	25
2.13.2 New service principal for Network File System server	26
2.14 Cryptographic enhancements	26
2.14.1 Cryptographic key management enhancements	27
2.14.2 Save/restore master key	28
2.14.3 4764 Cryptographic coprocessor	28

Chapter 3. Backup and recovery	31
3.1 File level backup for Linux	32
3.2 Save-while-active	32
3.2.1 Single synchronization point for multiple save commands	32
3.2.2 Saving storage spaces for active integrated servers	33
3.3 Deferring the restore of dependent objects	34
3.4 Saving and restoring private authorities	35
3.5 Alternate installation device	35
3.6 Recovering vital product data	35
3.7 Program conversion restoring to a different hardware type	36
3.8 Program conversion restoring to a different release	36
3.9 Restore user profile	36
3.10 User-defined file system backup	37
3.10.1 Omitting objects during a save	37
3.10.2 Saving a mounted user-defined file system	37
3.11 Restore actions for mounted user-defined file systems	38
3.12 Encrypted backups	38
3.13 Encrypted auxiliary storage pools	39
3.14 Restore the operation output file	39
3.15 Restoring external stored procedures	40
3.16 Backup and restore for IBM i on BladeCenter	40
3.17 Retrieve Disk Information (RTVDSKINF)	41
3.18 Backup, Recovery and Media Services	41
3.18.1 BRMS graphical interface	41
3.18.2 Command and display enhancements	52
3.18.3 Recovery enhancements	60
3.18.4 Media services enhancements	63
3.18.5 Hierarchical Storage Manager enhancements	66
3.18.6 Device enhancements	67
3.18.7 Network enhancements	69
3.18.8 Install enhancements	69
3.18.9 Other enhancements	69
3.19 Journal management	70
3.19.1 Journaling libraries and automatically starting journaling	70
3.19.2 Apply Journalled Changes (APYJRNCHG) enhancements	70
3.19.3 Using data port services with remote journaling	71
3.19.4 Additional remote journal options and information	71
3.19.5 Exit point for detaching journal receivers	71
3.19.6 Ending the journaling of opened files	71
3.19.7 Enhancements to start journaling and end journaling commands	71
3.19.8 System-managed access path protection	72
3.19.9 Replay Journal Entry (QjoReplayJournalEntry) API	72
3.19.10 Changing the journal recovery count and the cache wait time	72
3.20 Restrictions for previous release support	73
3.20.1 NetWare Enhanced Integration	73
Chapter 4. High availability	75
4.1 PowerHA for i	76
4.1.1 Geographic mirroring	77
4.1.2 Metro Mirror	78
4.1.3 Global mirror	78
4.1.4 FlashCopy	79
4.1.5 Using PowerHA for i and IBM Total Storage Copy Services functions	79

4.1.6	Further information	80
4.2	Clustering	80
4.2.1	Cluster management jobs	80
4.2.2	Switchable devices	81
4.2.3	Cluster-level and cluster resource groups-level messages.	81
4.2.4	Cluster administrative domain.	82
4.3	IBM iCluster for i	83
4.3.1	Complete integration with IBM i	84
4.3.2	Additional iCluster enhancements.	85
4.4	Independent auxiliary storage pool	85
4.4.1	Quiesce I/O for an auxiliary storage pool	85
4.4.2	Target side tracking for geographic mirroring	86
4.4.3	System base commitment control transactions.	86
4.4.4	Job queues	86
4.5	OptiConnect	87
4.5.1	Logical name support for DDM files	87
4.5.2	Work with OptiConnect Activity (WRKOPCACT) command	88
Chapter 5.	Database	89
5.1	Database performance and query optimization	90
5.1.1	SQL Query Engine	90
5.1.2	SQL runtime performance improvements.	91
5.1.3	Temporary distinct sorted list	92
5.1.4	Temporary values list	92
5.1.5	Optimizing performance using the plan cache	93
5.1.6	QAQQINI updates.	93
5.1.7	Derived key index	94
5.1.8	Full outer join	95
5.1.9	New database monitor format.	96
5.2	Distributed database programming	96
5.2.1	Distributed Relational Database Architecture	96
5.2.2	Distributed Data Management	99
5.3	OnDemand and availability	99
5.4	Commitment control and independent disk pools	99
5.4.1	System base commitment control transactions.	99
5.4.2	Quiesce processing	100
5.5	http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzakj/rzakjoptimize.htm Miscellaneous database enhancements	100
5.5.1	Run SQL Statements (RUNSQLSTM)	100
5.5.2	DB2 for i catalog views	101
5.6	IBM OmniFind Text Search Server for DB2 for i5/OS	102
5.7	System i Navigator database task	103
5.7.1	Create multiple indexes with Run SQL Scripts	103
5.7.2	Add a comment to a database object	107
5.7.3	Reset usage counts for indexes and materialized query tables	107
5.7.4	Show indexes for all tables in a schema.	109
5.7.5	Show indexes for a specific table	111
5.7.6	Show materialized query tables for all tables in a schema	113
5.7.7	Show materialized query tables for all views in a schema	114
5.7.8	Show materialized query tables for a specific table	115
5.7.9	Show materialized query tables for a specific view.	116
5.7.10	Show lock holders.	117

5.7.11	SQL details for jobs	118
5.7.12	Change the SQL plan cache size threshold	118
5.7.13	SQL plan cache statements	119
5.7.14	SQL plan cache event monitor	121
5.7.15	Working with an SQL plan cache event monitor	123
Chapter 6.	Performance	125
6.1	Java and WebSphere performance	126
6.1.1	IBM Technology for Java, J9	126
6.1.2	Classic JVM (64-bit)	127
6.1.3	IBM Technology for Java (32-bit)	127
6.1.4	IBM Technology for Java (64-bit)	127
6.1.5	Choosing which JVM to use	128
6.1.6	Comparing Java 5.0 to Java 6	128
6.1.7	Determining JVM heap size for IBM Technology for Java JVMs	128
6.2	Java Performance Tools	129
6.2.1	Collection Services integration	129
6.2.2	The Work with JVM Jobs (WRKJVMJOB) command	129
6.2.3	The Print JVM Job (PRTJVMJOB) command	132
6.2.4	The Generate JVM Dump (GENJVM DMP) command	133
6.3	IBM Performance Tools for i5/OS (PT1) packaging	134
6.3.1	Disk Watcher	136
6.3.2	Job Watcher	136
6.3.3	Commands moved to IBM i	137
6.4	IBM Performance Tools for i5/OS enhancements	138
6.4.1	Print System Report (PRTSYSRPT)	138
6.4.2	Wait buckets	138
6.5	IBM iDoctor for IBM i	138
6.5.1	Name changes	139
6.5.2	New components	139
6.5.3	Graphical user interface changes in IBM iDoctor	139
6.5.4	IBM iDoctor for IBM i Heap Analyzer changes	141
6.5.5	IBM iDoctor for IBM i Job Watcher changes	142
6.5.6	IBM iDoctor for IBM i PEX-Analyzer changes	144
6.6	IBM Systems Director Navigator for i	144
6.6.1	Collections	145
6.6.2	Investigate Data - Collection Services	148
6.7	Work with System Activity (WRKSYSACT)	180
6.8	Performance Explorer	180
6.8.1	Job name (JOB)	181
6.8.2	Maximum storage to use (MAXSTG)	181
6.8.3	Program events (PGMEVT)	181
6.8.4	Base events (BASEVT)	181
6.8.5	Lock events (LCKEVT)	181
6.8.6	Operating system events (OSEVT)	181
6.8.7	Synchronization events (SYNCEVT)	181
6.8.8	Virtual I/O events (VRTIOEVT)	181
6.8.9	Program bracketing events (PGMBKTEVT)	181
6.8.10	Convert Performance Collection (CVTPFCOL)	182
6.9	Collection Services	182
6.9.1	Create Performance Data (CRTPFDRDTA)	182
6.9.2	Configure Perf Collection (CFGPFRCOL)	183
6.10	Performance Management Agent	184

6.11 System i Navigator On Demand Performance Center enhancements	184
6.12 Workload Estimator.	185
6.12.1 Version availability	185
6.12.2 New and enhanced IBM Power system models	186
6.12.3 New System x models.	186
6.12.4 Energy estimation	186
6.12.5 Disk storage system sizing support for DS5000 and XIV	186
6.12.6 License and Web standard	186
6.12.7 Import customer performance data.	187
6.12.8 Workload Developer Toolkit enhancements	187
6.12.9 Workloads and sizing guides	187
6.12.10 Support for IBM i and Linux on Power 595 (9119-FHA)	187
6.12.11 New System x models.	187
6.12.12 Sizing capability for disk storage systems.	187
6.12.13 WebSphere Process Server and WebSphere Enterprise Service Bus.	188
6.12.14 Performance Management interface to Workload Estimator	188
6.12.15 System Planning Tool linkage.	188
6.12.16 Ability to tailor the PDF contents.	188
6.12.17 Ability to limit sizing recommendations to POWER6 models	188
6.12.18 Workload Developer toolkit enhancements.	188
6.12.19 New IBM Power System servers.	189
6.12.20 Converged sizing enablement for System i and System p	189
6.12.21 New POWER6 blade for IBM BladeCenter JS12 Express	189
6.12.22 New IBM System x models.	189
6.12.23 New Windows Server 2008 support	189
6.12.24 New disk attachment types, new SAS storage	189
6.12.25 Updates to POWER6 virtualization for IBM Power Systems.	190
6.12.26 Increased session time-out	190
6.12.27 Performance management for Power Systems	190
6.12.28 Workload Developer enhancements.	190
6.12.29 IBM Sizing Guides for Power Systems, System i, System p, System x	190
6.12.30 Workload updates	190
6.12.31 New IBM System p POWER6 models	191
6.12.32 New IBM System x models	191
6.12.33 IBM i Support on the POWER6 blade.	191
6.12.34 New disk models.	191
6.12.35 IBM i 6.1 and Virtual I/O Server client support	191
6.12.36 Logical Partitioning and Virtual I/O Server algorithm updates.	191
6.12.37 Restore and quick validation.	192
6.12.38 Performance Management quality improvement	192
6.12.39 Workload Developer	192
6.12.40 IBM i Workloads	192
6.12.41 3Com Telephony workload	192
6.12.42 Disk adapters and I/Os per second ratings.	192
Chapter 7. Logical partitioning	195
7.1 PowerVM.	196
7.1.1 Hypervisor.	197
7.1.2 Shared processor pools	197
7.1.3 Multiple Shared-Processor Pools	197
7.1.4 Dedicated processor sharing	201
7.1.5 Dynamic logical partitioning	201
7.1.6 Virtual I/O Server	202

7.1.7	Integrated Virtualization Manager	203
7.2	IBM i 6.1 host and client partitions	207
Chapter 8. Networking		217
8.1	TCP/IP	218
8.1.1	IPv6 support enhancements	218
8.1.2	TCP/IP configuration enhancements	218
8.1.3	IBM i command enhancements	219
8.1.4	System i Navigator enhancements	219
8.1.5	Systems Director Navigator for IBM i Network interfaces	220
8.2	Virtual IP enhancements	222
8.3	Open Shortest Path First routing protocol	223
8.4	Domain Name System	224
8.4.1	Berkeley Internet Name Domain version 9	224
8.4.2	New configuration commands	224
8.4.3	New query and update utilities	225
8.5	Netserver	226
8.6	Virtual Private Network	227
8.7	FTP client and server extended FTP commands	227
8.8	Mail support enhancements	228
8.8.1	Simple Mail Transfer Protocol (SMTP)	228
8.8.2	Post Office Protocol (POP)	230
8.9	IBM Tivoli Directory Server for i5/OS	230
8.9.1	Replication conflict resolution	230
8.9.2	The ldapmodify command	230
8.9.3	Multi-threaded replication	230
8.9.4	Password encryption	230
8.9.5	The IBMAttributeTypes attribute	230
8.9.6	Schema changes	231
8.9.7	Distributed directory	231
8.9.8	ldapmodrdn	231
8.9.9	Using Trace TCP/IP Application to help find problems	231
8.9.10	Read access to projected users	231
8.9.11	Multiple server instances	231
8.9.12	Password policy	231
8.9.13	Directory Server equivalence	231
8.9.14	Migrating from IBM i 5.4 or i5/OS V5R3 to IBM i 6.1	231
8.9.15	Migrating data from a version prior to i5/OS V5R3 to IBM i 6.1	232
8.10	Advanced program-to-program communications	233
Chapter 9. Printing		235
9.1	Printing overview	236
9.2	Dates and times in Coordinated Universal Time	237
9.2.1	Data queue support on output queues	238
9.2.2	Environment variable data queue support	238
9.3	Adding a printer with System i Navigator	239
9.4	Generating PDF files with IBM i	245
9.4.1	New print transform support	245
9.4.2	New print transform prerequisites	246
9.4.3	Workstation Customizing object (WSCST) parameter	246
9.4.4	Customizing WSCST object of *CTXFORM transform type	246
9.4.5	Create printer files	247
9.4.6	PDF output considerations	249

9.5 Start Spool Reclaim command	251
9.6 Updated open feedback area	251
Chapter 10. Integration with BladeCenter and System x	255
10.1 IBM i Support for Power Blades	256
10.1.1 BladeCenter JS12, JS22, JS23, and JS43	256
10.1.2 IBM client to VIOS on supported BladeCenter systems	256
10.1.3 IBM BladeCenter JS12	256
10.1.4 IBM BladeCenter JS22	256
10.1.5 IBM BladeCenter JS23 Express	257
10.1.6 10.1.6. IBM BladeCenter JS43 Express	257
10.1.7 Overview of backup and restore for IBM i on Power blade	257
10.2 Changes to iSCSI-attached integrated servers	258
10.2.1 Configuring hardware resource names	258
10.2.2 Shared memory pools	259
10.2.3 Challenge Handshake Authentication Protocol	259
10.2.4 iSCSI direct connect	259
10.3 Changes to integrated Windows servers	259
10.3.1 Backing up active servers	260
10.4 VMware ESX Server support	261
10.5 Changes to integrated Linux servers	261
10.5.1 Supported versions of Linux	261
10.5.2 File-level backup	261
10.5.3 Backing up active Linux servers	262
10.5.4 Time synchronization	262
10.6 IXS and IXA-attached integrated Windows servers	262
10.6.1 Disabling user profiles	262
10.6.2 Time synchronization	262
10.6.3 Windows Server installation advisor	263
10.7 Linux running on IXS and IXA-attached hardware	263
Chapter 11. IBM Advanced Job Scheduler for i5/OS	265
11.1 IBM Advanced Job Scheduler for i5/OS	266
11.2 Advanced Job Scheduler Web interface	267
11.3 Error notification	271
11.4 Output queue monitoring	271
11.5 Command lists	272
11.6 Multiple scheduling environments	273
11.7 Report distribution	274
11.8 Log	274
11.9 Scheduling and running jobs	275
11.9.1 Scheduled jobs properties	275
11.9.2 Day of week option for subordinate jobs	276
11.9.3 Start Group using Job Scheduler (STRGRPJS)	276
11.9.4 Print Schedule using Job Scheduler (PRTSCDJS)	276
Chapter 12. Web serving	277
12.1 Apache Software Foundation Tomcat servlet engine	278
12.2 Large file support added	278
12.3 IPv6 support	278
12.4 New module mod_proxy_balancer	279
12.5 New module mod_filter	280
12.6 Integrated Web application server	280
12.7 New Web services server	282

Chapter 13. Web enablers	283
13.1 Integration with the Integrated Web application server	284
13.2 Web Performance Management tools	285
13.2.1 Web Performance Monitor	287
13.2.2 Web Performance Advisor	292
13.3 WebSphere Application Server Express V7.0	305
13.3.1 Standards	305
13.3.2 Configuration options	310
13.3.3 Ordering information	311
13.4 WebSphere Application Server performance data	312
13.5 Simplified use of SSL with HTTP Administration server	312
Chapter 14. Application development	317
14.1 IBM Developer Kit for Java	318
14.1.1 Release-to-release compatibility	318
14.1.2 Changes to adopted authority support	318
14.1.3 Java Cryptography Extension hardware	319
14.1.4 Java Virtual Machine Tool Interface	319
14.1.5 Java Secure Socket Extension	320
14.1.6 New Java CL commands	320
14.1.7 JDBC 4.0	321
14.1.8 Support for multiple IBM Developer Kit for Java options	322
14.1.9 Installing IBM Technology for Java Virtual Machine	324
14.1.10 List of Java system properties	325
14.1.11 Java Runtime.exec	326
14.1.12 Miscellaneous changes	327
14.2 IBM Toolbox for Java	327
14.2.1 Using IBM i enhancements in the IBM Toolbox for Java JDBC	328
14.2.2 New classes	328
14.2.3 Enhanced classes	330
14.2.4 Compatibility	330
14.3 Websphere Development Studio	331
14.3.1 ILE RPG compiler	333
14.3.2 ILE COBOL	336
14.3.3 ILE C and ILE C++	337
14.4 Application Development ToolSet	338
14.5 Debugger	338
14.6 Portable Application Solutions Environment	339
14.6.1 Derivation of PASE	339
14.6.2 Compilers	339
14.6.3 Snapcore	339
14.6.4 New or changed PASE runtime functions	340
14.7 Qshell and utilities	340
14.7.1 Qshell	341
14.8 XML toolkit	341
14.8.1 New version of XML Toolkit	341
14.9 Control language	342
14.9.1 New commands	342
14.9.2 Command enhancements	343
14.10 Structure Query Language	345
14.10.1 Call Level Interface	345
14.10.2 SQL programming	346
14.10.3 Embedded SQL programming	348

14.10.4	SQL Messages and codes	348
14.10.5	Grouping and On-Line Analytical Processing	349
14.10.6	SQL reference.	350
14.10.7	SQL troubleshooting	352
14.10.8	SQL Query Engine enhancements	353
14.11	Application programming interfaces	354
14.11.1	API changes	354
14.11.2	Cluster.	354
14.11.3	Problem Management.	355
14.11.4	Security	355
14.11.5	Data conversion APIs	355
14.11.6	Miscellaneous	356
14.11.7	API documentation updates in the Information Center	356
14.11.8	Machine Interface instructions	357
14.11.9	Cryptographic services	358
14.12	Globalization	359
14.13	Multithreaded applications	359
Chapter 15.	Rational products for IBM i	361
15.1	Rational Developer for i for SOA Construction	362
15.1.1	Comparison between RDi and RDi SOA.	362
15.2	Rational Developer for i.	364
15.2.1	Remote System Explorer	364
15.2.2	Application Diagram	365
15.2.3	Screen Designer	365
15.2.4	Screen Designer technology preview phase 2	366
15.2.5	Debug	366
15.2.6	IBM i Web services and Java tools	366
15.2.7	Integration with IBM Rational Application Developer	366
15.2.8	Change management	367
15.2.9	Other WebSphere Development Studio Client functions	367
15.2.10	Improved performance	367
15.2.11	Reduced resource requirements.	367
15.3	Using Rational Developer for i	367
15.3.1	Remote System Explorer	368
15.3.2	i5/OS Projects	372
15.3.3	Application Diagram	374
15.3.4	Screen Designer	376
15.3.5	Integrated IBM i Debugger	377
15.3.6	Web Services and Java Tools for IBM i	378
15.4	Rational Business Developer	379
15.4.1	Rational Business Developer enhancements	381
15.4.2	Enterprise Generation Language enhancements	381
15.5	Using Enterprise Generation Language	382
15.6	Rational Host Access Transformation Services for 5250	394
15.6.1	WebFacing and HATS	395
15.6.2	HATS interaction with WebFacing	395
15.6.3	HATS V7.5 enhancements	398
15.6.4	WebFacing V7.1 enhancements	400
15.7	Rational Team Concert for i	400
15.8	WebSphere Development Studio Client for System i	401
Chapter 16.	IBM Systems Director Navigator for i	403

16.1 IBM Systems Director Navigator for i	404
16.2 Systems Director Navigator for i and Systems i Navigator	405
16.2.1 Features included in Systems Director Navigator for i	406
16.2.2 Features not included in Systems Director Navigator for i	406
16.3 Working with IBM Systems Director Navigator for i	407
16.3.1 i5/OS Tasks	412
16.3.2 Exploring IBM Systems Director Navigator for i	412
Chapter 17. System i Access products	429
17.1 Windows	430
17.1.1 .NET Data Provider	430
17.1.2 OLE DB	430
17.1.3 Basic operations	431
17.1.4 Work management	431
17.1.5 Network	432
17.1.6 TCP/IP servers	432
17.1.7 Integrated server administration	434
17.1.8 Security	434
17.1.9 Users and groups	435
17.1.10 Structured query language	435
17.1.11 Integrated file system	436
17.1.12 Network File System	437
17.1.13 Backup Recovery and Media Services	438
17.1.14 Management Central	439
17.1.15 Configuration and service: Time management and time zone	440
17.1.16 Configuration and service: Hardware	440
17.1.17 Installation and setup	441
17.2 Linux	442
17.3 System i Access for Web	442
17.3.1 Database function	443
17.3.2 New portal environment support	445
17.4 Wireless	446
17.4.1 Version and release changes in the ODBC driver behavior	446
17.4.2 Developing System i Navigator plug-ins	446
17.4.3 Tasks on the Web	447
Chapter 18. Miscellaneous IBM i 6.1 changes	449
18.1 System values	450
18.1.1 Security auditing level	450
18.1.2 Additional active jobs	452
18.1.3 Additional total jobs	452
18.1.4 Active jobs	452
18.1.5 Total jobs	452
18.1.6 Uninterruptible power supply delay time	452
18.1.7 Security related system values	452
18.2 Integrated file system	452
18.2.1 Collecting and analyzing folder attributes with System i Navigator	452
18.2.2 Enhanced check-in and check-out functions	452
18.2.3 Automatic conversion to support additional characters in names	453
18.2.4 Support for IPv6	453
18.2.5 Unicode enablement of some CL commands	453
18.2.6 Dropping support for the QNetWare file system	453
18.3 Virtual optical and virtual tape support	453

18.3.1	Image catalog entry for virtual tape	453
18.3.2	Independent ASP support for virtual tape	453
18.3.3	Size ranges for image files	454
18.4	Service tools language	454
18.5	Watch for event function	454
18.6	Power Down System command change	455
18.7	Maximum capacities	455
18.8	Quiesce I/O for an Auxiliary Storage Pool	455
18.9	Time management	456
18.9.1	New time zone values	458
18.9.2	New time zones	458
18.10	Operations Console	459
18.10.1	Simplified console setup for network connected consoles	459
18.10.2	Auto-create service tools device IDs	459
18.10.3	Automatic management of the access password	460
18.10.4	Automatic discovery of systems	460
18.10.5	Planning your Operations Console installation or upgrade	461
18.10.6	Migrating to Operations Console before your server model upgrade	461
18.10.7	New models	461
18.10.8	Switching console control to another user	462
18.11	Models not supported	462
18.12	Common Information Model	462
Chapter 19. Upgrading IBM i	465
19.1	Installing or upgrading IBM i	466
19.1.1	Models supported	466
19.1.2	Planning statements	466
19.1.3	Program Temporary Fixes for upgrading to IBM i 6.1	466
19.1.4	Pre-upgrade verification tool	466
19.1.5	Object conversions	466
19.1.6	Spooled file conversions	467
19.1.7	Licensed program releases and sizes	467
19.1.8	Media labels and their contents	468
19.1.9	Optical media for IBM i 6.1	468
19.1.10	Virtual tape devices for release upgrades	468
19.1.11	Installing using virtual optical devices and remote virtual image files	468
19.1.12	Sharing optical resources	468
19.1.13	Electronic software delivery	468
19.1.14	Software keys available online	469
19.2	Upgrading from i5/OS V5R3 to IBM i 6.1	469
19.2.1	Requirements before you install IBM i 6.1 Licensed Internal Code	469
19.2.2	Server firmware	469
19.2.3	Operations Console	470
19.2.4	Common Information Model	470
19.3	Distributing software	470
19.3.1	Using virtual tape drives	470
19.3.2	Sharing optical resources	470
19.3.3	Duplicating virtual images to physical media	470
19.3.4	Using image catalogues	470
19.4	Program Temporary Fix enhancements	470
19.4.1	Cumulative Program Temporary Fix media	471
19.4.2	Multiple level support for groups	471
19.4.3	Group status	471

19.4.4 Target release field	471
19.4.5 Work with Program Temporary Fix Groups command	471
19.4.6 Compare and Update wizard	472
19.4.7 Display Program Temporary Fix command.	472
19.4.8 Program Temporary Fix orders.	472
Chapter 20. IVM and HMC feature summary.	473
20.1 IVM and HMC comparison at a glance	474
20.2 IVM and HMC feature summary	475
Related publications	479
IBM Redbooks	479
Online resources	479
How to get Redbooks.	481
Help from IBM	481
Index	483

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 on 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 product(s) and/or the program(s) 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

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

AFS®	i5/OS®	Print Services Facility™
AIX 5L™	IBM®	Rational Team Concert™
AIX®	iCluster®	Rational®
alphaWorks®	IMS™	Redbooks®
AS/400®	Integrated Language Environment®	Redpaper™
BladeCenter®	iSeries®	Redbooks (logo)  ®
ClearCase®	Jazz™	RPG/400®
DataPower®	Language Environment®	Sametime®
DB2 Universal Database™	Lotus®	Solid®
DB2®	Micro-Partitioning™	System i5®
Distributed Relational Database Architecture™	OmniFind®	System i®
Domino®	OpenPower®	System p5®
DRDA®	OS/400®	System p®
DS4000®	Passport Advantage®	System Storage™
DS6000™	POWER Hypervisor™	System x®
DS8000®	Power Systems™	System z®
EnergyScale™	POWER5™	Tivoli®
Enterprise Workload Manager™	POWER5+™	VisualAge®
eServer™	POWER6®	WebSphere®
FlashCopy®	PowerHA™	XIV®
Focal Point™	PowerVM™	xSeries®
	POWER®	

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Adobe, and Portable Document Format (PDF) are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Disk Magic, and the IntelliMagic logo are trademarks of IntelliMagic BV in the United States, other countries, or both.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Interchange, Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

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

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows NT, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, 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 and other countries.

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.

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

Preface

This IBM® Redbooks® publication introduces a technical overview of the main new features, functions and enhancements available in IBM i 6.1 (formerly called i5/OS® V6R1). It gives a summary and brief explanation of new capabilities and what has changed in the operating system, and also discusses many of the licensed programs and application development tools associated with IBM i.

Many other new and enhanced functions are described, such as virtualization of storage, security, Java™ performance, improved performance with IBM System Storage™ devices, backup and recovery, including base IBM i, Backup, Recovery and Media Services (BRMS). The book introduces the PowerHA™ product, IBM Systems Director-based system management and an easier Web enablement.

The information provided in this book will be useful for customers, Business Partners, and IBM service professionals involved with planning, supporting, upgrading, and implementing IBM i 6.1 solutions.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Jim Cook was a Senior Software Engineer Project Leader at the ITSO, Rochester Center, and was the original project leader for this publication. He led teams that produced IBM System i® announcement presentation sets that are maintained on the System i Technical Support Web sites, and presented at ITSO Forums internationally. Jim also produced IBM Redbooks publications about various System i, and then System p® based hardware, and IBM i (i5/OS) related topics. Jim is now happily retired. He expresses his most sincere thanks to the IBM i developers and the authors on the team that created this book.

Ahmed Mansour is a software engineer at the Cairo Technology Development Center, IBM Lab in Egypt. He holds a Bachelor's degree in Computer Science. He has been with IBM since 2006. Ahmed is a member of the IBM i/Rational® team, focusing on providing Arabic bidirectional scripts and globalization support for IBM products and platforms. He currently works on the development and support of DB2® Text Extender and CommonStore for SAP on IBM i. You can reach Ahmed at amansour@eg.ibm.com.

Giancarlo Omati is an IT Specialist for System i at the Technical Support in IBM Italy. He has 12 years of experience in networking, communications (TCP/IP and SNA) and the High Availability field. His areas of expertise include problem determination, planning and implementation of HA software solutions at customer sites. He holds a diploma in Information Technology.

Ana Flavia Pavan has been working at IBM Brazil for five years as IT Specialist for IBM System i. She has 10 years of experience in technical support on IBM System i and 14 years of experience in the computing field. She has taught IBM System i classes at IBM Brazil and also has contributed to the content of IBM System i EAD (eLearning) Oficina do Futuro offered by IBM Brazil and education partners, for instructing future IBM operators. She holds a Bachelor of Computer Science degree from Universidade Municipal de Sao Caetano do Sul - USCS (Brazil). She holds a postgraduate degree in computer networks from Universidade

Estadual de Campinas - UNICAMP (Brazil) and a postgraduate degree in Internet consulting from Faculdades Associadas de Sao Paulo - FASP (Brazil).

Mark Roy is an IBM i specialist based in Melbourne, Australia. He was previously with a large Australian bank, designing, installing and enhancing their IBM i (a.k.a. CPF, OS/400® and i5/OS) environment. He was responsible for the automation and global consolidation of their infrastructure, implementing robust disaster recovery and high availability solutions to support the centralized environment. Mark is an author of an IBM Redbooks publication on IBM i diagnostic tools and was a member of the IBM Rochester Performance Advisory Council. He recently established Sysarb, a company providing freelance consulting and contracting services to IBM i customers and service providers. He specializes in IBM i infrastructure architecture, technical support, performance, and systems management. Mark can be contacted at Mark.Roy@sysarb.com.au.

The project that produced this publication was managed by:

Scott Vetter, PMP
IBM Austin

Thanks to the following people for their contributions to this project:

Lindsay S Avers, Hernando Bedoya, David Bhaskaran, Adam T Clark, Jenny Dervin, Robert Dick, Bruno Digiovani, Scott Elliott, Livio Teixeira Filho, Arv Fisher, Larry Hall, Kristi Harney, Terry Hennessy, David C Hinkemeyer, Brian Jongekryg, Gary Konicek, Octavian Lascu, Dave Legler, Peg Levering, Dawn May, Greg McIntire, Jeff Meaden, Kent Milligan, Timothy D Mullenbach, Vess Natchev, Christy L Norman, Bill Schmidt, Oliver Stadler, MJ Thompson, Randy Tomfohrde, Scott Urness, Kevin Vette, Steve Will

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will 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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

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

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:
redbooks@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



IBM i 6.1 introduction and summary

This chapter summarizes new functions and enhancements available in IBM i 6.1 (formerly called i5/OS V6R1). It highlights new capabilities in the operating system and also many of the licensed programs and application development tools associated with IBM i.

The following chapters provide additional overview information. As a technical overview publication, detailed “how to do” information is beyond the scope of this publication, which consolidates into one place most of the information about IBM i 6.1 and related new or enhanced capabilities. For additional details, reference other sources such as:

- ▶ The IBM i Information Center for IBM i 6.1 Web site at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>
- ▶ The *i5/OS V6R1 Technical Overview* presentation set, located in the Support for IBM i Web site at:
<http://www.ibm.com/systems/support/i/news/v6r1techoverview.html>

These presentations contain some details not included in this IBM Redbooks publication but do not include capabilities or information that became available during the second half of 2008.

If you are currently running on an IBM i release level older than 5.4, and wish to upgrade to IBM i 6.1, we suggest you also use these sources:

- ▶ The IBM i Information Center for IBM i 6.1 Web site at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/index.jsp>
- ▶ The IBM Redbooks publication *IBM System i5 V5R4 Technical Overview*, SG24-72711, which can be found at:
<http://www.ibm.com/redbooks>

1.1 Overview of what is new in IBM i 6.1

In January 2008 IBM announced a new major release of the i5/OS operating system - IBM i 6.1. During 2008 i5/OS was renamed to IBM i and the notations 5.4 and 6.1, to indicate operating system release levels V5R4 and V6R1, were introduced. Some user documentation, Web page links, and programmed interfaces use IBM i terminology while others still use the i5/OS nomenclature. In this publication we use IBM i terminology but occasionally also use i5/OS, typically where it is part of a product name or appears on a panel.

IBM i 6.1 and licensed programs associated with it contain a very wide range of new functions and enhancements. In the area of IBM WebSphere® application development tools, enhancements are also repackaged into tools identified with the Rational brand name. For example, the “new next release” of WebSphere Development Tools for iSeries® Client (WDSC) is named Rational Developer for i.

From a functional and performance viewpoint, IBM i 6.1 should be installed and utilized as soon as possible, especially if you are on a release prior to IBM i 5.4.

IBM i 6.1 runs on all POWER5™ and POWER6® technology systems, as well as previous technology 9406 models 800, 810, 825, 870, and 890.

IBM i 5.4 runs on all POWER5 and POWER6 technology systems, as well as previous technology 9406 models 800, 810, 820, 825, 830, 840, 870, 890, SB2 and SB3.

This publication is focused on IBM i and associated software. For detailed information about hardware, refer to the system builder IBM Redbooks, which can be found by searching for *IBM i system builder* on the IBM Redbooks site at:

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

The key new and enhanced functions in IBM i 6.1 are in the following areas:

- ▶ Security, with expanded base security options and greatly expanded network intrusion detection options.
- ▶ Encryption of data during disk and tape backups, including management under the Backup Recovery and Media Services product.
- ▶ Backup, recovery and system availability, including more save-while-active options, and journaling and clustering options that enable more flexible save and restore operations and high availability system configurations.
- ▶ Enhanced backup and recovery products, including Backup Recovery and Media Services, and new with IBM i 6.1, the IBM High Availability Solutions Manager (HASM) product. HASM was rebranded during 2008 as PowerHA for i.
- ▶ Virtualization of IBM i storage, including IBM i as a client using virtual I/O, hosted by either another IBM i partition or a Virtual I/O Server partition. Client IBM i virtualization also enables IBM i to run on IBM POWER6 BladeCenter® servers.
- ▶ Expanded number of BladeCenter servers under IBM i System x® and BladeCenter integration support.
- ▶ IBM DB2 for i functional and performance enhancements, along with lower cost usage options of the DB2 Web Query for i product.
- ▶ Enhanced Java performance under the IBM Technology for Java 32-bit (introduced with IBM i 5.4) and 64-bit (new in IBM i 6.1) JVMs and new JDK levels support.

- ▶ Expanded Web enablement made easier capabilities. This includes the following support with the objective of making it easier to get started running Web applications under IBM i:
 - Integrated Web Applications Server for i (introduced with IBM i 5.4), for applications with less sophisticated Java requirements than those available through the robust capabilities of IBM WebSphere Application Server servers. The integrated Web application server requires less system resources than a corresponding Web Application Server instance.
 - Integrated Web Services for i server, available to less sophisticated applications using Web services functions, for those wanting to get started with Web services.
- ▶ Expanded application development tool capabilities for IBM i applications. These include:
 - Enhancements to CL, C++, RPG, and COBOL, such as more flexible source statement importing and RPG multiple thread support and other enhancements with a focus on easier enabling of applications in a Web services environment.
 - Repackaged and easier to install and use workstation-based application development tool products, such as Rational Development for i (RDi) and Rational Developer for i for SOA Construction (RDi SOA). This is the “next release” replacement for the familiar WebSphere Development Studio for iSeries client (WDSC) development tool.
- ▶ An new Web-based system management interface to IBM i, called IBM Systems Director Navigator for i. This is a browser-based interface to many of the same functions available using the System i Navigator interface from a Windows® client workstation. The IBM Systems Director Navigator for i also offers new graphical analysis of collected IBM i performance data.

This interface to a single IBM i partition is the same technology as that used by the multiplatform suite of IBM Systems Director products, for example, IBM Systems Director 6.1. IBM Systems Director 6.1 is available on IBM servers running Linux® on x86, Linux on System z®, Linux on Power, AIX®, as well as Windows operating environments, and can manage other operating systems, such as IBM i.

IBM Systems Director Navigator for i can be accessed from IBM Systems Director 6.1.

Note: Some readers may be confused by the significance of the version numbers of IBM 6.1 and IBM Systems Director 6.1. There is no relationship between the version numbers, it is coincidental.

IBM i 6.1 or 5.4 can function as an IBM Systems Director 6.1 *agent* but not as an IBM Systems Director 6.1 *server*. IBM Systems Director 5.2 (predecessor to IBM Systems Director 6.1) does run as server or agent under IBM i 5.4 and 6.1.

This publication does not contain information about IBM Systems Director 6.1. For more information about IBM Systems Director 6.1, refer to the following:

- ▶ IBM Systems Software Information Center at:
 - <http://publib.boulder.ibm.com/infocenter/systems/index.jsp>
 - Select System Software and expand *IBM Director 5.n* or *IBM Systems Director V6.1*.
- ▶ Announcement letter 208-280, *IBM Systems Director for Power Systems V6.1*.
- ▶ The IBM Redbooks publication *Implementing IBM Systems Director 6.1*, SG24-76944.

Note: Information Center for IBM i 6.1 and 5.4 contains significant overview and more detailed level information and includes IBM i 5.4 and 6.1 versions of the Memo to Users. These publications document key differences between IBM i 5.4 and i5/OS V5R3, and IBM i 6.1 and IBM i 5.4. Each of the topics in the Memo to Users provides important compatibility information that you should become familiar with before you install the new release. Some of these topics have a larger impact on more customers than others, and you should review these topics carefully. A list of these topics can be found in the section titled “Key messages in this document”.

Use this URL:

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

and select a release to find this information,



Security enhancements

This chapter discusses the following topics:

- ▶ “Intrusion detection and prevention” on page 6
- ▶ “Encryption of data at rest on disk” on page 11
- ▶ “Encrypted backups” on page 12
- ▶ “New security system values” on page 13
- ▶ “Changed security system values” on page 18
- ▶ “Restricting changes to security settings” on page 19
- ▶ “User profile enhancements” on page 19
- ▶ “Service Tools user profile and IBM i profile” on page 21
- ▶ “Object signing and signature verification” on page 23
- ▶ “Digital Certificate Manager” on page 24
- ▶ “Single sign-on” on page 24
- ▶ “Enterprise Identity Mapping” on page 24
- ▶ “Network Authentication Service” on page 25
- ▶ “Cryptographic enhancements” on page 26

2.1 Intrusion detection and prevention

An intrusion detection and prevention system (IDS) notifies you of attempts to hack into, disrupt, or deny service to the system. IDS also monitors for potential extrusions, where your system might be used as the source of the attack. These potential intrusions and extrusions are logged as intrusion monitor audit records in the security audit journal and displayed as intrusion events in the Intrusion Detection System graphical user interface (GUI). You can configure IDS to prevent intrusions and extrusions from occurring.

Important: The term intrusion detection is used two ways in IBM i documentation. In the first sense, intrusion detection refers to the prevention and detection of security exposures. For example, a hacker might be trying to break into the system using a user ID that is not valid, or an inexperienced user with too much authority might be altering important objects in system libraries. In the second sense, intrusion detection refers to the intrusion detection function that uses policies to monitor suspicious traffic on the system.

Intrusion detection in IBM i 6.1 now does the following:

- ▶ Provides a built-in prevention method with variable dynamic throttling.
- ▶ Provides real time notification (e-mail, pagers).
- ▶ Provides improved auditing with additional intrusion events being audited.
- ▶ Recognizes well-known attacks such as “Smurf”, “Fraggle”, ACK storms, Address Poisoning (both IPv4 ARP poisoning, and IPv6 neighbor discovery poisoning), and Ping-Of-Death.
- ▶ Is no longer dependant on the QoS server.

To manage the intrusion detection environment we can use both the System i Navigator and the IBM Systems Director Navigator for i GUIs. With these, we can manage security policies and display intrusion events, as an alternative to viewing the audit journal.

2.1.1 Variable dynamic throttling

Variable dynamic throttling, provided in IBM i 6.1, is a prevention method that automatically starts if certain intrusion event thresholds are met. Throttling stays active until thresholds are no longer exceeded for an interval of time. You can choose to throttle network traffic from all or specific ports and IP addresses. You also can specify the slow and fast scan thresholds, the maximum event message thresholds, or use the default values for those thresholds in your IDS policies. Throttling is activated once a threshold for that policy has been exceeded, and stays active for either a user-defined or system-defined time interval. If the threshold is exceeded at any time during the interval, the throttling is increased immediately and the time interval is reset. Throttling could eventually lead to denying all packets from a given interface. This process continues until the number of offending packets no longer exceeds the thresholds for an entire time interval. When the number of packets drops below the thresholds, throttling is deactivated and normal packet flow resumes.

Consider using variable dynamic throttling when a specific intrusion detection policy is generating a lot of intrusion events to prevent overloading the system with what might be a denial-of-service (DoS) attack. You can configure variable dynamic throttling for each IDS policy. Throttling detects all types of intrusions and extrusions.

2.1.2 Setting up intrusion detection

The following figures show some of the steps to set up the new IDS capabilities using IBM Systems Director Navigator for i. The System i Navigator panels provide equivalent functions.

First, access the IBM Systems Director Navigator for i, expand the IBM i Management section, select **Security** and then **Intrusion Detection**. On the Intrusion Detection Management panel, you can now start to configure intrusion detection policies by accessing the link Manage Intrusion detection policies, as shown in Figure 2-1.

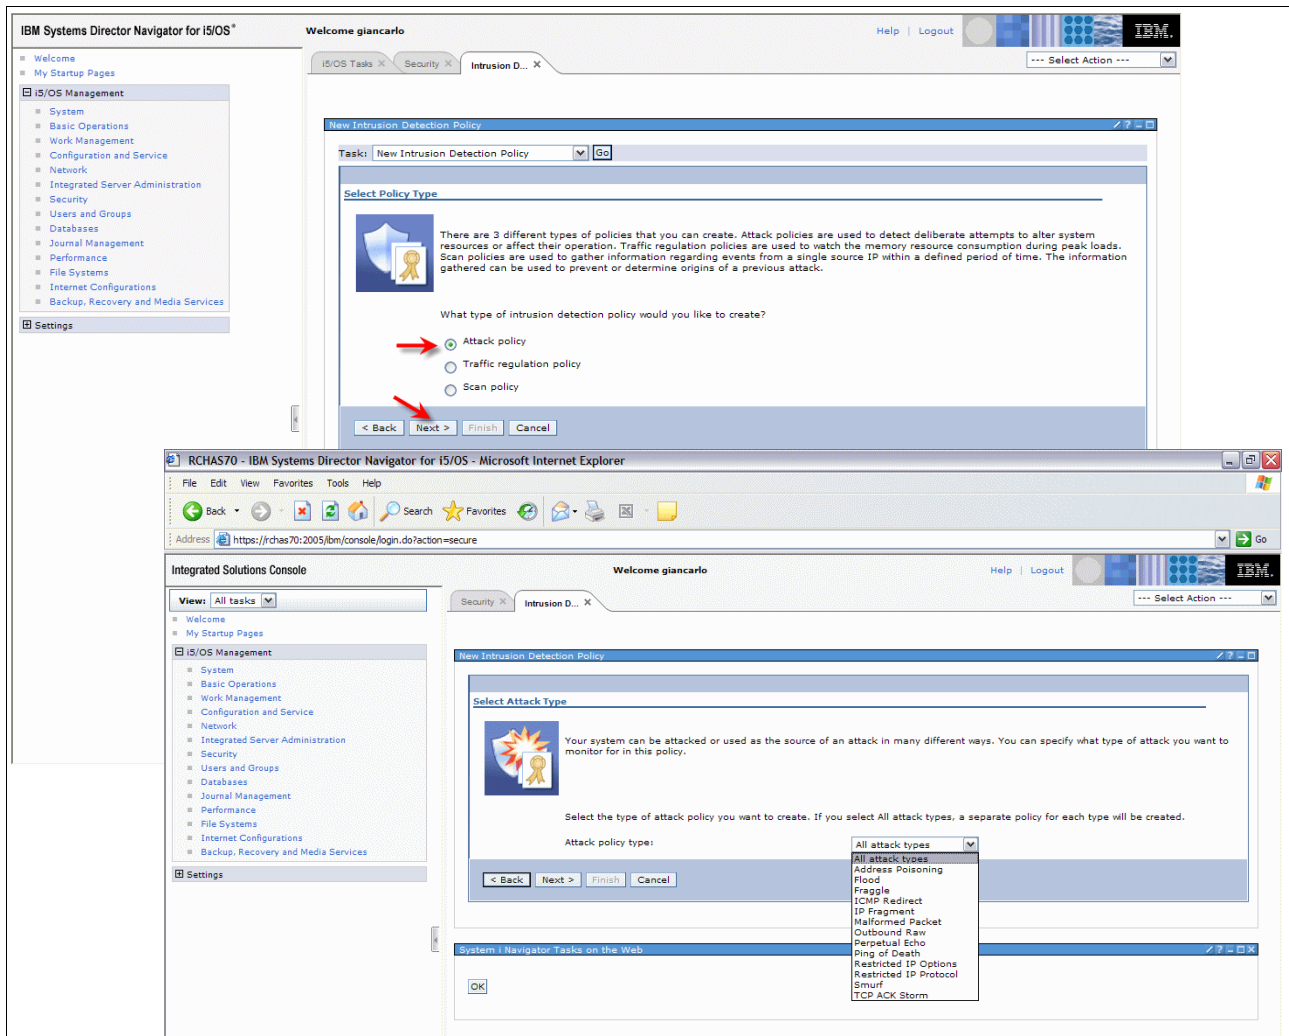


Figure 2-1 Starting to define an Attack Policy

Once you have activated your policies, you can view any associated intrusion events. As shown in Figure 2-2 on page 8, there are several possible actions you can select, including defining a new policy, selecting an individual event, filtering and sorting event information shown, specifying the time range of events to be viewed, and others.

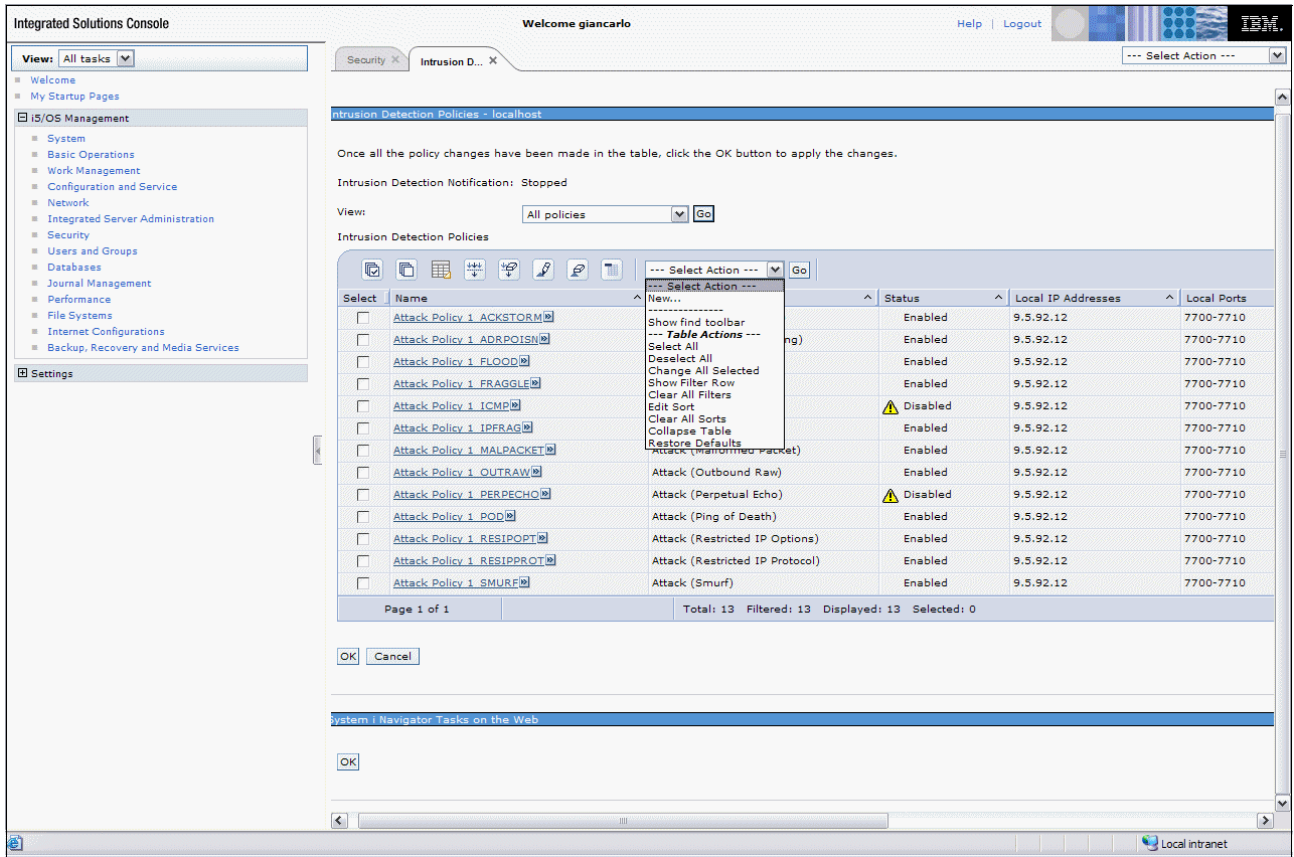


Figure 2-2 Manage intrusion detection policies

Figure 2-3 shows examples of the high-level summary information for each event shown after some filtering has been specified. You see Scan and Traffic regulation policy events.

The screenshot displays the 'Intrusion Detection Events - localhost' window. It features a table with the following data:

Select	Date and Time	Type	Direction	Local IP Address	Local Port
<input type="radio"/>	Oct 24, 2008 3:02:48 PM	Scan	Inbound	9.5.92.12	7701
<input type="radio"/>	Oct 24, 2008 3:02:50 PM	Scan	Inbound	9.5.92.12	7701
<input type="radio"/>	Oct 24, 2008 3:02:52 PM	Scan	Inbound	9.5.92.12	7701
<input type="radio"/>	Oct 24, 2008 3:02:54 PM	Scan	Inbound	9.5.92.12	7701
<input type="radio"/>	Oct 24, 2008 3:02:56 PM	Scan	Inbound	9.5.92.12	7701
<input type="radio"/>	Oct 24, 2008 3:04:10 PM	Traffic regulation (TCP)	Inbound	9.5.92.12	1456
<input type="radio"/>	Oct 24, 2008 3:04:10 PM	Traffic regulation (TCP)	Inbound	9.5.92.12	1457
<input type="radio"/>	Oct 24, 2008 3:04:11 PM	Traffic regulation (TCP)	Inbound	9.5.92.12	1458
<input type="radio"/>	Oct 24, 2008 3:04:16 PM	Traffic regulation (TCP)	Inbound	9.5.92.12	1459
<input type="radio"/>	Oct 24, 2008 3:04:16 PM	Traffic regulation (TCP)	Inbound	9.5.92.12	1460

Page 1 of 1 Total: 10 Filtered: 10 Displayed: 10 Selected: 0

Figure 2-3 Event summary

In Figure 2-4 we show the event details for the Scan event that occurred at 3:02:52 PM.

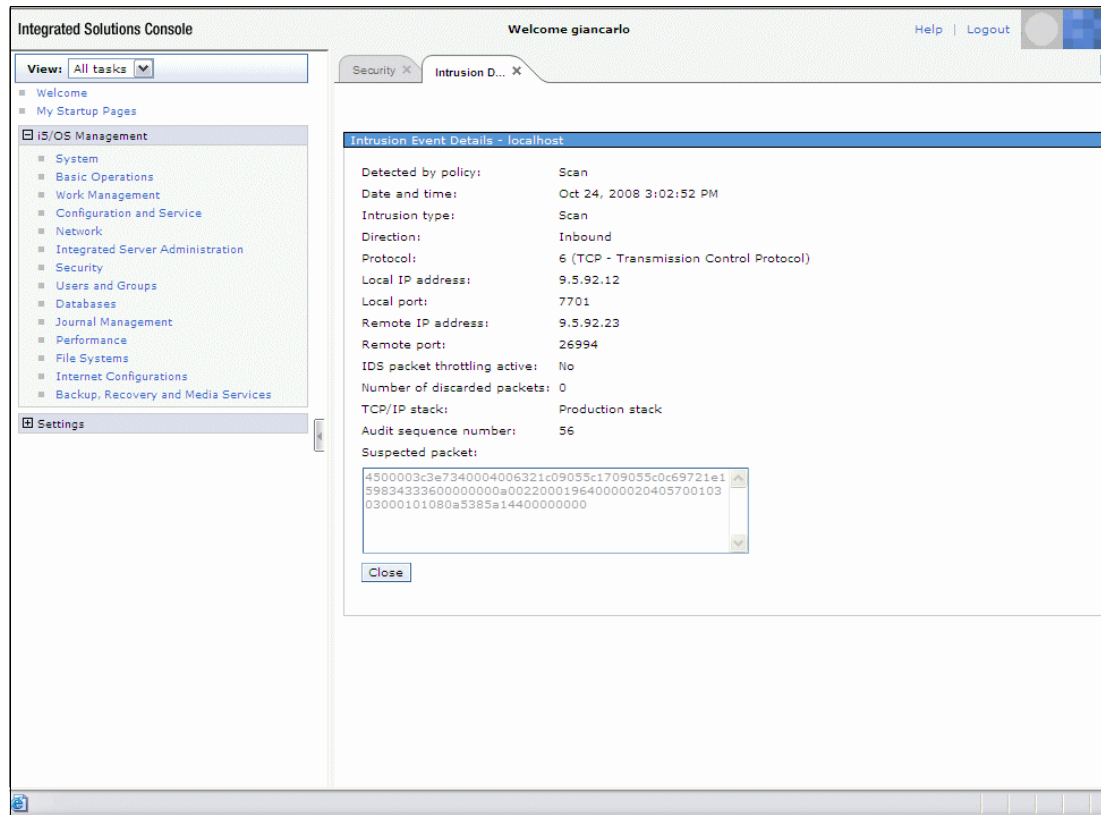


Figure 2-4 Event details

In the area identified as the suspected packet, an experienced TCP/IP person may be able to review the data to determine more details.

To learn more about intrusion detection and how to configure it, refer to the following information sources:

- ▶ The IBM i Information Center, selecting the path **Security** → **Intrusion Detection**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680

2.1.3 Remove TCP/IP Table (RMVTCPTBL)

Support has also been added to the Remove TCP/IP Table (RMVTCPTBL) command for those times when you absolutely must deactivate IDS in case you have locked yourself out from your Web server port. To do this, use the Remove TCP/IP Table (RMVTCPTBL) command, selecting *IDS in the table parameter as shown in Figure 2-5.

```
Remove TCP/IP Table (RMVTCPTBL)

Type choices, press Enter.

Table . . . . . > *IDS          *ALLRULE, *ALL, *IPFTR...
```

Figure 2-5 Remove TCP/IP Table command

2.2 Encryption of data at rest on disk

Disk encryption allows you to encrypt data stored in basic (dependent) disk pools, also known as Auxiliary Storage Pools (ASPs). You can also have encrypted data in an independent disk pool (Independent ASP). Data in the system ASP cannot be encrypted.

Encryption is specified only on the create ASP (disk pool) interface. All data placed onto the disks within an encrypted disk pool is encrypted. Once a disk pool has been created as encrypted, it cannot be changed to not use encryption. An existing un-encrypted disk pool cannot be changed to be encrypted.

Disk encryption protects data from a number of different threats:

- ▶ Protects data transmission to and from the disk drive (important in a SAN environment).
- ▶ Protects data transmission in a cross-site mirroring environment (only when the data being mirrored is on an encrypted independent disk pool).
- ▶ Protects data in the case of theft of the disk drive.
- ▶ Protects data in the case of return or resale of a disk drive (reduces the need to sanitize the disk drive).

Using disk pool encryption may reduce the need for application-implemented data encryption.

Other characteristics of IBM i disk pool encryption are:

- ▶ Uses the Advanced Encryption Standard (AES) symmetric key algorithm. IBM i internally handles encryption master keys. You do not need to create and manage a master key for this function with auxiliary storage pools. With regards to an independent ASP, the data keys are managed automatically but stored on the ASP and encrypted under the *ASP master key that was added in IBM i 6.1. The *ASP master key needs to be set.
- ▶ Cryptographic keys, including master keys, are generated internally by IBM i. They are stored in software but protected by “isolated” storage and master keys.
- ▶ Encrypted disk pool data is separate and distinct from software-based tape data encryption implemented by IBM i 6.1 Backup Recovery and Media Services (BRMS).

For example, assume you use either a BRMS restore function or a standard IBM i restore command to restore tape media data that was encrypted by BRMS. The first step of the restore process looks for the Master key information used to create the encrypted tape

data. If the master key information is not on this target system, the restore function fails. If that same master key information is on the target system, the encrypted tape data is first un-encrypted.

The tape data is then restored on the target system according to the encryption or no encryption attributes of the disk pool you are restoring into. If the target disk pool has been defined to be encrypted, then the tape data is re-encrypted as it is placed into the target system's encrypted disk pool. If the target pool is not defined to be encrypted, then un-encrypted tape data is placed on the disks in that disk pool.

In order to use disk encryption, you must have 5761-SS1 Option 45 - Encrypted ASP Enablement installed. The option to enable encryption is available when you create a disk pool or independent disk pool.

The encryption option is available on the "create disk pool (ASP)" interface within:

- ▶ IBM Systems Director Navigator for i under **Configuration and Service** → **New Disk Pool**
- ▶ System i Navigator under **Configuration and Service** → **Hardware** → **Disk units** → **Disk Pools**
- ▶ Dedicated Service Tools (DST) or System Service Tools (SST) under **Work with disk units** → **Work with disk configuration**

Note: Encryption of user ASPs can be performed through all above interfaces. However, encryption of independent ASPs can only be performed through graphical user interfaces (IBM Systems Director Navigator for i and System i Navigator).

When you set up an encrypted disk pool, the system generates a data key, which encrypts the data written to that storage pool and decrypts data read from that storage pool. The data keys for independent storage pools is kept with the storage pool and is protected with the ASP master key. Basic ASPs are protected with a data key that is stored in the Licensed Internal Code. The ASP master key is not required for creating an encrypted basic ASP.

Applications running in a partition with dedicated disks are able to access their data independent of its encryption status. IBM i performs necessary decryption on input operations and encryption on output operations. Encrypting the data does not protect the data from IBM i users. Standard IBM i security facilities must be used to protect or limit access to the objects and data within the objects stored in the encrypted disk pool.

For more information about using the Disk Encryption, refer to *Security Guide for IBM i V6.1*, SG24-7680.

2.3 Encrypted backups

You can encrypt backups to tape media to prevent the loss of personal customer information or confidential data if the media is lost or stolen. There are two methods for performing encrypted backups:

- ▶ Hardware encryption using an encryption-capable tape drive. You can use native save/restore commands or Backup Recovery and Media Services (BRMS) with the encrypting tape drive. These functions have been available since OS/400 V5R2.
- ▶ Software encryption using BRMS, available in IBM i 6.1.

Consider the following factors in making your decision about your encryption media and method:

- ▶ Choose the hardware encryption method using an encrypting tape drive if you want the best performance for doing save and restore operations, especially a full-system save or restore operation. You do not need host-based encryption of data or the use of specialized encryption appliances to use the encrypting tape drive.
- ▶ Choose the software encryption method if you want a low-cost solution. This solution is ideal for backing up individual objects that contain customers' personal information or confidential data. Customers with sufficient system resources and a large enough backup window can also encrypt the backup without impacting their business. You can use any tape drive or tape library model with software encryption. However, the performance is not as good as with hardware encryption.

For detailed information about encrypting backups, refer to *Security Guide for IBM i V6.1*, SG24-7680.

2.4 New security system values

This section provides an overview of the new system values.

2.4.1 Password rules (QPWDRULES)

The Password Rules (QPWDRULES) system value specifies the rules used to check whether a password is formed correctly. You can specify more than one value for the QPWDRULES system value, unless you specify *PWDSYSVAL.

Possible values for the QPWDRULES system value are:

*PWDSYSVAL This value specifies that the QPWDRULES system value is ignored and the other password system values are used to check whether a password is formed correctly. These other password system values include QPWDLMTAJC, QPWDLMTCHR, QPWDLMTREP, QPWDMAXLEN, QPWDMINLEN, QPWDPOSDIF, and QPWDQDDGT.

Note: If any value other than *PWDSYSVAL is specified for QPWDRULES, the QPWDLMTAJC, QPWDLMTCHR, QPWDLMTREP, QPWDMAXLEN, QPWDMINLEN, QPWDPOSDIF, and QPWDQDDGT system values are ignored when a new password is checked to see if it is formed correctly. In addition, any attempt to change these system values will be rejected as long as the QPWDRULES system value contains a value other than *PWDSYSVAL.

*CHRLMTAJC The value specifies that a password cannot contain 2 or more occurrences of the same character that are positioned adjacent to each other. This value performs the same function as specifying a value of 2 for the QPWDLMTREP system value. If the *CHRLMTREP value was specified, this value cannot be specified.

*CHRLMTREP The value specifies that a password cannot contain 2 or more occurrences of the same character. This value performs the same function as specifying a value of 1 for the QPWDLMTREP system

	value. If the *CHRLMTAJC value was specified, this value cannot be specified.
*DGTLMTAJC	The value specifies that a password cannot contain 2 or more adjacent digit characters.
*DGTLMTFST	The value specifies that the first character of a password cannot be a digit character. If *LTRLMTFST and *SPCCHRLMTFST values were specified, this value cannot be specified. If the system is operating at password level 0 or 1, the system functions like the *DGTLMTFST value is specified.
*DGTLMTLST	The value specifies that the last character of the password cannot be a digit character. If *LTRLMTLST and *SPCCHRLMTLST values were specified, this value cannot be specified.
*DGTMAXn	The value specifies the maximum number of digit characters that can occur in the password. The n is a number from 0 to 9. Only one *DGTMAXn value can be specified. If a *DGTMINn value is also specified, the n value specified for *DGTMAXn must be greater than or equal to the n value specified for *DGTMINn.
*DGTMINn	The value specifies the minimum number of digit characters that must occur in the password. The n is a number from 0 to 9. Only one *DGTMINn value can be specified. If a *DGTMAXn value is also specified, the n value specified for *DGTMAXn must be greater than or equal to the n value specified for *DGTMINn.
*LMTSAMPOS	The same character cannot be used in a position corresponding to the same position in the previous password. This value performs the same function as the QPWDPOSDIF system value. When the password is set by the Change User Profile (CHGUSRPRF) or Create User Profile (CRTUSRPRF) command, this password rule cannot be checked because the previous password value is not supplied.
*LMTPRFNAME	The uppercase password value cannot contain the complete user profile name in consecutive positions.
*LTRLMTAJC	The value specifies that a password cannot contain 2 or more adjacent letter characters.
*LTRLMTFST	The value specifies that the first character of the password cannot be a letter character. If *DGTLMTFST and *SPCCHRLMTFST values were specified, this value cannot be specified. If the system is operating with a QPWDVLV value of 0 or 1, *LTRLMTFST and *SPCCHRLMTFST cannot both be specified.
*LTRLMTLST	The value specifies that the last character of the password cannot be a letter character. If *DGTLMTLST and *SPCCHRLMTLST values were specified, this value cannot be specified.
*LTRMAXn	The value specifies the maximum number of letter characters that can occur in the password. The n is a number from 0 to 9. Only one *LTRMAXn value can be specified. If a *LTRMINn value is also specified, the n value specified for *LTRMAXn must be greater than or equal to the n value specified for *LTRMINn. If a *MIXCASEn value is also specified, the n value specified for *LTRMAXn must be greater than or equal to 2 times the n value specified for *MIXCASEn.

*LTRMINn	<p>The value specifies the minimum number of letter characters that must occur in the password. The n is a number from 0 to 9.</p> <p>Only one *LTRMINn value can be specified. If a *LTRMAXn value was specified, the n value specified for *LTRMAXn must be greater than or equal to the n value specified for *LTRMINn.</p>
*MAXLENnnn	<p>The value specifies the maximum number of characters in a password. The nnn is a number from 1 to 128 (without leading zeros). This value performs the same function as the QPWDMAXLEN system value.</p> <p>If the system is operating at QPWLVL 0 or 1, the valid range is from 1 to 10. If the system is operating at QPWLVL 2 or 3, the valid range is from 1 to 128.</p> <p>The nnn value specified must be large enough to accommodate all *MIXCASEn, *DGTMAXn, *LTRMAXn, *SPCCHRMAXn, first and last character restrictions, and nonadjacent character requirements.</p> <p>If *MINLENnnn is also specified, the nnn value specified for *MAXLENnnn must be greater than or equal to the nnn value specified for *MINLENnnn.</p> <p>If no *MAXLENnnn value is specified, a value of *MAXLEN10 is assumed if the system is operating with a QPWLVL value of 0 or 1 or a value of *MAXLEN128 is assumed if the system is operating with a QPWLVL value of 2 or 3.</p>
*MINLENnnn	<p>The value specifies the minimum number of characters in a password. The nnn is a number from 1 to 128 (without leading zeros).</p> <p>If the system is operating at QPWLVL 0 or 1, the valid range is from 1 to 10. If the system is operating at QPWLVL 2 or 3, the valid range is from 1 to 128.</p> <p>If *MAXLENnnn is also specified, the nnn value specified for *MAXLENnnn must be greater than or equal to the nnn value specified for *MINLENnnn.</p> <p>If no *MINLENnnn value is specified, a value of *MINLEN1 is assumed.</p>
*MIXCASEn	<p>The value specifies that a password must contain at least n uppercase and n lowercase letters. The n is a number from 0 to 9. This value is rejected if the system is operating with a QPWLVL value of 0 or 1 because passwords are required to be uppercase.</p> <p>Only one *MIXCASEn value can be specified.</p> <p>If a *LTRMAXn value was specified, the n value specified for *LTRMAXn must be greater than or equal to two times the n value specified for *MIXCASEn.</p>
*REQANY3	<p>The value specifies that a password must contain characters from at least three of the following four types of characters:</p> <p><i>Uppercase letters</i> <i>Lowercase letters</i> <i>Digits</i> <i>Special characters</i></p> <p>When the system is operating with a QPWLVL of 0 or 1, *REQANY3 has the same effect as though *DGTMIN1, *LTRMIN1, and *SPCCHRMIN1 were all specified.</p>
*SPCCHRLMTAJC	<p>The value specifies that a password cannot contain 2 or more adjacent (consecutive) special characters. A character is considered as a special character if its equivalent unicode character has the property of not being a letter nor a digit.</p>

*SPCCHRLMTFST	<p>The value specifies that the first character of the password cannot be a special character. A character is considered as a special character if its equivalent unicode character has the property of not being a letter nor a digit.</p> <p>If *DGTLMTFST and *LTRLMTFST values were specified, this value cannot be specified. If the system is operating with a QPWDVLV value of 0 or 1, *LTRLMTFST and *SPCCHRLMTFST cannot both be specified.</p>
*SPCCHRLMTLST	<p>The value specifies that the last character of the password cannot be a special character. A character is considered as a special character if its equivalent unicode character has the property of not being a letter nor a digit.</p>
*SPCCHRMAXn	<p>The value specifies the maximum number of special characters that may occur in the password. The n is a number from 0 to 9. A character is considered as a special character if its equivalent unicode character has the property of not being a letter nor a digit.</p> <p>Only one *SPCCHRMAXn value can be specified. If a *SPCCHRMINn value was specified, the n value specified for *SPCCHRMAXn must be greater than or equal to the n value specified for *SPCCHRMINn.</p>
*SPCCHRMINn	<p>The value specifies the minimum number of special characters that must occur in the password. The n is a number from 0 to 9. A character is considered as a special character if its equivalent unicode character has the property of not being a letter nor a digit.</p> <p>Only one *SPCCHRMINn value can be specified. If a *SPCCHRMAXn value was specified, the n value specified for *SPCCHRMAXn must be greater than or equal to the n value specified for *SPCCHRMINn.</p>

For a list of samples about how to set up password rules, access the IBM i Information Center and select the path **Security** → **Security reference** → **Security system values** → **System values that apply to passwords** → **Password Rules** → (QPWDRULES).

2.4.2 Password expiration warning (QPWDEXPWRN)

The password expiration warning (QPWDEXPWRN) system value controls the number of days prior to a password expiring, to begin displaying password expiration warning messages on the Sign-on Information display.

2.4.3 Block password change (QPWDCHGBLK)

The block password change (QPWDCHGBLK) system value specifies the time period during which a password is blocked from being changed following the prior successful password change operation. This system value does not restrict password changes made by the Change User Profile (CHGUSRPRF) command.

2.4.4 Secure Sockets Layer cipher specification list (QSSLCSL)

If you specify the Use user-defined (*USRDFN) option for the Secure Sockets Layer cipher control (QSSLCSLCTL) system value, you can define the Secure Sockets Layer cipher specification list (QSSLCSL) system value. If the QSSLCSLCTL system value is system defined, the QSSLCSL system value is read-only.

The System SSL uses the sequence of the values in the QSSLCSL system value to order the default cipher specification list. The default cipher specification list entries are system-defined and can change with different releases. If a default cipher suite is removed from the QSSLCSL system value, the cipher suite is removed from the default list. The default cipher suite is added back to the default cipher specification list when it is added back into the QSSLCSL system value. You cannot add other cipher suites to the default list beyond the set that the system defines for the release.

You cannot add a cipher suite to the QSSLCSL system value if the required SSL protocol value for the cipher suite is not set for the Secure Sockets Layer protocols (QSSLPCL) system value.

This system value can have the following values:

- ▶ ***RSA_AES_128_CBC_SHA**
Use the RSA encoding algorithms for the Advanced Encryption Standard (AES) cipher with cipher block chaining (CBC) and 128-bit keys. Use Secure Hash Algorithm (SHA) for generating message authentication codes (MAC).
- ▶ ***RSA_RC4_128_SHA**
Use the RSA encoding algorithms for Rivest Cipher 4 (RC4) and 128-bit keys. Use SHA for generating MAC.
- ▶ ***RSA_RC4_128_MD5**
Use the RSA encoding algorithms for the RC4 cipher and 128-bit keys. Use message digest algorithm 5 (MD5) for generating MAC.
- ▶ ***RSA_AES_256_CBC_SHA**
Use the RSA encoding algorithms for the AES cipher with CBC and 256-bit keys. Use SHA for generating MAC.
- ▶ ***RSA_3DES_EDE_CBC_SHA**
Use the RSA encoding algorithms for the Triple Data Encryption Standard (3DES) cipher with the encrypt/decrypt/encrypt (EDE) and CBC modes and 168-bit keys. Use SHA for generating MAC.
- ▶ ***RSA_DES_CBC_SHA**
Use the RSA encoding algorithms for the Data Encryption Standard (DES) cipher with the CBC mode and 56-bit keys. Use SHA for generating MAC.
- ▶ ***RSA_EXPORT_RC2_CBC_40_MD5**
Use the RSA encoding algorithms for Rivest Cipher 2 (RC2) with the CBC mode and 40-bit keys. Use MD5 for generating MAC.
- ▶ ***RSA_EXPORT_RC4_40_MD5**
Use the RSA encoding algorithms for the RC4 cipher and 40-bit keys. Use MD5 for generating MAC.
- ▶ ***RSA_NULL_SHA**
Use the RSA encoding algorithms but do not use any cipher. Use SHA for generating MAC.
- ▶ ***RSA_NULL_MD5**
Use the RSA encoding algorithms but do not use any cipher. Use MD5 for generating MAC.

- ▶ *RSA_RC2_CBC_128_MD5
Use the RSA encoding algorithms for the RC2 cipher with the CBC mode and 128-bit keys. Use MD5 for generating MAC.
- ▶ *RSA_3DES_EDE_CBC_MD5
Use the RSA encoding algorithms for the 3DES cipher with the EDE and CBC modes and 168-bit keys. Use MD5 for generating MAC.
- ▶ *RSA_DES_CBC_MD5
Use the RSA encoding algorithms for the DES cipher with the CBC mode and 56-bit keys. Use MD5 for generating MAC.

2.4.5 Secure Sockets Layer protocols (QSSLPCL)

The Secure Sockets Layer protocols (QSSLPCL) system value specifies the Secure Sockets Layer (SSL) protocols supported by the System SSL.

The values allowed for the QSSLPCL system value are as follows:

- *OPSYS Determined by the system.
- *TLSV1 Transport Layer Security version 1.0 will be supported.
- *SSLV2 Secure Sockets Layer version 2.0 will be supported.
- *SSLV3 Secure Sockets Layer version 3.0 will be supported.

2.4.6 Secure Sockets Layer cipher control (QSSLCSLCTL)

The Secure Sockets Layer cipher control (QSSLCSLCTL) system value specifies whether the system or the user controls the Secure Sockets Layer cipher specification list (QSSLCSL) system value.

The values allowed for the QSSLCSLCTL system value are as follows:

- *OPSYS The values contained in the QSSLCSL (SSL cipher specification list) system value are automatically modified to contain the list of supported cipher suites as determined by the operating system release.
- *USRDFN The QSSLCSL system value can be modified.

2.5 Changed security system values

This section describes the changes made in IBM i 6.1 to system values that were available in previous releases.

2.5.1 Limit device sessions (LMTDEVSSN)

The limit device sessions (QLMTDEVSSN) parameter controls the number of device sessions a user can sign on to. In previous releases, valid values were more than one device sessions (0) or one device session (1). In IBM i 6.1, this has been changed to:

- 0 Users are not limited to a specific number of device sessions.
- 1-9 Indicates maximum number of concurrent device sessions.

This parameter has been changed on the Create User Profile (CRTUSRPRF) and the Change User Profile (CHGUSRPRF) commands.

2.5.2 Verify object on restore (QVFYOBJRST)

The Verify Object on Restore (QVFYOBJRST) system value determines whether objects are required to have digital signatures in order to be restored to your system. You can prevent anyone from restoring an object, unless that object has a correct digital signature from a trusted software provider. This value applies to objects of types *PGM, *SRVPGM, *SQLPKG, *CMD, and *MODULE. It also applies to *STMF objects, which contain Java programs.

Program, service program, and module objects that are created or converted on a system with a release before IBM i 6.1 are treated as unsigned when they are restored to an IBM i 6.1 or later system. Likewise, program, service program, and module objects that are created or converted on an IBM i 6.1 or later release are treated as unsigned when they are restored to a system before IBM i 6.1.

Objects that have the system-state attribute and objects that have the inherit-state attribute are required to have a valid signature from a system-trusted source. Objects in Licensed Internal Code PTFs are also required to have a valid signature from a system-trusted source. If these objects do not have a valid signature, they cannot be restored, regardless of the value of the QVFYOBJRST system value.

To learn more about Security system values, refer to the following information sources:

- ▶ The IBM i Information Center, selecting the path **Security** → **Security Reference** → **Security system values**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680

2.6 Restricting changes to security settings

You can restrict users from changing security-related system values. System service tools (SST) and dedicated service tools (DST) provide an option to lock these system values. By locking the system values, you can prevent even a user with *SECADM and *ALLOBJ authority from changing these system values with the CHGSYSVAL command. In addition to restricting changes to these system values, you can also restrict adding digital certificates to digital certificate store with the Add Verifier API and restrict password resetting on the digital certificate store.

2.7 User profile enhancements

This section details the changes made to some of the commands used to manage user profiles.

2.7.1 Block password change (PWDCHGBLK)

This new parameter specifies the time period during which a password is blocked from being changed, following the prior successful password change operation. This value can be used to prevent users from reusing the same expired password value by simply changing their password numerous times to get back to the expired password value. Block password change (1-99 hours).

This parameter has been added to the Create User Profile (CRTUSRPRF) and the Change User Profile (CHGUSRPRF) commands.

2.7.2 Limit device sessions (LMTDEVSSN)

The limit device sessions (QLMTDEVSSN) parameter controls the number of device sessions a user can sign on to. In previous releases, valid values were more than one device sessions (0) or one device session (1). In IBM i 6.1, this has been changed to:

- 0 Users are not limited to a specific number of device sessions.
- 1-9 Indicates maximum number of concurrent device sessions.

This parameter has been changed on the Create User Profile (CRTUSRPRF) and the Change User Profile (CHGUSRPRF) commands.

2.7.3 Display User Profile (DSPUSRPRF)

There are also some new additions and changes for the Display User Profile (DSPUSRPRF) command:

- ▶ Password verifications not valid

The number of password verifications that were not valid for the specified user since the previous successful password verification. Passwords are verified by sign-on operations, various servers such as File Transfer Protocol (FTP), and API calls such as Get Profile Handle (QSYGETPH).

Note: “Sign-on attempts not valid” has been changed to “Password verifications not valid” for IBM i 6.1.

- ▶ Days before password expires

The number of days left before the password expires. This field is not always displayed.

- ▶ Password set expired by command

Specifies whether the password is set to expired for this user by a system administrator using the Create User Profile (CRTUSRPRF) or Change User Profile (CHGUSRPRF) command.

Note: “Set password to expired” has been changed to “Password set expired by command” for IBM i 6.1.

- ▶ Creation date/time

The date and time the user profile was created.

- ▶ Change date/time

The date and time of the last change to the user profile. This includes automatic updates to the user profile when it is updated for operations such as when one of its owned objects is deleted.

- ▶ Last used date

The date the user profile was used last.

- ▶ Restore date/time

The date and time the user profile was last restored.

2.8 Service Tools user profile and IBM i profile

In IBM i 6.1 there are two new capabilities associated with Dedicated Service Tools or System Service Tools (DST or SST) user IDs and IBM i user profiles:

- ▶ The IBM i Display Service Tools User ID (DSPSSTUSR) command
- ▶ Association between a service tools ID and an IBM i user profile

2.8.1 Display Service Tools User ID (DSPSSTUSR)

The new Display Service Tools User ID (DSPSSTUSR) command provides the ability to view information about a specific, or all, service tools user IDs without having to use DST or SST interfaces. Figure 2-6 shows the output from a DSPSSTUSR USRID(*ALL) command.

```
Display Service Tools User IDs                                     System:  RCHAS70
Type options, press Enter.
 5=Display

Opt  DST/SST ID  Status   Linked   Description
-----
QSECOFR  *ENABLED  QSECOFR  QSECOFR
QSRV     *ENABLED
11111111 *ENABLED  11111111
22222222 *ENABLED  22222222

F3=Exit  F5=Refresh  F12=Cancel  F17=Top  F18=Bottom                                     Bottom
```

Figure 2-6 Display Service Tools User IDs (all users)

If a specific user is selected for display, information similar to that in Figure 2-7 is shown.

```

                                Display Service Tools User ID
                                System:   RCHAS70
Service tools user ID . . . . . : QSECOFR
Previous sign-on . . . . . : 07/24/09 08:07:33
Password verification not valid . . . . . : 0
Status . . . . . : *ENABLED
Linked user profile . . . . . : QSECOFR
  Linked user profile exists . . . . . : Yes
Date password last changed . . . . . : 07/24/09
Date password expires . . . . . : None
Password is set expired . . . . . : No
Description . . . . . : QSECOFR
Privileges:
  Disk units - operations . . . . . : *GRANTED
  Disk units - administration . . . . . : *GRANTED
  Disk units - read only . . . . . : *GRANTED
  System partitions - operations . . . . . : *GRANTED
  System partitions - administration . . . . . : *GRANTED
  Partition remote panel key . . . . . : *GRANTED
  Operator panel functions . . . . . : *GRANTED
  Operating system initial program load
  (IPL) . . . . . : *GRANTED
  Install . . . . . : *GRANTED
  Performance data collector . . . . . : *GRANTED
  Hardware service manager . . . . . : *GRANTED
  Display/Alter/Dump . . . . . : *GRANTED
  Main storage dump . . . . . : *GRANTED
  Product activity log . . . . . : *GRANTED
  Licensed Internal Code log . . . . . : *GRANTED
  Licensed Internal Code fixes . . . . . : *GRANTED
  Trace . . . . . : *GRANTED
  Dedicated service tools (DST) environment: *GRANTED
  Remote service support . . . . . : *GRANTED
  Service tools security . . . . . : *GRANTED
  Service tools save and restore . . . . . : *GRANTED
  Debug . . . . . : *GRANTED
  System capacity - operations . . . . . : *GRANTED
  System capacity - administrator . . . . . : *GRANTED
  System security . . . . . : *GRANTED
  Start service tools . . . . . : *GRANTED
  Take over console . . . . . : *GRANTED

F3=Exit   F12=Cancel   F15=Display linked user profile

```

Figure 2-7 Display Service Tools User IDs (specific user)

2.8.2 Service tools user ID and IBM i user profile association

You can associate a service tools user ID and password with an IBM i user ID and password. There are two methods for doing this association:

- ▶ In DST or SST, explicitly link a service tool user ID and password with an IBM i user profile and password.

If you explicitly link a service tool user ID to an IBM i user profile, the IBM i user can have the privileges of the service tools user ID while using an IBM i interface.

- ▶ Logically link a service tool user ID and password by specifying a service tool ID and password that exactly matches a corresponding IBM i user profile and password. In this case you do not have to explicitly link the service tool ID and IBM i user profile. However, there are specific requirements that must be followed for this linkage to function:
 - In the service tools, change the password level to level 2, which uses Secure Hash Algorithm (SHA) encryption. With this setting all the service passwords become case sensitive. When defining the next service tool user ID's password, the password you enter is case sensitive.
 - Later, when signing directly on to the service tools, you are required to enter the password in exactly the same case as originally defined (or later changed in the service tools). This upper/lower case sensitivity is in both DST and SST.
 - If the IBM i password level system value QPWDVLV remains at 0 or 1, the IBM i password processing treats upper and lower case letters the same (not case sensitive).

When using IBM Systems Director Navigator for i, the Configuration and Service links for the disk functions have a specific requirement for the IBM i user profile and a matching service tools ID. The service tools ID and IBM i user profile must have the same password with the added qualification the service tools ID's password must be in uppercase. If these rules are not followed, the attempt to use a disk management function from IBM Systems Director Navigator for i will fail.

For more information about the new DSPSSTUSR command and linking service tools user IDs to user profiles, refer to:

- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ The IBM i Information Center, selecting the path **Security** → **Security Reference** → **Security system** → **values**

2.9 Object signing and signature verification

In OS/400 V5R2, IBM i shipped with a code checking function that you can use to verify the integrity of signed objects on your system, including all operating system code that IBM ships and signs for your system. Beginning in i5/OS V5R3, you can use the new Check System Application Programming Interface (API) to verify the integrity of the code checking function itself, as well as key operating system objects.

In IBM i 6.1, IBM signs the Licensed Internal Code (LIC) and you can either use the Check System (QydoCheckSystem) API or the Check Object Integrity (CHKOBJITG) command to verify the LIC. The Check System API reports error messages for verification failures and other errors or verification failures to the job log. The CHKOBJITG command stores the corresponding information in a nominated output file.

To learn more about object signing and signature verification, refer to the following information sources:

- ▶ The IBM i Information Center, selecting the path **Security** → **Object signing and signature verification**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680

2.10 Digital Certificate Manager

IBM i provides extensive digital certificate support that allows you to use digital certificates as credentials in a number of security applications. Digital Certificate Manager (DCM) allows you to manage digital certificates for your network and use Secure Sockets Layer (SSL) to enable secure communications for many applications.

2.10.1 Starting Digital Certificate Manager

In IBM i 6.1, DCM is still started by selecting Digital Certificate Manager from the i5/OS Tasks Page. However, access to the i5/OS Tasks Page has changed. It is now access from the Welcome page of IBM Systems Director Navigator for i. Refer to 16.1, “IBM Systems Director Navigator for i” on page 404 for more information about the functions in IBM Systems Director Navigator for i.

2.11 Single sign-on

Many of the new or enhanced single sign-on functions are a result of new and enhanced function for EIM and network authentication service, the two technologies that make up the IBM i single sign-on solution.

To learn more about single sign-on and EIM, refer to:

- ▶ “Enterprise Identity Mapping” on page 24
- ▶ *Security Guide for IBM i V6.1*, SG24-7680

2.12 Enterprise Identity Mapping

In previous releases of IBM i, Enterprise Identity Mapping (EIM) only supported mapping to one local user identity per system. In IBM i 6.1 EIM supports selecting from multiple local user identity mappings for the same system, using the IP address of the target system to select the correct local user identity mapping on that system. Single Sign-on has also been updated to take advantage of this new support.

Now you can use Enterprise Identity Mapping Java classes instead of APIs to perform identity mapping operations.

To learn more about EIM and how to configure it on the IBM i platform, refer to the following information sources:

- ▶ The IBM i Information Center, selecting the path **Security** → **Enterprise Identity Mapping**

- ▶ *Windows-based Single Signon and the EIM Framework on the IBM eServer iSeries Server, SG24-6975*

2.13 Network Authentication Service

2.13.1 New Kerberos control language commands

In IBM i 6.1, the following Kerberos CL commands have been added:

- ▶ Add Kerberos Keytab Entry (ADDKRBKTE)
- ▶ Add Kerberos Ticket (ADDKRBTKT)
- ▶ Change Kerberos Password (CHGKRBPWD)
- ▶ Delete Kerberos Credentials Cache File (DLTKRBCCF)
- ▶ Display Kerberos Credentials Cache File (DSPKRBCCF)
- ▶ Display Kerberos Keytab Entries (DSPKRBKTE)
- ▶ Remove Kerberos Keytab Entry (RMVKRBKTE) command

For more information about these commands, refer to the following Infocenter topics:

- ▶ Control language
- ▶ Changing Kerberos Passwords
- ▶ Deleting expired credentials cache files
- ▶ Displaying credentials cache
- ▶ Obtaining or renewing ticket-granting tickets
- ▶ Managing keytab files

2.13.2 New service principal for Network File System server

With IBM i 6.1, you can add and update key table entries for the Network File System server as shown in Figure 2-8, taken from the Network Authentication Service (NAS) configuration wizard of the System i Navigator for Windows.

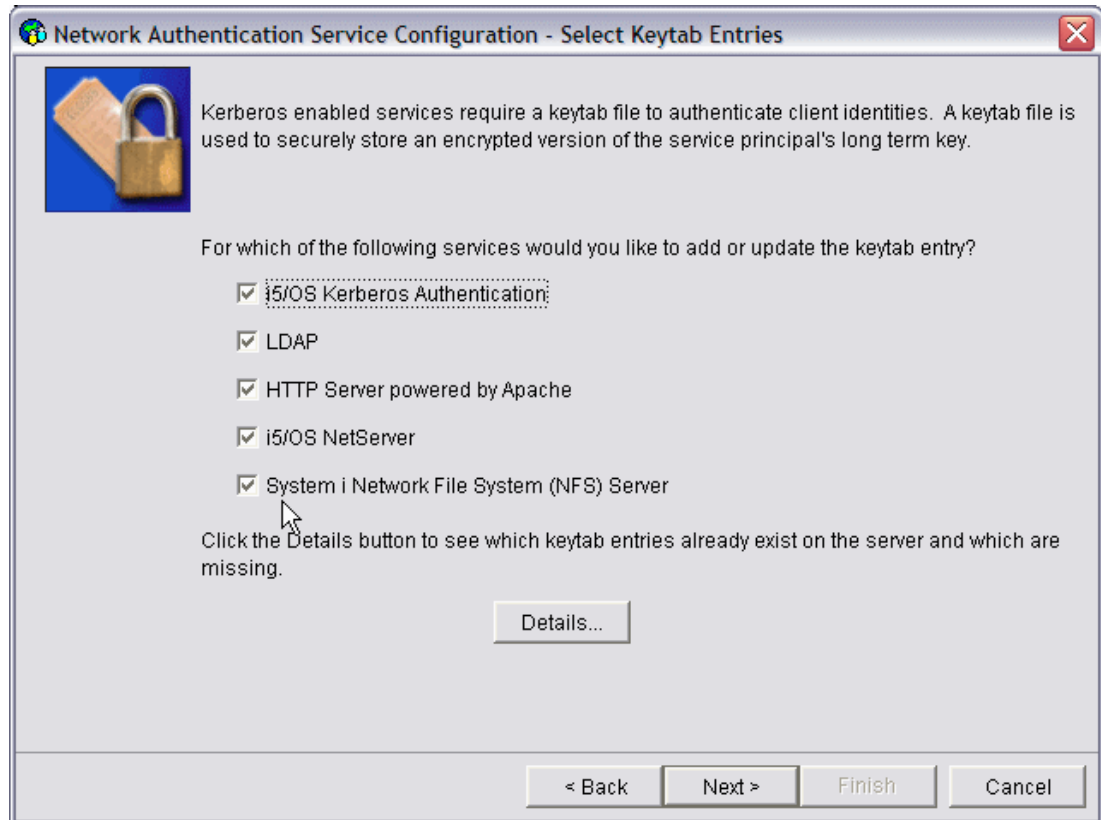


Figure 2-8 Network Authentication Service Configuration Wizard

Depending on the operating system on which you have configured a Kerberos server, the steps for adding the IBM i principal are different. We provide instructions on adding the IBM i principals to a Kerberos server in IBM i PASE or a Windows 2000 domain. If you have optionally created service principals for IBM Tivoli® Directory Server for IBM i (LDAP), IBM i NetServer, Network File System (NFS) Server, or HTTP Server, you must also add those service principals to the Kerberos server.

2.14 Cryptographic enhancements

IBM offers several IBM i cryptography solutions. A comprehensive cryptography solution is an important part of a successful security strategy. IBM offers both software cryptography and a family of cryptographic hardware options for protecting data and for securing transaction processing.

Encryption involves an algorithm that transforms understandable information (plaintext) into an unintelligible piece of data (ciphertext), and decryption involves the transformation of the unintelligible data back into understandable information.

IBM i 6.1 provides encryption in SQL column specifications, encryption of data stored on tape under Backup Recovery and Media Services (BRMS), and encryption of all data within a disk pool that has been created as encrypted. There are also IBM i APIs that you can use to perform data encryption and decryption.

There are various algorithms for encryption and decryption, all of which are controlled by encryption keys and typically defined by some government or international standards organization.

New in IBM i 6.1 is that the SQL Column encryption can make use of Advanced Encryption Standard (AES), a standard adopted by the U.S. government for encryption.

Additional information can be found in the following:

- ▶ The IBM i Information Center, selecting the path **Security** → **Cryptography**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ *IBM System i Security: Protecting i5/OS Data with Encryption*, SG24-7399

2.14.1 Cryptographic key management enhancements

Cryptographic services key management for the IBM i operating system enables you to store and manage master keys and key stores. Since you are exchanging sensitive data to manage master keys and key stores, it is recommended that you use a secure session.

Cryptographic Services supports a hierarchical key system. At the top of the hierarchy is a set of master keys. These keys are the only key values stored in the clear (unencrypted). Cryptographic services securely stores the master keys in the Licensed Internal Code (LIC), also known as machine code.

Eight general-purpose master keys are used to encrypt other keys which can be stored in key store files, which are database files. Any type of key supported by cryptographic services can be stored in a key store file, for example AES, RC2, RSA, and SHA1-HMAC.

In addition to the eight general-purpose master keys, cryptographic services supports two special-purpose master keys.

- ▶ The ASP master key is used for protecting data in the independent Auxiliary Storage Pool (in the Disk Management GUI it is known as an Independent Disk Pool).
- ▶ The save/restore master key is used to encrypt the other master keys when they are saved to media using a Save System (SAVSYS) operation.

New easy-to-use interfaces were added for Cryptographic services key management. Key management can now be performed through a set of control language (CL) commands and through the Cryptographic Services Key Management graphical interface, added as part of System i Navigator and IBM Systems Director Navigator for i.

For complete information about cryptographic key management, refer to the following information sources:

- ▶ 2.14.2, “Save/restore master key” on page 28
- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ The IBM i Information Center, selecting the path **Security** → **Cryptography** → **Cryptographic services key management**

2.14.2 Save/restore master key

Master keys are used to encrypt other keys. You can load, set, and test master keys. You can clear a master key only after you have set it.

In this release master keys are now saved in encrypted form by the Save System (SAVSYS) command. The master keys are encrypted under a new “master key” called the save/restore master key. This helps with disaster recovery. The master keys are automatically restored during an install of LIC, provided that the default save/restore key is used. It is recommended, however, that users set their own save/restore master key. In that case the master keys are not restored until the save/restore master key is first entered. This prevents misuse of the master keys in case the install media is lost or falls into the wrong hands.

The master key management functions also support an ASP master key. This master key protects the data keys that are used for encrypting Independent ASPs.

For more information about save/restore of the master key, refer to the following information sources:

- ▶ The IBM i Information Center, selecting the path **Security** → **Cryptography** → **Cryptographic services key management** → **Saving and restoring master keys**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680

2.14.3 4764 Cryptographic coprocessor

IBM offers a Cryptographic Coprocessor, which is available on a variety of system models. Cryptographic Coprocessors contain hardware engines, that perform cryptographic operations used by IBM i application programs and IBM i SSL transactions.

The 4764 Cryptographic Coprocessor can run in the CCA and the accelerator modes simultaneously. While the accelerator mode of operation uses clear keys only as opposed to the CCA mode that uses only secure keys, it processes cryptographic operations at a much higher rate than CCA mode. Accelerator mode is used exclusively for RSA private key operations during SSL session establishment provided that the private key was not generated by the 4764 operating in CCA mode.

IBM i provides the IBMJCECCAI5OS implementation that extends Java Cryptography Extension (JCE) and Java Cryptography Architecture (JCA) to add the capability to use hardware cryptography, by using the IBM Common Cryptographic Architecture (CCA) interfaces.

The IBMJCECCAI5OS provider takes advantage of hardware cryptography in the existing JCE architecture and gives Java 2 programmers the significant security and performance advantages of hardware cryptography with minimal changes to existing Java applications. As the complexities of hardware cryptography are taken care of within the normal JCE, advanced security and performance using hardware cryptographic devices are made easily available.

The IBMJCECCAI5OS provider plugs into the JCE framework in the same manner as the current providers. For hardware requests, the CCA APIs are invoked through the new native methods. The IBMJCECCAI5OS provider stores CCA RSA key labels in a new Java key store type of JCECCAI5OSKS.

Notes:

- ▶ The 4758 Cryptographic Coprocessor and the 2058 Cryptographic Accelerator are no longer available but are still supported.
- ▶ IBM i Cryptographic Device Manager, 5733-CY1, is replaced by Cryptographic Device Manager, 5733-CY2, in IBM i 6.1.

For more information about using the Cryptographic Coprocessor, refer to the following information sources:

- ▶ The IBM i Information Center, following the path **Security** → **Cryptography** → **4764 Cryptographic Coprocessor**
- ▶ The IBM i Information Center, selecting the path **Programming** → **Java** → **IBM Developer Kit for Java** → **Java Security** → **Java Cryptography Extension** → **Using hardware cryptography**
- ▶ *Security Guide for IBM i V6.1*, SG24-7680



Backup and recovery

This chapter discusses the following topics:

- ▶ “File level backup for Linux” on page 32
- ▶ “Save-while-active” on page 32
- ▶ “Deferring the restore of dependent objects” on page 34
- ▶ “Saving and restoring private authorities” on page 35
- ▶ “Alternate installation device” on page 35
- ▶ “Recovering vital product data” on page 35
- ▶ “Program conversion restoring to a different hardware type” on page 36
- ▶ “Program conversion restoring to a different release” on page 36
- ▶ “Restore user profile” on page 36
- ▶ “User-defined file system backup” on page 37
- ▶ “Restore actions for mounted user-defined file systems” on page 38
- ▶ “Encrypted backups” on page 38
- ▶ “Encrypted auxiliary storage pools” on page 39
- ▶ “Restore the operation output file” on page 39
- ▶ “Restoring external stored procedures” on page 40
- ▶ “Backup and restore for IBM i on BladeCenter” on page 40
- ▶ “Retrieve Disk Information (RTVDSKINF)” on page 41
- ▶ “Backup, Recovery and Media Services” on page 41
- ▶ “Journal management” on page 70
- ▶ “Restrictions for previous release support” on page 73

3.1 File level backup for Linux

File level backup for Linux partitions and Linux running on an integrated BladeCenter or System x server is now available with IBM i 6.1. The function is available using the Save Object (SAV) and Restore Object (RST) commands and also Backup Recovery and Media Services.

Some setup is required in IBM i and Linux to enable file level backups:

- ▶ IBM i
 - Add file shares to the save list file (QUSRSYS/QAZLCSAVL).
 - At least one share per volume to back up files contained therein.
- ▶ Linux
 - Version 3 or above of Samba installed.
 - OpenSSL package installed (used for some libraries).
 - File shares defined on volumes.
 - Userid on Linux must be the same and have same password as the IBM i userid.
 - Backup program needs to run as root to have full access to all files to enable backup.

Note: The IBM Extended Integrated Server Support for i5/OS licensed program (5761-LSV) must be installed to use this feature.

For more information, refer to the relevant pages in the IBM Systems Information Center, starting with *Backing up and recovering individual files and directories for Linux servers running in logical partitions* at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzalm/rzalmba/ckinguplinuxfiles.htm>

3.2 Save-while-active

IBM i 6.1 brings several enhancements to the save-while-active set of functions.

3.2.1 Single synchronization point for multiple save commands

In previous releases, save-while-active created a separate synchronization point each time a save was started. In backup strategies where different parts of the system are saved at different times, it was not possible to restore a system to a single point in time using save-while-active backups.

The new command, Start Save Synchronization (STRSAVSYNC), starts a synchronized checkpoint for more than one save-while-active operation. This allows you to save objects in one or more libraries, and objects in one or more directories, and have all the objects reach a checkpoint together.

In addition to the benefit of having a single checkpoint for multiple saves, it improves system performance for save-while-active because there is only one checkpoint.

To create a single checkpoint for multiple saves, the STRSAVSYNC command is issued prior to the SAV nnn commands, for example:

1. Issue the STRSAVSYNC command specifying:
 - Synchronization ID (SYNCID) - A valid name for your synchronization ID
 - Number of operations (NUMSYNC) - Number of individual saves that will be performed
 - Start save wait time (STRSAVWAIT) - Specify the time in seconds to wait for all of the participating save operations to be started. If the participating save operations specified for the number of operations in NUMSYNC do not start within the specified time, any operations that do start within that time will be ended.
2. Start the first save, referencing the Synchronization ID that was specified on the STRSAVSYNC command. The save can be started from the same job that issued the STRSAVSYNC command or another job. This action starts the QRSYNCM job in QSYSWRK, which manages the synchronization tasks. The save command that was issued is then interrupted, awaiting the start of the other saves for the same Synchronization ID.
3. Start the other saves, referencing the Synchronization ID that was specified on the STRSAVSYNC command.
4. When the number of operations (saves) from the STRSAVSYNC command is reached, a common synchronization point is established for all the pending save commands and then the saves commence.

For more information, refer to the relevant pages in the IBM Systems Information Center, starting with *Using save-while-active to synchronize the saved data* at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzaiu/rzaiusavewhileactivesynch.htm?>

3.2.2 Saving storage spaces for active integrated servers

The save-while-active function can be used to save the network server storage space while keeping the integrated server online. The save-while-active function reduces or eliminates the outage for save operations.

Network storage spaces for the following client partition operating systems can be saved while they are active:

- ▶ IBM i
- ▶ Windows
- ▶ Linux
- ▶ AIX
- ▶ VMware

The save-while-active function can be used when the Network Server Storage Space is linked to *ISCSI and *GUEST Network Server Descriptions, but not when the disk is linked to *IXSVR.

Note: The IBM Extended Integrated Server Support for i5/OS licensed program (5761-LSV) must be installed to use this feature for Linux.

The following example shows how to save the storage space for an active integrated server.

This command saves the objects that are associated with the 'QFPNWSSTG/MYDISK storage space whose associated network server description (*NWS) is currently varied on.

```
SAV DEV('/QSYS.LIB/MYLIB.LIB/MYSAVF.FILE') OBJ('/QFPNWSSTG/MYDISK')
SAVACT(*YES) SAVACTOPT(*NWSSTG)
```

If you do not specify the SAVACT (save active) and SAVACTOPT (save active option) parameters on the SAV command, you must shut down the integrated server to save the storage space.

3.3 Deferring the restore of dependent objects

You can defer the restore of database logical files, indexes, and SQL-materialized query tables (MQTs), that have dependencies on other files.

In releases before IBM i 6.1, restoring files that resided in a different library than the based-on files was difficult to manage. If the library with a logical file was restored before the library with the underlying physical file, the logical file did not restore. To restore the logical file, you needed to attempt to restore from the media one or more times after restoring the physical file. Materialized query tables (MQTs) are physical files that have dependencies on based-on files. An MQT could be restored if one of its based-on files was missing.

You can now defer the restore of logical files and MQTs whose based-on files or members are missing. A deferred logical file is cached and restored from the cache after its based-on files are available. A deferred MQT is created, so that its data spaces can be restored from the media on the first pass. By deferring the restore of logical files and MQTs until the based-on files become available, you can complete the restore operation in one pass through the media, which saves time.

To defer a restore of database files whose based-on files are missing, specify the Defer ID (DFRID) parameter on the Restore Library (RSTLIB) or Restore Object (RSTOBJ) command. Specifying the same defer ID for multiple restore operations provides an association between the set of deferred files.

When the based-on files become available, use the Restore Deferred Objects (RSTDFROBJ) command with the same defer ID to complete the restore of deferred objects if the based-on files are available. If you restore the missing based-on files with the same defer ID that you used for the dependent, deferred files, the system automatically attempts to complete the restore of the deferred files at that time.

Note: Deferred objects can be logical files or MQTs. A deferred logical file is not created until the restore operation is complete. A deferred MQT is created, but until the restore operation is complete, you cannot perform functions on the MQT that require access to the based-on files.

After the deferred restore operation completes, use the Remove Defer ID (RMVDFRID) command to remove information about a deferred file from the cache. Use the RMVDFRID command if you have created or restored a file with the same name using a different defer ID or no defer ID.

If you have save system (*SAVSYS) special authority and use the RSTLIB command with the following parameters, the system uses the DFRID parameter value of QRSTLIB to defer the restore of objects that are missing based-on files:

- ▶ Specify the *NONSYS, *ALLUSR, or *IBM values for the Saved Library (SAVLIB) parameter.
- ▶ Specify *DFT for the DFRID parameter or omit the DFRID parameter.

In this case, the system manages the deferred files. In all other cases, if you omit the DFRID parameter, you need to manage the dependent database files.

3.4 Saving and restoring private authorities

In previous releases, the private authorities to an object were only saved by the Save System (SAVSYS) and Save Security Data (SAVSECDA) commands. Reapplying the private authorities was a two-step process: first the private authorities were restored to the system using the Restore User Profiles (RSTUSRPRF) command, and then “applied” using the Restore Authority (RSTAUT) command.

IBM i 6.1 provides an option to save the private authorities to an object when the object is saved. They can also be restored if they were previously saved. Many SAV nnn and RST nnn commands have a new private authorities (PVTAUT) parameter to use this new support.

Saving private authorities will increase the amount of time it takes to save the objects, but it can simplify the recovery of an object or a group of objects. It will not simplify the recovery of an entire system.

3.5 Alternate installation device

You can do an IPL-type D from a tape device attached by a Fibre Channel interface only when it is attached to an adapter without an input/output processor (IOP) on a POWER6 model. This includes tape devices supported on Fibre Channel adapters #5735, #5749, and #5774.

For other system models and configurations, a tape device attached by a Fibre Channel interface needs to be set up as an alternate installation device. For more information, see the *Alternate installation device: Overview* section in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmaidov.htm>

3.6 Recovering vital product data

The system’s vital product data (VPD) information contains information about hardware features on the system. Disk VPD includes which disk unit is the load-source disk unit. If the VPD information is lost, the partition might not IPL in A-mode or B-mode.

The procedure for recovering VPD from the flexible service processor (FSP) during a scratch installation (IPL in D-mode) is documented in Infocenter in the *Recovering the vital product data information if the partition does not IPL in A-mode or B-mode* section at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmvdp.htm>

3.7 Program conversion restoring to a different hardware type

The system can now create *MODULE, *PGM, and *SRVPGM objects that take advantage of the latest hardware features of the system on which they are created or converted. New hardware features can be used immediately. In IBM i 5.4 and previous releases, the system did not use new processor features until all models supported by the current release of the operating system had the new feature. For more information, see the *Adaptive Code Generation* section in the *Advanced Optimization Techniques* chapter of *ILE Concepts*. *ILE Concepts* can be found on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/books/sc415606>

Programs that use the latest hardware features might require conversion when moved to a different system that is at the same or different (earlier or later) release level, but running on an older version of the hardware. The program might require conversion because it uses at least one feature not supported by the system to which the object is being restored. To restore a program without converting it, set the Force conversion on restore system value (QFRCCVNRST) to 0 before restoring the program and set the parameters of your restore command to not force conversion during restore.

For IBM i 6.1 information on object (including programs) conversions, see *IBM i5/OS Program Conversion: Getting Ready for i5/OS V6R1*, REDP-42933.

3.8 Program conversion restoring to a different release

The internal object format for program objects is different between IBM i 6.1 and earlier releases. When you move a program object (*MODULE, *PGM, *SRVPGM, *SQLPKG) between a release earlier than IBM i 6.1 and IBM i 6.1 or later, the system must convert the program object using the information that is stored with the program. If you are restoring an older format program onto IBM i 6.1 or later, the settings of the Force conversion on restore (QFRCCVNRST) system value and the Force object conversion (FRCOBJCNV) parameter on the restore command control whether the program object gets converted during the restore operation. If you set the QFRCCVNRST system value to 2 or greater, or use the FRCOBJCNV (*YES *ALL) parameter on the restore command, the program object gets converted on the restore.

In contrast, an IBM i 6.1-format program object is always converted when restored to a system running a release earlier than IBM i 6.1, regardless of the system value and restore command settings. For more information, refer to the *Restoring programs to a different release* topic in Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmrp_rel.htm

See also *IBM i5/OS Program Conversion: Getting Ready for i5/OS V6R1*, REDP-42933.

3.9 Restore user profile

A new parameter has been added to the Restore User Profiles (RSTUSRPRF) command, Omit security data (OMITSECDA). This new parameter is useful when merging systems.

Four special values are supported:

*NONE	No security information is omitted.
*AUTL	Authorization list (*AUTL) and authority holder (*AUTHLR) objects are omitted. However, for any of these objects that already exist on the system, any specific user authorities will be restored. You need to run the Restore Authority (RSTAUT) command to complete the restore of these authorities.
*DCM	The internal objects required by Digital Certificate Manager (DCM) are omitted.
*FCNUSG	Function usage information is omitted. However, for any function identifiers that already exist on the system, any specific user settings will be restored. You need to run the Restore Authority (RSTAUT) command to complete the restore of these settings.

Note: *AUTL, *DCM, and *FCNUSG can be used separately or in any combination on a single restore command.

3.10 User-defined file system backup

IBM i has supported user-defined file systems (UDFSs) for several releases in its embedded integrated file system support. A UDFS is a hierarchical directory structure similar to PC or UNIX®-based operating systems such as DOS, Microsoft® Windows, AIX, and, Linux, and general UNIX-based file systems.

A UDFS can be defined in any supported ASP, including an iASP, independent auxiliary storage pool (ASP). IBM i enables you to create and manage these file systems.

In a UDFS, as in the root (/) and QOpenSys file systems, you can create directories, stream files, symbolic links, local sockets, and *FIFO objects.

A single block special file object (*BLKSF) represents a UDFS. As you create UDFSs, you also automatically create block special files. The block special file is only accessible to the user through the integrated file system generic commands, APIs, and the QFileSvr.400 interface.

In the context of backup and recovery, there are two new IBM 6.1 enhancements for a UDFS.

3.10.1 Omitting objects during a save

You can use two methods to omit objects from an unmounted UDFS during a save operation: You can either use the *OMIT option on the OBJ parameter or the PATTERN parameter on the SAV command. By omitting certain objects, such as Domino® data or temporary objects, for example, you can reduce your backup window while saving an unmounted UDFS.

3.10.2 Saving a mounted user-defined file system

In previous releases, you had to unmount the user-defined file systems (UDFSs) before performing a save operation if you wanted to save both the UDFS information and the objects in the UDFS. If you saved a mounted UDFS, only the objects in the UDFS were saved, but not the UDFS information. In this release, both the UDFS structure and objects within the UDFS are saved. You can now restore both the UDFS information and the objects in the UDFS.

3.11 Restore actions for mounted user-defined file systems

In IBM i 6.1, different actions occur automatically during the restore operation when RBDMFS(*UDFS) is specified on the RST command. The UDFSs might not be mounted, or might be mounted over the same directory or a different directory, depending on the status of the directory being restored to. Table 3-1 provides the restore actions for mounted user-defined file systems.

Table 3-1 Restore actions for mounted user-defined file systems

Directory being restored over	UDFS that was mounted over the directory when saved			
	UDFS does not exist	UDFS not mounted	UDFS mounted over same directory	UDFS mounted over different directory
Directory being restored over does not exist	<ul style="list-style-type: none"> ▶ Create directory ▶ Create UDFS ▶ Mount UDFS ▶ Restore tree 	<ul style="list-style-type: none"> ▶ Create directory ▶ Mount UDFS ▶ Restore tree 	Not applicable	The restore operation fails with message CPD37B8
No UDFS is mounted over the directory	<ul style="list-style-type: none"> ▶ Create UDFS ▶ Mount UDFS ▶ Restore tree 	<ul style="list-style-type: none"> ▶ Mount UDFS ▶ Restore tree 	Not applicable	The restore operation fails with message CPD37B8
Same UDFS is mounted over the directory	Not Applicable	Not Applicable	Restore tree ^a	Not Applicable
Different UDFS is mounted over the directory	<ul style="list-style-type: none"> Create UDFS Mount UDFS Restore tree 	<ul style="list-style-type: none"> ▶ Mount UDFS ▶ Restore tree 	Not Applicable ^a	The restore operation fails with message CPD37B8

a. Suppose the target UDFS is mounted over the target directory and one or more UDFSs are mounted over it. This situation is treated as if the target UDFS is mounted over a different directory. The restore of that mount point will fail with a message that indicates that the UDFS is in use.

For further information, refer to the *Restore actions for mounted user-defined file systems* topic in Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmre/storemountedudfs.htm?tocNode=int_184095

3.12 Encrypted backups

With IBM i 6.1 you can use software-implemented encrypted data backups to tape media to prevent the loss of personal customer information or confidential data if the media is lost or stolen. There are two methods for performing encrypted backups:

- ▶ Hardware encryption using an encryption-capable tape drive. You can use native save/restore commands or Backup Recovery and Media Services (BRMS) with the encrypting tape drive (OS/400 V5R2 or later).
- ▶ Software encryption using BRMS (IBM i 6.1 or later).

Consider the following factors in making your decision regarding your encryption media and method:

- ▶ Choose the hardware encryption method using an encrypting tape drive if you want the best performance for doing save and restore operations, especially a full-system save or restore operation. You do not need host-based encryption of data or the use of specialized encryption appliances to use the encrypting tape drive.
- ▶ Choose the software encryption method if you want a low-cost solution. This solution is ideal for backing up individual objects that contain customers' personal information or confidential data. Customers with sufficient system resources and a large enough backup window also can encrypt the backup without impacting their business. You can use any tape drive or tape library model with software encryption. However, the performance is not as good as using hardware encryption.

Refer to *Security Guide for IBM i V6.1*, SG24-7680, for more information about data encryption on tape media.

3.13 Encrypted auxiliary storage pools

With IBM i 6.1, disk encryption allows you to encrypt data stored in basic user auxiliary storage pools (ASPs) and independent ASPs (iASPs).

Disk encryption protects data from a number of different threats:

- ▶ Protects data transmission to and from the disk drive (important in a SAN environment).
- ▶ Protects data transmission in the cross-site mirroring environment (only when the data being mirrored is on an encrypted independent ASP).
- ▶ Protects data in the case of theft of the disk drive.
- ▶ Protects data in the case of return or resale of a disk drive (reduces the need to sanitize the disk drive).

In order to use encryption, the system or partition must have 5761-SS1 option 45 - Encrypted ASP Enablement installed. The option enables encryption when a new user disk pool or independent disk pool is created.

Data stored in encrypted ASPs can be saved in the same way as data in unencrypted ASPs. However, if the data in the system ASP or independent ASP is lost, you need to perform additional recovery steps.

Data is encrypted only while it resides on the ASP. When you read the data, it is decrypted. When doing a save operation, the data is decrypted as it is read for the save operation. The data is encrypted on the save media only if you are doing an encrypted backup using either an encrypting tape drive or the software solution.

Refer to *Security Guide for IBM i V6.1*, SG24-7680, for more information about saving and restoring encrypted data in ASPs and iASPs.

3.14 Restore the operation output file

Several restore commands and the Restore Object List (QSRRSTO) API use the QASRRSTO (restore) output file. New fields have been added to the model file to support the

new Private authority and Deferred restore functions. For more information, see *Restore operation output file information* on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmrestoreopoutputfile.htm>

3.15 Restoring external stored procedures

If you move data to a new system, you might need to restore external stored procedures. For more information, refer to the IBM Techdoc *Stored Procedures Restore Overview* located at:

http://www-01.ibm.com/support/docview.wss?rs=0&dc=DB520&dc=D900&dc=D800&dc=DA900&dc=DA800&q1=Restoring+External+Stored+AND+AS400KBXXYYZZRCH&uid=nas10bf6be811065714f8625707e0073a26b&loc=en_US&cs=UTF-8&lang=all

3.16 Backup and restore for IBM i on BladeCenter

Backing up and restoring an IBM i client logical partition on a BladeCenter blade server works differently than the typical procedure for backing up and restoring a logical partition.

IBM i does not have direct access to a physical tape drive for backup and restore. This changes the save/restore process as a whole, although the steps performed in IBM i remain the same (for the most part).

Performing an IBM i backup is a 2-stage process on the Power blade:

- ▶ First, a standard IBM i save command or BRMS is used to perform a save on a writable optical device that contains an optical volume. The writable optical device is a file-backed virtual optical drive created in VIOS and the optical volume is a virtual media image, which is a file on disk in VIOS.
- ▶ Second, the virtual media image, containing the IBM i save, is written out to a Serial Attached SCSI (SAS) attached tape drive using a VIOS command.

Performing a restore follows the same 2-stage process in reverse:

- ▶ First, the virtual media image file is restored from the SAS tape drive onto a VIOS disk using a VIOS command. The image file is then mounted on the correct virtual optical drive assigned to the IBM i partition and becomes available as a volume from which to restore.
- ▶ Second, a standard IBM i restore is performed from the volume using an IBM i restore command or BRMS. A full system restore can be executed by first performing a D-mode IPL from the virtual optical image, provided the image contains a full system save performed previously.

For more information on these backup and recovery considerations, see these topics in the Infocenter:

Restoring logical partitions at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmrstlpar.htm#rzarmrstlpar>

i5/OS client partition considerations at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahc/rzahcb1ade5osvirtualpartitionsoverview.htm>

i5/OS on POWER® Blade Read-me First at:

http://www.ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf

3.17 Retrieve Disk Information (RTVDSKINF)

New functionality was added to the Retrieve Disk Information (RTVDSKINF) command in IBM i 6.1. Some of the Reclaim Storage (RCLSTG) command's functions were implemented on the RTVDSKINF command.

The command will now also check that certain objects are addressable through a library. If the object is not attached to a library, it will be moved to the proper library or the QRCL library (for objects in the system or basic auxiliary storage pools) or library QRCLnnnn (for objects in an independent auxiliary storage pool where nnnn is the number associated with the independent ASP). Objects found to be damaged or unusable will be deleted. A message will be sent for each object deleted by RTVDSKINF.

For more information about how RTVDSKINF recovers objects, refer to the *Reclaim Storage (RCLSTG) command* topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzarm/rzarmrc1stg.htm>

Note: RTVDSKINF only tries to recover a subset of the object types that a RCLSTG tries to recover, so the new RTVDSKINF function does not replace the function of RCLSTG.

3.18 Backup, Recovery and Media Services

The following sections describe the Backup, Recovery and Media Services (BRMS), including the graphical interface, command and display enhancements, recovery enhancements, media services enhancements, hierarchical storage manager (HMS) enhancements, device enhancements, network enhancements, install enhancements, and others.

3.18.1 BRMS graphical interface

This section describes the functions that were changed or added to the BRMS graphical interface.

Note: BRMS graphical interface refers to both the System i Navigator BRMS plug-in and the IBM Systems Director Navigator for i5/OS Web environment.

The corresponding IBM i commands and displays have also been enhanced to enable these new features but are not shown in the following sections. Refer to 3.18.2, "Command and display enhancements" on page 52 for information about the command and display changes.

IBM Systems Director Navigator for i

IBM Systems Director Navigator for i is a Web console interface for IBM i administration where you can work with the Web-enabled tasks of System i Navigator.

For more detailed information about IBM Systems Director Navigator for i in general, refer to Chapter 16, “IBM Systems Director Navigator for i” on page 403.

To get to IBM Systems Director Navigator for i, access the following URL from a Web browser where hostA is your System i name:

http://hostA:2001

You will be presented with a Web page similar to that shown in Figure 3-1.

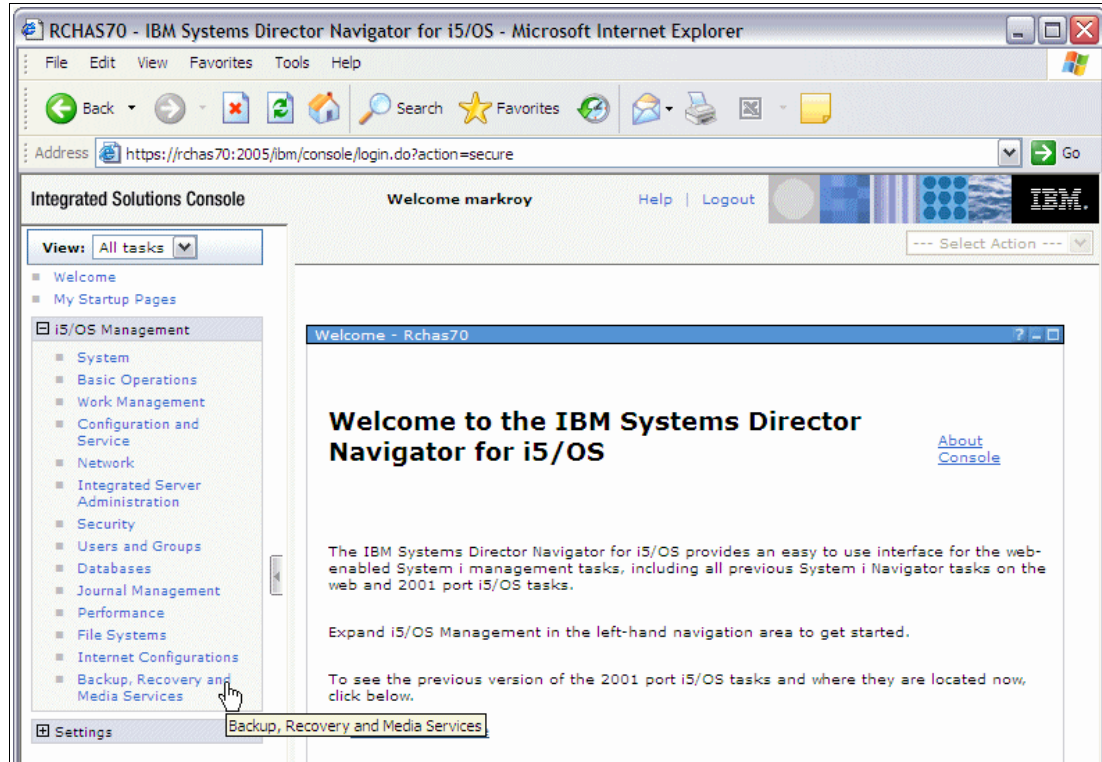


Figure 3-1 Systems Director Navigator for i5/OS welcome page

On the welcome page, select the Backup, Recovery and Media Services link in the left pane. You will be presented with the Backup, Recovery, and Media Services for i5/OS page. Select the Advanced button to view the list of available tasks, as shown in Figure 3-2.

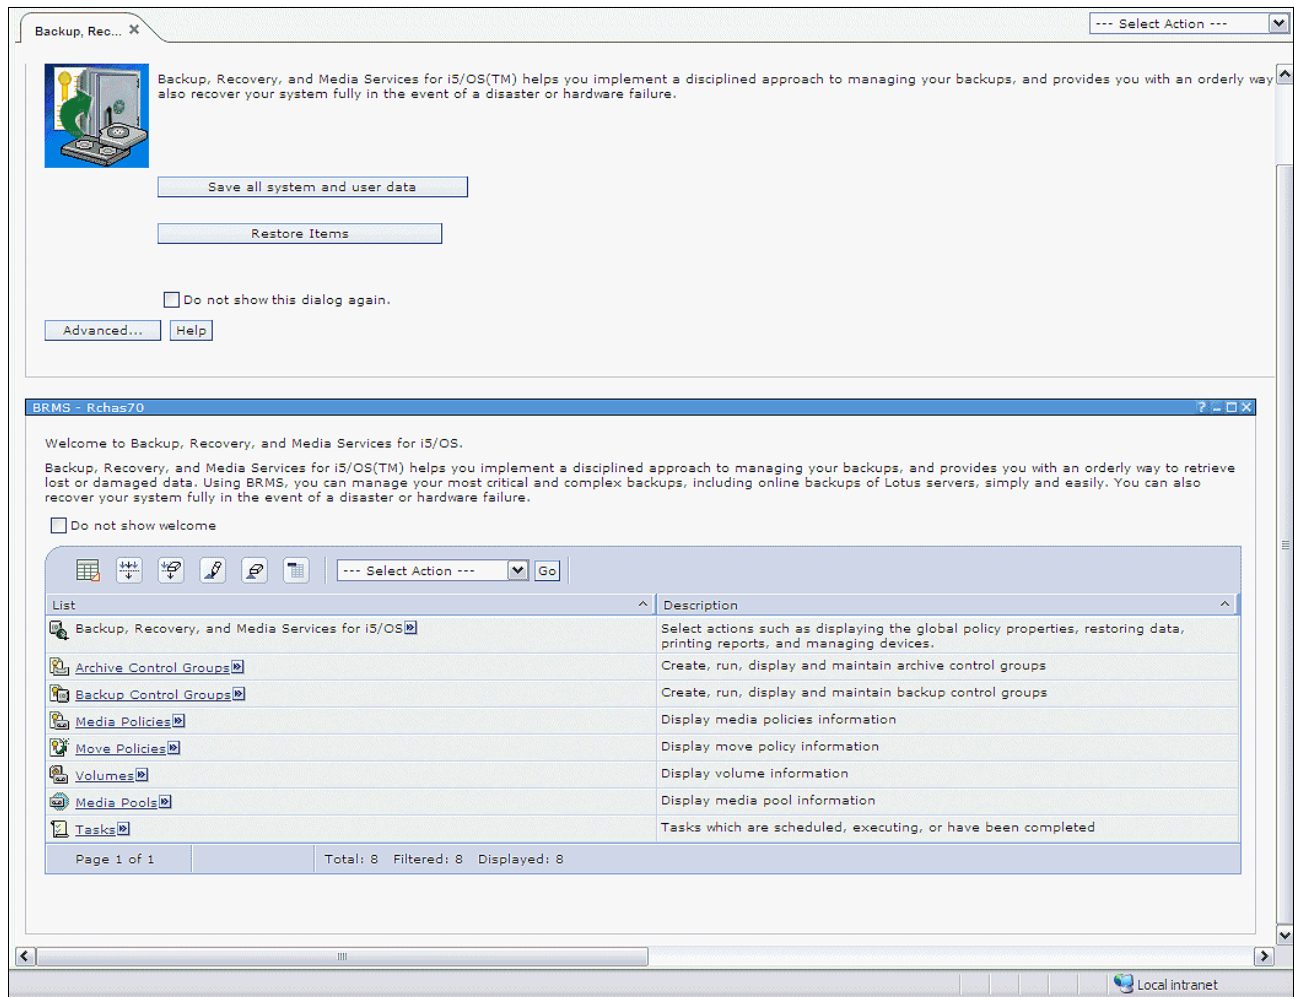


Figure 3-2 BRMS tasks in IBM Systems Director Navigator for i

Note: When the HTTP ADMIN servers are started (automatically through the STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN) command), Systems Director Navigator for i5/OS is also started. During its startup processing, it detects the presence of BRMS and the Backup, Recovery and Media Services link is added to the i5/OS Management links.

Media policies

Previous releases hid media policies. The ability to create and manage media policies has now been added. They can be created, displayed, edited, and added to Control Groups.

Note: In this section the media policies are shown using both System i Navigator and Systems Director Navigator for i5/OS. To minimize the size of this publication we show only the System i Navigator windows in the following sections. All these functions are also available in Systems Director Navigator for i5/OS.

Figure 3-3 shows the media policies list available on System i Navigator.

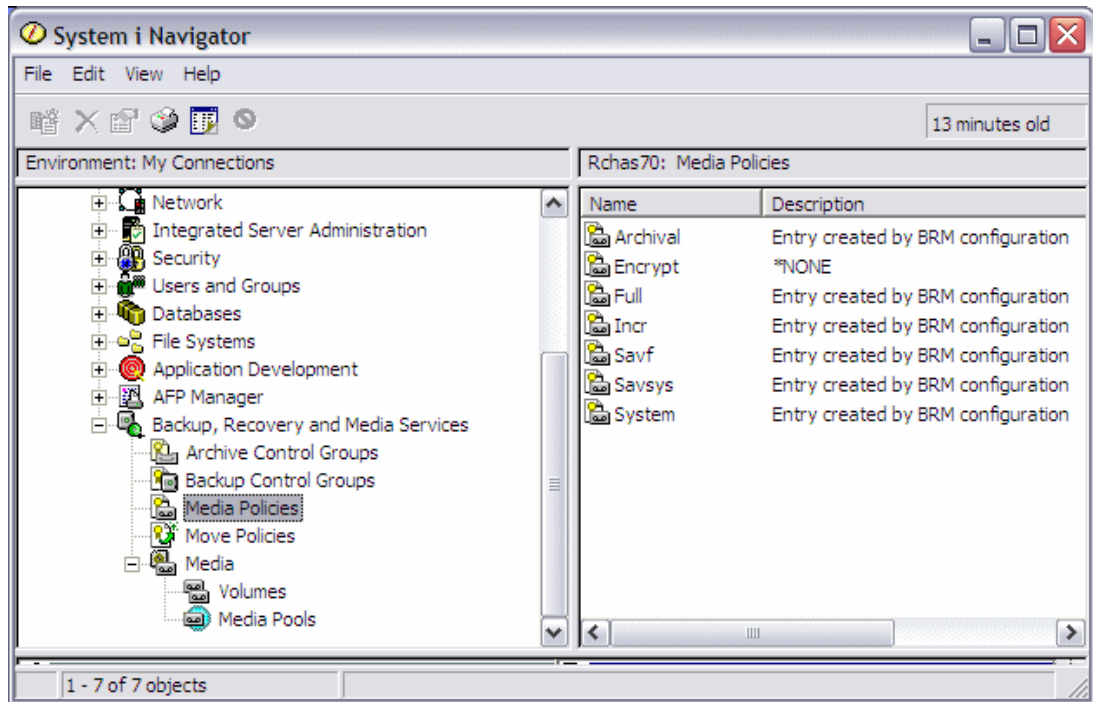


Figure 3-3 Media policies in System i Navigator

Figure 3-4 shows the media policies list available in Systems Director Navigator for i5/OS.

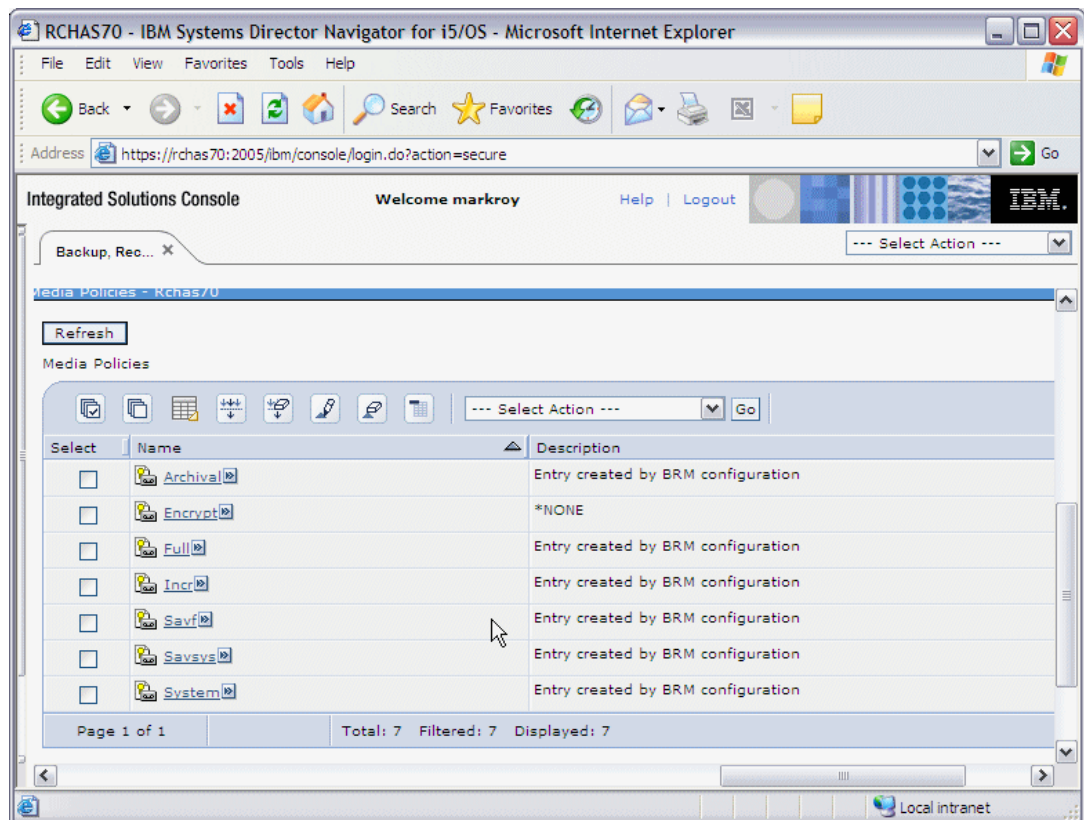


Figure 3-4 Media policies in Systems Director Navigator for i5/OS

DVD, CD-ROM, optical libraries, and virtual optical devices

These devices are now supported. Figure 3-5 shows a system with DVDs and a virtual optical device.

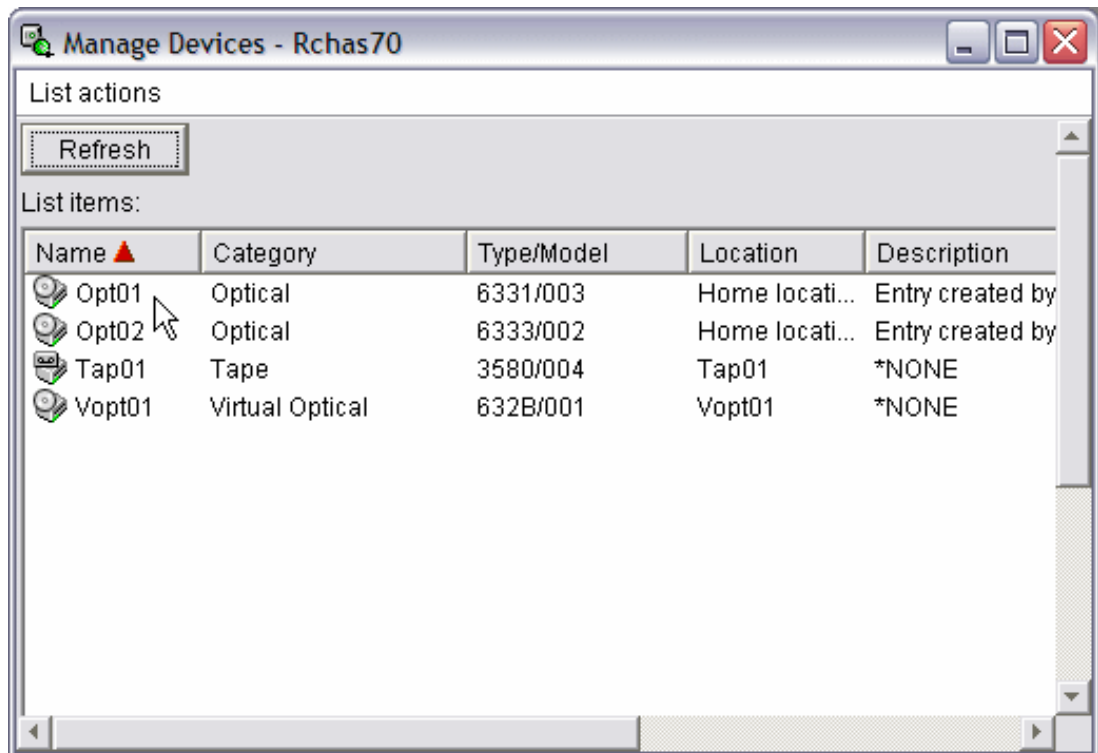


Figure 3-5 Managing DVD, CD-ROM, optical libraries, and virtual optical devices

Object level backup of Linux or Windows guest operating systems

Prior to IBM 6.1, file and object level backups of Windows guest operating systems was available through the IBM i command interface. This support allows saves and restores of individual objects residing on the Windows or Linux server through the QNTC file system. For example:

```
SAV DEV('/QSYS.LIB/TAP01.DEVD') OBJ('/QNTC/servername/sharename')
```

IBM i 6.1 introduced the same functionality for Linux guest operating systems.

Additionally, the BRMS GUIs have enhanced function in 6.1 to perform file and object level backups of both operating systems, as an alternative to the command interface.

Save-while-active

IBM i 6.1 introduces save-while-active for guest operating systems, as described in 3.2.2, “Saving storage spaces for active integrated servers” on page 33. The BRMS GUIs have been updated to include this new function, shown in Figure 3-6.

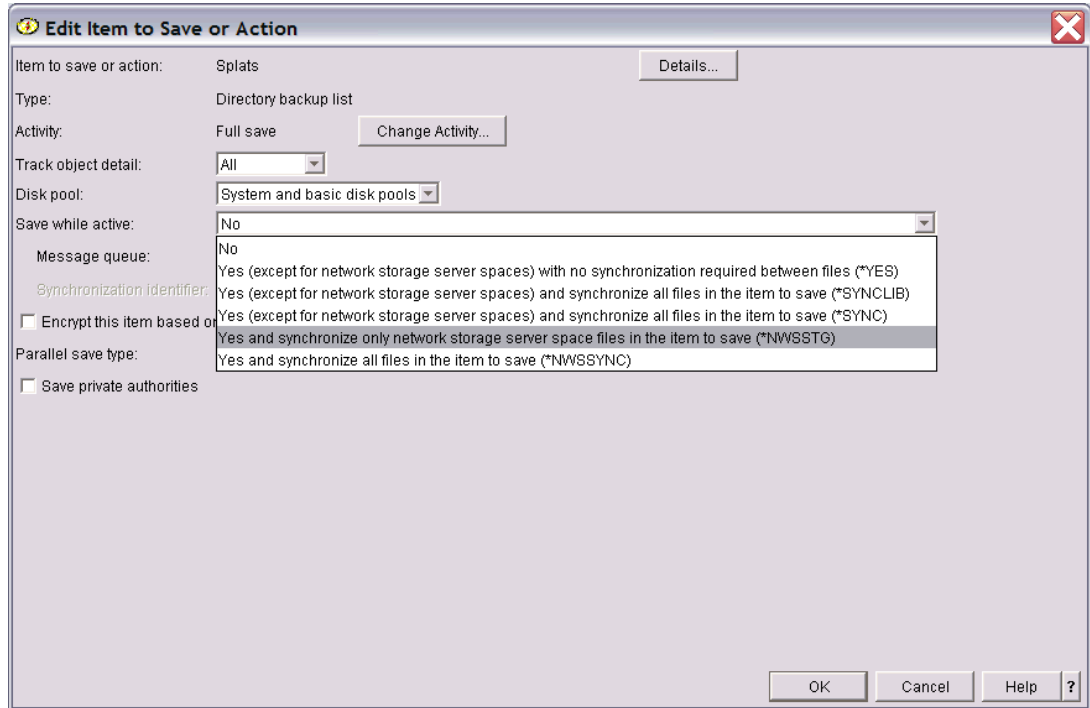


Figure 3-6 Save-while-active: Edit Item to Save or Action

Naming changes

Backup policy and archive policy have been renamed as backup control group and archive control group, to align them with the nomenclature used in the IBM i command interface to BRMS.

Managing and selecting media

The ability to manage and select media policies from the Backup Control Group and Archive Control Group Wizards has been added. These options are also available from the Properties page. Figure 3-7 shows the System i Navigator properties for a control group called Test.

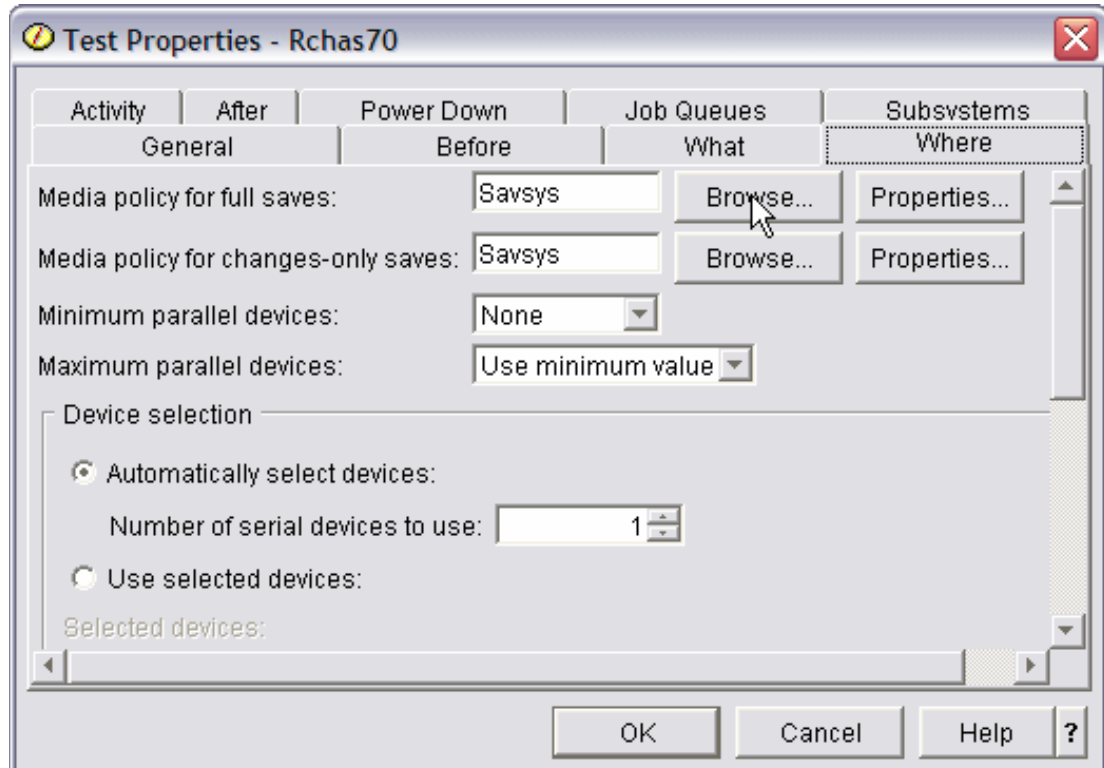


Figure 3-7 Control group properties

Using the Browse button presents a window of available policies that you can select from. The Properties button displays the properties of the selected media policy.

Saving and restoring private authorities

IBM i 6.1 provides an option to save the private authorities to an object when the object is saved. They can also be restored if they were previously saved. For more information about this new support, refer to 3.4, “Saving and restoring private authorities” on page 35. The example in Figure 3-8 shows where to change the setting for an object in its control group properties.

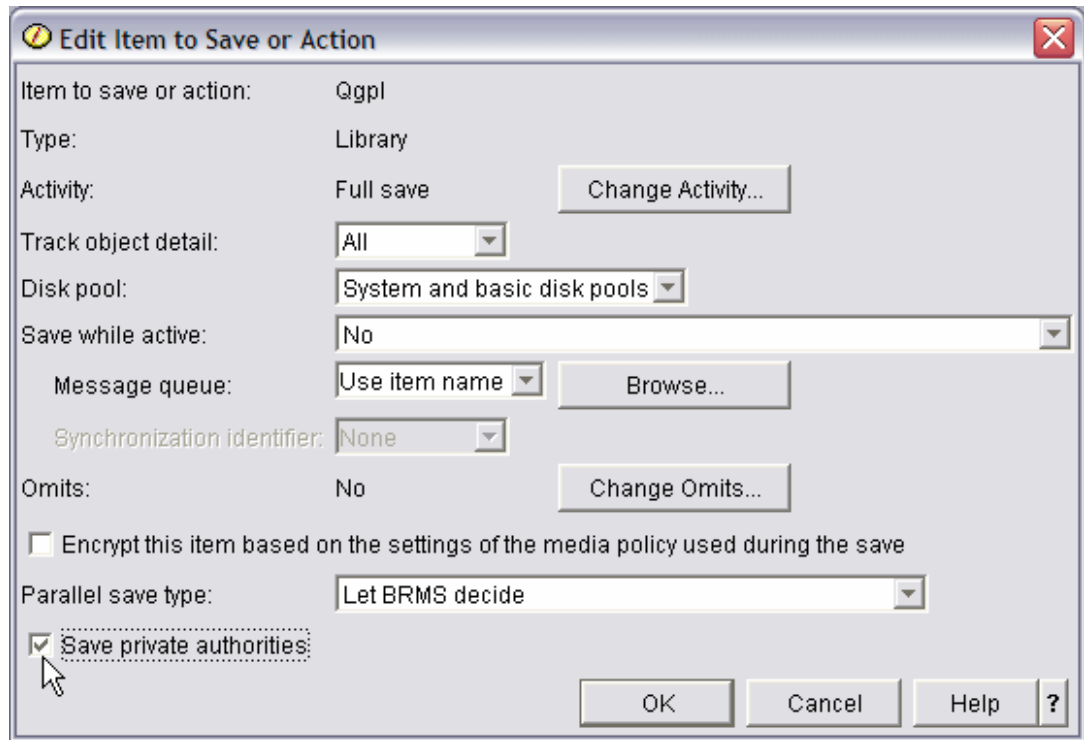


Figure 3-8 Changing the “Save private authorities” of an object in a BRMS Control Group

Disk pools and generic libraries

Backup control group and archive control group properties now support disk pools (auxiliary storage pools) and generic libraries.

Encryption

You can encrypt backups to tape media to prevent the loss of personal customer information or confidential data if the media is lost or stolen. There are two methods for performing encrypted backups:

- ▶ Hardware encryption using an encryption-capable tape drive. This type of encryption is performed without the knowledge of IBM i. You can use native save/restore commands or Backup Recovery and Media Services with the encrypting tape drive.
- ▶ Software encryption using Backup Recovery and Media Services

For more information, refer to:

- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ *Backup, Recovery, and Media Services for i5/OS* document on Infocenter at:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/books/sc415345.pdf>

Global policy properties

You can specify global settings for your system through the global policy. These settings are used, in conjunction with values that you specify in your control groups, to enable you to control how your saves are done.

Global policies, at the system level, have been extended in 6.1 to include support for:

- ▶ Restore operations from alternate media
- ▶ Use of SNA networking
- ▶ Management of restricted state interfaces

SNA networking and restricted state interfaces are on the Network tab as shown in Figure 3-9.

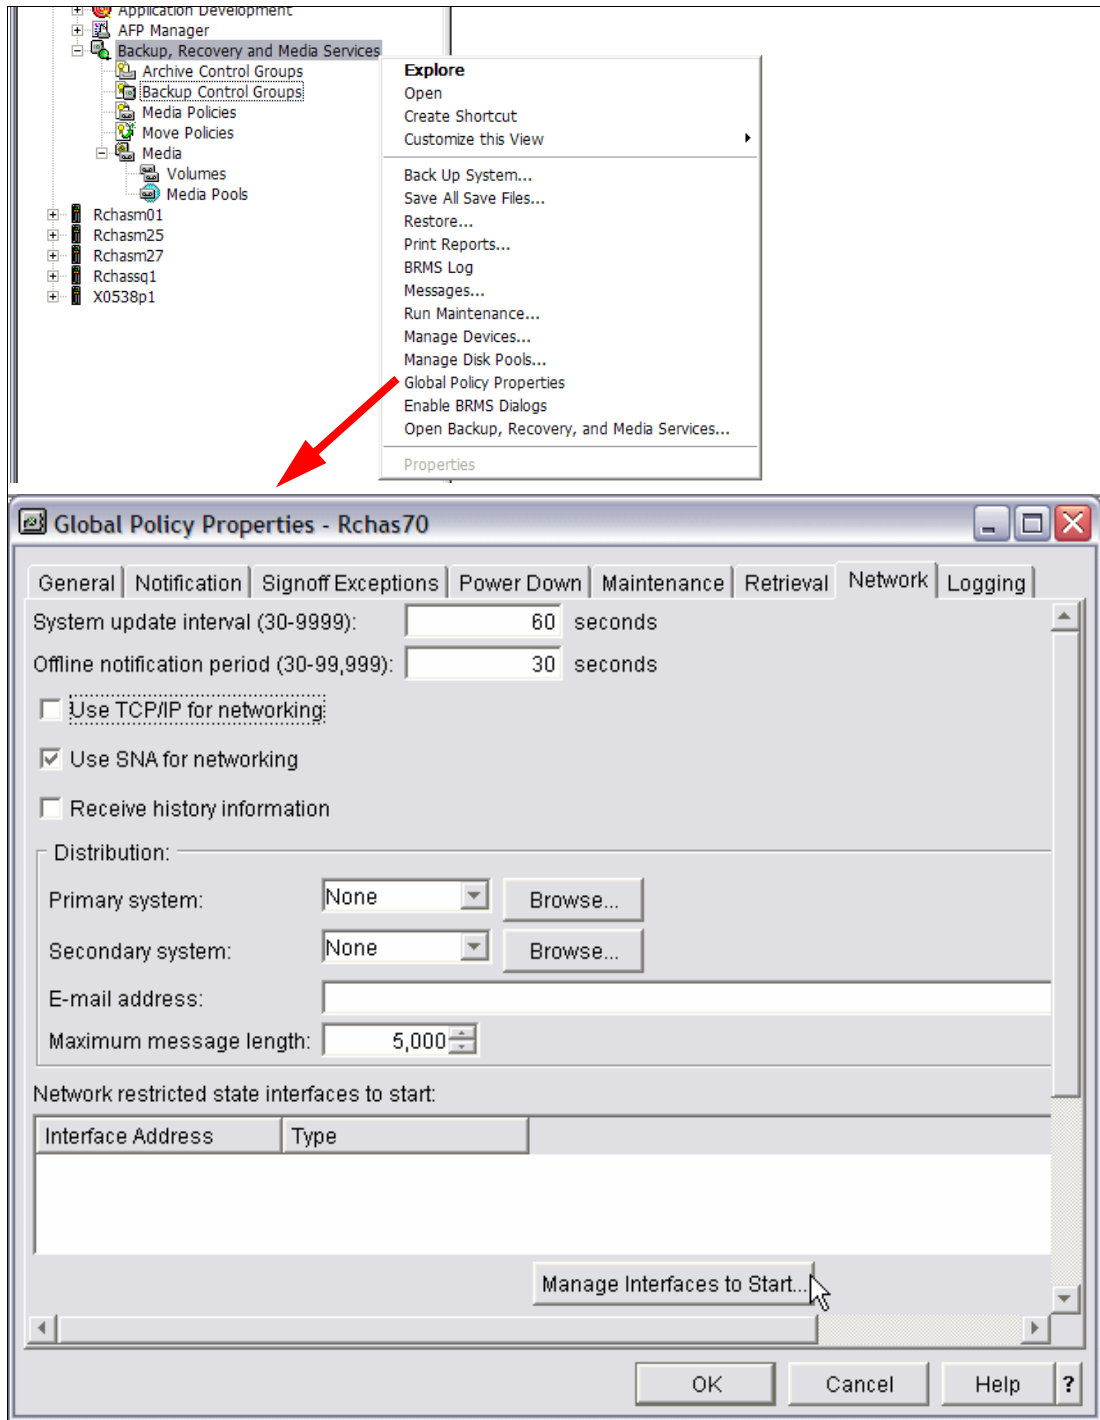


Figure 3-9 Global policy properties

In the Global Policy Properties dialog, under the General tab there is a check box that indicates that all restores are from an alternate media..

Balancing media ownership

Media ownership can be balanced to match the required number of media needed by systems within a BRMS network (network systems must be running IBM i 6.1 or later and have 5761BR1, Option 1, BRMS - Network Feature installed). The examples shown in the following figures depict how to access this feature and the options that are available. Figure 3-10 shows how to get the Manage Volume Balancing option.

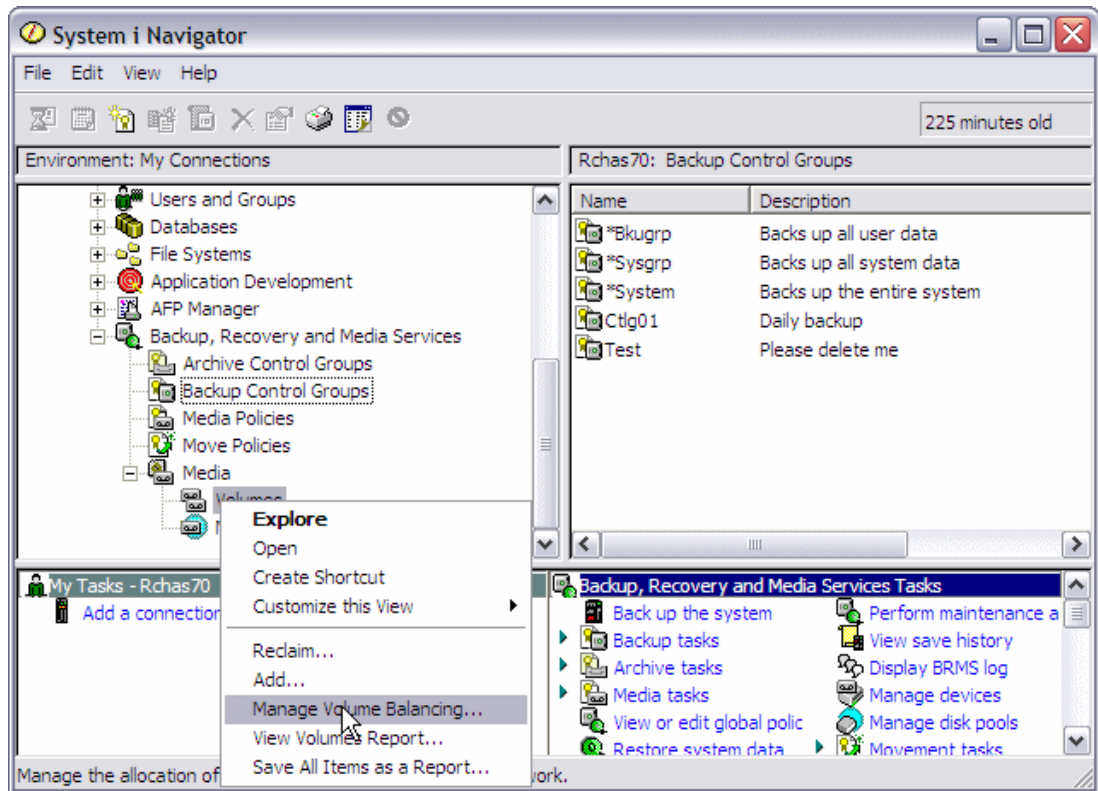


Figure 3-10 Accessing volume balancing in BRMS

Figure 3-11 shows all actions available on the Manage Volume Balancing option.

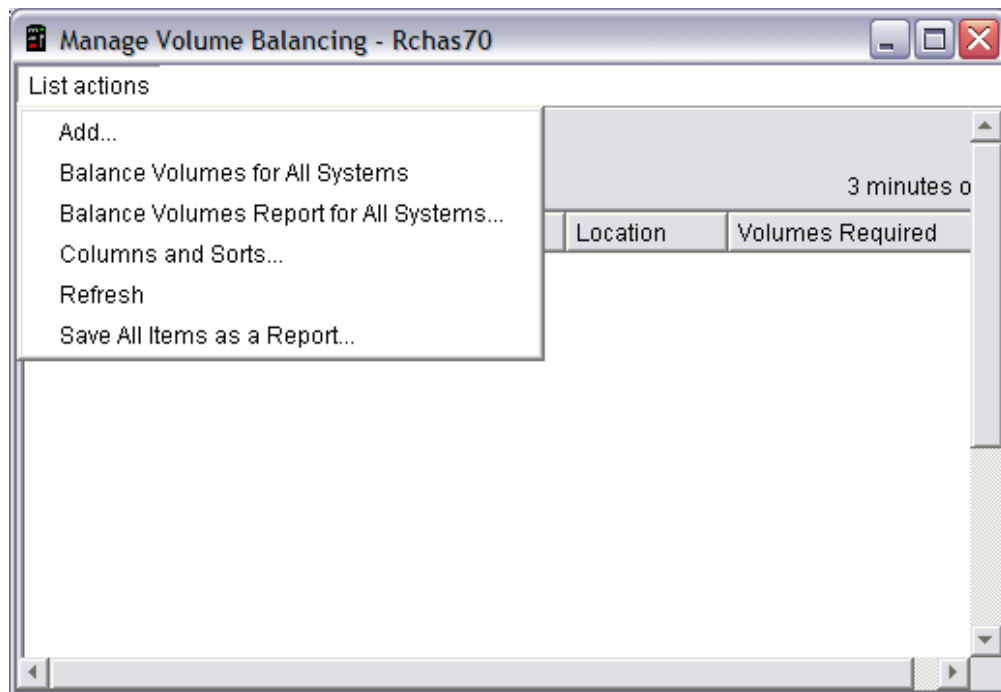


Figure 3-11 Volume balancing actions in BRMS

The Balance Volumes for All Systems action will include media from all systems in the BRMS network. It produces a report that shows the scratch media currently available as well as the scratch media requirements for BRMS operations. It can be used to determine what actions will be taken when a balance is performed.

3.18.2 Command and display enhancements

The following discusses the command and display enhancements.

Print Report using BRM

In 6.1, BRMS introduced a new command, Print Report using BRM (PRTRPTBRM).

The command provides the ability to produce a backup statistics report. The report provides summarized information about all objects saved by BRMS within a specified period. The information provided includes:

- ▶ Control Group
- ▶ Start Date
- ▶ Start Time
- ▶ End Date
- ▶ End Time
- ▶ Duration
- ▶ Number of Objects Saved
- ▶ Number of Objects Not saved
- ▶ Save Size (MB)
- ▶ Save GB/HR
- ▶ Media Class
- ▶ Volume
- ▶ Save Status

Encrypting backups or archives

BRMS now offers the ability to encrypt your data to a tape device.

This encryption solution is hardware-independent, meaning no need for any encryption device. To use the encryption function, you need to have the BRMS Advanced feature (5761-BR1 Option 2) and Cryptographic Service Provider (5761-SS1 Option 44) installed on the operating system.

Note: IBM i supports Library Managed Encryption (LME), sometimes also referred to as Transparent Encryption. With LME, the encrypting tape (LTO 4 or 3592 E05) must be in a library such as the 3584, 3577, 3576, or 3573 for encryption to be available. The library and drive work together with the required Encryption Key Manager (EKM) component that is available on the hardware to provide data encryption without any host involvement. Essentially, neither IBM i nor BRMS is aware of encryption or decryption being performed.

This is the best solution compared to software encryption (IBM i and BRMS) on System i, because there is no CPU utilization consumed by this hardware solution. However, feedback from small to medium-sized customers indicates that the hardware encryption solution is currently cost prohibitive. Thus, IBM i 6.1 provides software-based encryption support for backup under BRMS.

If you have the appropriate master key in the IBM i directory structure in your IBM i partition's keystore file Q1AKEYFILE in library QUSRBRMS, the basic IBM i restore commands detect the encryption information on the tape media being restored. Assuming no other object incompatibility or security constraints, the restore will complete successfully.

Figure 3-12 shows an example of the encryption parameters in the user-created Encrypt policy.

```
Change Media Policy

Media policy . . . . . : SAVSYS

Type choices, press Enter.

Encrypt Data . . . . . *NO      *NO, *YES
Key store file . . . . . Q1AKEYFILE Name
Key store library. . . . . QUSRBRM  Name
Key record label . . . . .

F3=Exit  F5=Refresh  F12=Cancel

Bottom
```

Figure 3-12 Change Media Policy - encryption

For more information about data encryption on tape media, refer to *Security Guide for IBM i V6.1*, SG24-7680 and the *Backup, Recovery, and Media Services for i5/OS* document on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/books/sc415345.pdf>

Generic libraries

The Save Object using BRM (SAVOBJBRM) command now supports generic library names on the Library parameter. This can be used to save the same object(s) from multiple libraries in the save operation. Figure 3-13 shows SAVOBJBRM command parameters.

Save Object using BRM (SAVOBJBRM)		
Type choices, press Enter.		
Library		Name, generic*
Object		Name, generic*, *ALL
	+ for more values	
Device		Name, *NONE, *MEDCLS
	+ for more values	
Object type	*ALL	*ALL, *ALRTBL, *BNDDIR, ...
	+ for more values	
Member	*ALL	Name, generic*, *ALL
Media policy		*SYSPCY, *NONE, ARCHIVAL, E..
Parallel device resources:		
Minimum resources	*NONE	1-32, *NONE, *AVAIL
Maximum resources	*MIN	1-32, *MIN, *AVAIL
Save active	*NO	*NO, *LIB, *SYSDFN, *SYNCLIB

Figure 3-13 Save Object using BRM (SAVOBJBRM)

Ability to save private authorities at an object level

The SAV nnn BRM commands have a new Private authorities (PVTAUT) parameter added to allow saves of private authorities. Figure 3-14 shows an example of the SAVBRM command where the Private authorities parameter is set to *NO.

Save Object using BRM (SAVBRM)		
Type choices, press Enter.		
Save active message queue . . .	*NONE	Name, *NONE, *WRKSTN
Library	*LIBL	Name, *LIBL, *CURLIB
Use optimum block size	*BKUPCY	*BKUPCY, *DEV, *NO, *YES
Type of save	*FULL	*FULL, *CUML, *INCR
End of media option	*REWIND	*REWIND, *LEAVE, *UNLOAD
Sequence number	*END	1-16777215, *END
Target release	*CURRENT	*CURRENT, *PRV, VxRxMx
Clear	*NONE	*NONE, *AFTER, *ALL
Object pre-check	*NO	*NO, *YES
Data compression	*DEV	*DEV, *NO, *YES
Data compaction	*DEV	*DEV, *NO
Private authorities	*NO	*NO, *YES

Figure 3-14 Save Object using BRM (SAVBRM)

Overriding backup control group settings

New parameters on the Start Backup using the BRM (STRBKUBRM) command can be used to override the device, parallel resources, media class, and move policy that is set in a control

group at the time when running the backup control group. Figure 3-15 shows the new parameters available on the STRBKUBRM command.

```

                                Start Backup using BRM (STRBKUBRM)

Type choices, press Enter.

Device . . . . . *CTLGRPATR   Name, *CTLGRPATR, *MEDCLS
      + for more values
Parallel device resources:
  Minimum resources . . . . . *CTLGRPATR   1-32, *CTLGRPATR, *NONE...
  Maximum resources . . . . . *MIN         1-32, *MIN, *AVAIL
Media class . . . . . *CTLGRPATR   *CTLGRPATR, *SYSPCY, SAVSYS..
Move policy . . . . . *CTLGRPATR   *CTLGRPATR, *NONE, OFFSITE
Omits . . . . . *PROCESS        *PROCESS, *IGNORE
  
```

Figure 3-15 New parameters on Start Backup using BRM (STRBKUBRM)

Parallel type of entries in a control group

The Edit Backup Control Group Entries display has been changed to allow the selection of the type of parallel save to run with a specific entry in a control group. Figure 3-16 shows an example for Edit Backup Control Group Entries.

```

                                Edit Backup Control Group Entries                                RCHAS70

Group . . . . . : TEST
Default activity . . . . .
Text . . . . . ITS0 Residency Test

Type information, press Enter.

      Backup   List Parallel   Private
Seq  Items    Type  Type      Authorities  Encrypt

  10 *EXIT
  20 QGPL          *PARALLEL *NO
  30 QRPG          *SERIAL  *NO
  40 QUSRSYS      *DEFAULT *NO
  50 *EXIT
  
```

Figure 3-16 Edit Backup Control Group Entries

FlashCopy

BRMS now supports the IBM Storage Solutions external disk FlashCopy® environment. You use FlashCopy to create and manage FlashCopy logical drives. FlashCopy provides an almost instant, point-in-time copy of a logical volume (LUN) within independent disk pools (IASPs) on external storage, which can reduce the time it takes to complete daily backups. Point-in-time copy functions give you an instantaneous copy, or view, of what the original data looked like at a specific point in time. The target copy is totally independent of the source independent disk pool and is available for both read and write access once the FlashCopy command has been processed. FlashCopy uses a bitmap to represent the relationship between a source volume (LUN) and target volume (LUN). This mechanism enables a

FlashCopy operation to complete in a few seconds. The target volume is available for use immediately after the FlashCopy operation has completed.

Changes have been made to the Initialize BRMS (INZBRM) command, the Change System Policy display, and the Change Network Group display.

On the Initialize BRMS (INZBRM) command, the Option (OPTION) parameter can be used to access the new State (STATE) parameter. The State parameter allows you to set the FlashCopy state. The help text for the State parameter provides more detail on the four State special values. Figure 3-17 shows the INZBRM command panel.

```

                                Initialize BRMS (INZBRM)

Type choices, press Enter.

Option . . . . . > *FLASHCOPY   *DATA, *CHGSYSNAM, *DEVICE,..
State  . . . . .                *ENDBKU, *ENDPRC, *STRBKU...

.....
:                               State (STATE) - Help                               :
:                                                                                   :
:      information will appear as if the backup had been                          :
:      performed on the production system.                                         :
:                                                                                   :
: *STRPRC                                                                           :
:      Start the FlashCopy process on this production system.                      :
:      This notifies BRMS that the production system's data                       :
:      is being copied via ESS FlashCopy and the backup is                        :
:      being performed on the clone system.This step is                          :
:      required prior to performing the ESS FlashCopy                            :
:      function. During this state, there should be no BRMS                      :
:                                                                                   :
:                                                                                   :
:      F2=Extended help   F10=Move to top   F11=Search Index                       :
F3=Exit  F4= : F12=Cancel   F20=Enlarge    F24=More keys                         :
F24=More keys :
:.....

```

Figure 3-17 Initialize BRMS (INZBRM)

Missed object policies

You can have BRMS record library objects and members or directories and files that could not be saved when running backup control groups. BRMS records the unsaved objects in a missed object policy, which is nothing more than another backup control group containing one or more backup object lists or backup link lists containing these missed objects. You can then use the missed object policy to review which objects did not get saved, or to run the missed object policy in an attempt to re-save the missed objects.

Note: Missed object policies do not support folders and documents or spooled files.

Missed object policies are defined in Control Groups, as shown in Figure 3-18.

```

Change Backup Control Group Attributes

Group . . . . . : *SYSTEM

Type information, press Enter.

IPL after backup . . . . . *BKUPCY      *YES, *NO, *BKUPCY
  How to end . . . . . *BKUPCY      *CNTRLD, *IMMED, *BKUPCY
  Delay time, if *CNTRLD . . . . . *BKUPCY      Seconds, *NOLIMIT
  Restart after power down . . . . . *BKUPCY      *YES, *NO, *BKUPCY
  IPL source . . . . . *BKUPCY      *PANEL, A, B, *BKUPCY
  IPL restart type . . . . . *IPLA      *FULL, *IPLA, *SYS ...

Save active wait time:
  Object locks . . . . . 120          0-99999, *NOMAX
  Pending record changes . . . . . 120      0-99999, *NOCMTBDY,
*NOMAX
  Other pending changes . . . . . 120      0-99999, *NOMAX
Backup item exit program . . . . . *BKUPCY      Name, *NONE, *BKUPCY
  Exit program library . . . . .          Name
  Exit program format . . . . .          BKUI0100
Missed object policy . . . . . *BKUPCY      Name, *NONE, *BKUPCY
  Clear missed objects . . . . . *NO      *NO, *YES

F3=Exit  F12=Cancel

More...

```

Figure 3-18 Change Backup Control Group Attributes - missed objects policy

Save-while-active

The save-while-active function has been available for several releases, with each new release containing enhancements. In IBM i 6.1, support was added for creating a single synchronization point for multiple save commands. Refer to 3.2.1, “Single synchronization point for multiple save commands” on page 32 for more details about this feature. Various changes have been made to BRMS to access the new function. The changes include:

- ▶ The Monitor Save While Active (MONSWABRM) command

Synchronization ID (SYNCID), Number of operations (NUMSYNC) and Start save wait time (STRSAVWAIT) parameters have been added, to start synchronized checkpoint for more than one save-while-active operation.

Figure 3-19 shows the options available on the MONSWABRM command.

```

Monitor Save While Active (MONSWABRM)

Type choices, press Enter.

Saved library . . . . . >           Name, *MSGQ
Command to run . . . . . >

Job description . . . . . *USRPRF   Name, *USRPRF
  Library . . . . .           Name, *LIBL, *CURLIB
Job queue . . . . . *JOBDB         Name, *JOBDB
  Library . . . . .           Name, *LIBL, *CURLIB
Time limit in seconds . . . . . 3600 1-999999, *NOMAX
Multiple save synchronization:
  Synchronization ID . . . . . *NONE   Name, *NONE
  Number of operations . . . . . 2     2-32
  Start save wait time . . . . . 600   1-99999, *NOMAX

```

Figure 3-19 Monitor Save While Active (MONSWABRM)

- ▶ Edit Backup Control Group Entries display

Special values *SYNC, *NWSSTG, and *NWSSYNC added for the Save While Active field.

- *SYNC

Objects can be saved and used at the same time. All of the object checkpoints occur at the same time.

- *NWSSTG

Objects can be saved and used at the same time. The object checkpoints can occur at different times. Allows network server storage spaces in directory /QFPNWSSTG to be saved when they are active.

- *NWSSYNC

Objects can be saved and used at the same time. All of the object checkpoints occur at the same time. Allows network server storage spaces in directory /QFPNWSSTG to be saved when they are active.

- Edit Backup Control Group Entries, Display Backup Control Group Entries, and Create Control Group Entries displays added the Sync ID parameter for specifying a Synchronization ID to use for a synchronized checkpoint with the Save While Active.

Figure 3-20 shows an example of setting Edit Backup Control Group Entries.

```

                                Edit Backup Control Group Entries                                RCHAS70
Group . . . . . : TEST
Default activity . . . . .
Text . . . . .

Type information, press Enter.

      Backup      List ASP      Weekly   Retain Save   SWA
Seq  Items       Type Device   Activity Object While  Message      Sync
                                           SMTWTFS Detail Active Queue         ID

  10 *EXIT
  20 QGPL          *SYSBAS   FFFFFFFF *YES  *NO
  30 QRPG          *SYSBAS   *DFTACT  *ERR  *NO
  40 IFS          *LNK *SYSBAS   *DFTACT  *YES  *SYNC   IFSSAVE     ITSOTEST
  50 *EXIT
  
```

Figure 3-20 Sync ID on the Edit Backup Control Group Entries display

Journal receiver support

Two new special values have been added to the Backup Items in Control Groups:

- ALLCHGRVC** Save all journal receivers in user libraries that do not start with the letter Q that have never been saved or that have changed since they were last saved.
- ALLDTCRCV** Save all detached journal receivers in user libraries that do not start with the letter Q that have never been saved or that have changed since they were last saved.

3.18.3 Recovery enhancements

IBM i 6.1 includes recovery enhancements.

Logical/physical file restore order independence

You can defer the restore of database logical files, indexes, and SQL materialized query tables (MQTs), that have dependencies on other files.

In releases before IBM i 6.1, restoring files that resided in a different library than the based-on files was difficult to manage. If the library with a logical file was restored before the library with the underlying physical file, the logical file did not restore. To restore the logical file, you needed to attempt to restore from the media one or more times after restoring the physical file. Materialized query tables (MQTs) are physical files that have dependencies on based-on files. An MQT cannot be restored if one of its based-on files is missing.

For further information about the new defer support in IBM i 6.1, refer to 3.3, “Deferring the restore of dependent objects” on page 34.

Any logical/physical file dependencies are now handled by BRMS. A BRMS-defined ID of Q1ARSTID will be used for the Defer ID (DFRID) parameter on RSTLIB and RSTOBJ

commands. This parameter allows you to restore all objects in a set of libraries when the libraries with dependent objects are restored before the libraries with the objects they depend on.

Private authority support

IBM i 6.1 provides an option to save the private authorities to an object when the object is saved. They can also be restored if they were previously saved. For more information about this new support, refer to “Saving and restoring private authorities” on page 48. A new Private authorities (PVTAUT) parameter has been added to the RST nnn BRM commands and to the Change Recovery Policy display. Figure 3-21 shows an example of setting Change Recovery Policy.




Figure 3-21 Change Recovery Policy display - Private authorities

Remote journal support

IBM i 6.1 provides an option to save journal receivers to any system in the Backup, Recovery and Media Services (BRMS) network. If you need to apply any journal changes to an object, it is possible even from a remote journal receiver.

Priority sequencing of saved items for STRRCYBRM

There are times where the sequence of when you restore a specific library becomes important. Many times it is because of dependencies that objects have across libraries. To assist in these situations, BRMS has provided a way to prioritize these libraries as well as list them ahead in the recovery report, before all the other user libraries. To be able to use this useful function you will need to generate a list of these libraries (recovery order list) in the order that you wish them to be. Once that list has been made, use it on the Start Recovery using the BRMS (STRRCYBRM) command.

Create an object list using the Work with Lists using the BRM (WRKLBRM) command. For further information about creating a recovery list, refer to the “Priority Sequencing Your Recovery Report” section in *Backup, Recovery, and Media Services for i5/OS*, SC41-5345, available in the Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/books/sc415345.pdf>

Location omit for Start Recovery using BRM (STRRCYBRM)

Prior to IBM i 6.1, the Start Recovery using BRM (STRRCYBRM) command only had a parameter to omit locations. In 6.1 a new parameter, Volume locations to omit (OMITLOC),

has been added to allow the omission of up to 10 locations. Figure 3-22 shows an example of using the STRRCYBRM command.

```

Start Recovery using BRM (STRRCYBRM)

Type choices, press Enter.

Additional Parameters

Volume locations to include . . LOC          *ALL
                               + for more values
Volume locations to omit . . . . OMITLOC    *NONE
                               + for more values
From system . . . . . FROMSYS          *LCL
Print system information . . . . PRTSYSINF *NO
Use duplicate media . . . . . USEDUPMED   *NO
User recovery information . . . USRRCYINF  *NONE
Recovery order list . . . . . RCYORDLST   *NONE

```

Figure 3-22 Start Recovery using BRM (STRRCYBRM) - Omit locations

Start Recovery using BRM (STRRCYBRM) report enhancements

The Start Recovery using BRM (STRRCYBRM) report has added the following enhancements:

- ▶ Optical
- ▶ Encryption
- ▶ Location omit

Figure 3-23 shows the encrypted save key information from a Recovery Volume Summary Report.

```

*****
Volume Identifier Information
*****
Volume Volume          Media      Creation  Expiration  Storage  System
Serial Identifier      Status   Class     Date       Date       Location  ID
SAH001 SCRATCH20      *ACT     FMTOPTMLB 8/11/09    8/25/09   OPTMLB01 REESES
*****

-----
Encrypted Save Key Information
-----
Key store  Key store  Key
File      File      Record
Volume  Library  Name      Label
*****
LN4621  QUSRBRM  Q1AKEYFILE  FRED
Total number of key record label required . . . . . : 1
***** END OF LISTING *****

```

Figure 3-23 Encrypted Save Key information

Use of Date or Time on RST nnn BRM commands

The Save level time reference (SAVDATE) parameter has been added to RST nnn BRM commands. This new parameter contains two elements, Save date and Save time. These elements allow you to specify a more precise save level to restore from. Figure 3-24 shows the RSTOBJBRM command having the Save level parameter set to *SAVDATE and the Save level time reference parameter set to *LATEST.

```

Restore Object using BRM (RSTOBJBRM)

Type choices, press Enter.

Object . . . . . > Name, generic*, *ALL
      + for more values
Library . . . . . > Name
Device . . . . . > Name, *MEDCLS
      + for more values
Parallel device resources:
  Minimum resources . . . . . *SAV 1-32, *SAV, *NONE, *AVAIL
  Maximum resources . . . . . *MIN 1-32, *MIN, *AVAIL
  Object type . . . . . *ALL *ALL, *ALRTBL, *BNDDIR, ...
      + for more values
Save level . . . . . > *SAVDATE 1-99, *CURRENT, *SAVDATE
Save level time reference:
  Save date . . . . . Date
  Save time . . . . . *LATEST Time, *LATEST
  End of media option . . . . . *REWIND *REWIND, *LEAVE, *UNLOAD
  Option . . . . . *ALL *ALL, *NEW, *OLD, *FREE
  
```

Figure 3-24 Restore Object using BRM (RSTOBJBRM) - Save level time reference

3.18.4 Media services enhancements

The media services were also enhanced.

Work with Media using BRM (WRKMEDBRM)

The Select volumes (TYPE) parameter has three new special values:

- *ALL Displays or prints all media volumes.
- *ERR Displays or prints only media flagged in error.
- *INZ Displays or prints only media requiring initialization.

Media Balancing across systems

The new Start Media Balancing for BRM (STRBALBRM) command gives you the ability to balance a number of scratch volumes owned by each system in the BRMS network. With the scratch volumes balanced, a restricted state save will be less likely to run out of scratch media, and the system will not need to remotely ask other systems for more media. It is recommended that you run this when there is no other BRMS activity going on.

The Media priority (MEDPTY) and Number of media required (MEDREQ) parameters can be used to influence the outcome of the balancing.

For further information about balancing media, refer to the “Media Balancing Across Systems” section in *Backup, Recovery, and Media Services for i5/OS*, SC41-5345, available on the Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/books/sc415345.pdf>

Figure 3-25 shows the Start Media Balancing for BRM command.

```

Start Media Balancing for BRM (STRBALBRM)

Type choices, press Enter.

Action . . . . . > *SET          *REPORT, *BALANCE, *SET...
Media class . . . . . *ALL        *ALL, SAVSYS
Volume locations to include . . *ALL    *ALL, *HOME, TAP01, VAULT, ..
System . . . . . *NETGRP
      + for more values
Media priority . . . . .          0-9999
Number of media required . . . . 0-9999
  
```

Figure 3-25 Start Media Balancing for BRM (STRBALBRM)

Automated duplication of media

Automatic duplication can be used to duplicate data after it has been saved. In a control group backup, this will happen at each save done within the control group, so multiple duplications will occur in asynchronous jobs. The job log from the backup will contain messages indicating which jobs are performing the duplications. The backup will finish and return control to the user before the duplications have completed. Each duplication job will issue completion and error messages to the BRMS log. Use DSPLOGBRM to view these messages. This new feature is most ideal for ensuring that your virtual media are being duplicated off to physical media.

Enhancements to the Change Media Policy, Display Media Policy, and Create Media Policy displays include an Automatic duplication prompt, which includes these parameters:

- ▶ From devices
 - The device that you would like to duplicate from
- ▶ From device end option parameter
 - The end option you would like the From device to use
- ▶ To devices
 - The devices that you would like to duplicate to
- ▶ To device end option
 - The end option for the To device
- ▶ To media policy
 - The media policy that you would like to use for duplicating to
- ▶ Save media information
 - To choose to have the BRMS media information saved after all the automatic duplications have completed

Figure 3-26 shows Change Media Policy parameters.

```

Change Media Policy

Media policy . . . . . : ARCHIVAL

Type choices, press Enter.
Automatic duplication:
  From devices . . . . . Name, *NONE, *SAME, *MEDCLS.

  From device end option . . . . *REWIND, *LEAVE, *UNLOAD
  To devices . . . . . Name, *SAME, *MEDCLS, F4

  To device end option . . . . . *REWIND, *LEAVE, *UNLOAD
  To media policy . . . . . *SAME, *SYSPCY, F4
  Save media information . . . . *SAME, *BKUPCY, *LIB, *OBJ
  *NONE

```

Figure 3-26 Automatic duplication in Change Media Policy

DUPMEDBRM completion message

The Duplicate Media using BRM (DUPMEDBRM) command now issues a message upon completion indicating what output volumes were used, that is, data has been moved from media ABC to media XYZ.

Append rules for volume selection

The media append rules have been changed to select based on the most current expiration date so BRMS will not pick a tape that is close to expiring, unless there is nothing else to select.

Move Media retention capability

Move Media retention capability, which enables you to move media to a location where it is needed, and then after a specified amount of time, move it back to its original location without affecting its movement schedule.

Container action in move policies

A new field has been added to the media policy displays, Container action. The Container action column defines the behavior of containers when they arrive at the specified location.

Note: This column is only active when the Use container prompt has been set to *YES.

The possible values are:

- *OPEN The container will be opened when it arrives at this location.
- *CLOSE The container will be closed when it arrives at this location.
- *NONE The container's status will not be changed.

Figure 3-27 shows an example of setting the Use container parameter.

```

Change Move Policy
RCHAS70

Move policy . . . . . : OFFSITE
Home location . . . . . *SYSPCY      Name, *SYSPCY, *ORIGIN, F4 list
Use container . . . . . *YES          *YES, *NO
Verify moves . . . . . *YES          *YES, *NO
Calendar for working days . . . *ALLDAYS  Name, *ALLDAYS, F4 for list
Calendar for move days . . . . *ALLDAYS  Name, *ALLDAYS, F4 for list
Text . . . . . Entry created by BRM configuration

Type choices, press Enter.

Seq   Location   Duration   Container Action
10    VAULT        *EXP      *OPEN

```

Figure 3-27 Container action

3.18.5 Hierarchical Storage Manager enhancements

Hierarchical Storage Management (HSM) is a data storage technique that provides an automatic way of managing and distributing data between different storage layers (high-cost and low-cost) in order to meet the users' needs for accessing data while minimizing the overall cost. The concept of HSM involves the placement of data items in such a way as to store it efficiently.

In a typical HSM scenario, data files that are frequently used are stored on disk drives (faster devices), but the files can be migrated to slower storage devices (such as optical discs, or tape drives) if the files are not used for a certain period of time (typically a few months). If a data file that is stored on those slower devices is reused, it will be automatically moved back to disk storage. BRMS provides an interface to utilize this feature.

To find out more about how to set up BRMS for HSM, refer to *Hierarchical Storage Management*, SC41-5351.

Migration

Through the use of BRMS migration control groups, you can select which data to move from faster, relatively expensive DASD to slower but less costly DASD. You might also set these control groups up to move the data in the opposite direction if it becomes more frequently accessed.

Archive

Through the use of BRMS archive control groups, you can define criteria to archive less frequently used data from disk to less expensive media (tape or optical). The criteria can be based on inactivity limit, size, or even days used per month. Once the data meets the criteria, it is then saved to tape and deleted from the system, no longer taking up that much-needed space. The archive feature now supports Independent ASPs and Encryption. *Hierarchical Storage Management*, SC41-5351 does not document these capabilities. Restrictions and setup are similar to BRMS's backup control groups.

Dynamic Retrieval

Besides having the ability to archive the data off to removable media, it is just as important to recover the data. A user can easily recover the data using the normal BRMS recovery process; however, BRMS also provides Dynamic Retrieval of the data. This allows the recovery action to be prompted by a single access of the data. For example, if a database file has been archived, but an application actually tries to access this file, BRMS will then be asked to restore the file. After the file is restored, the application continues on. There is no major disruption in the application other than a delay for the restore of the file.

3.18.6 Device enhancements

IBM i 6.1 includes device enhancements.

Enhancements for 35xx tape libraries

35xx tape libraries do not have a library manager so the library is unable to advise other systems when a tape volume is ejected from one system. For example, System A might eject volume VOL01. System B would not be aware of this and still has the status of VOL01 set to *INSERT. This caused many issues with tape movements. To alleviate these problems, BRMS now acts like a library manager on behalf of the 35xx libraries.

DVD/Optical library support

The Work with Devices using BRM (WRKDEVBRM), Work with Media using BRM (WRKMEDBRM) and Work with Media Library Media (WRKMLMBRM) displays have changed to include optical devices, as shown in the following examples.

Figure 3-28 shows output after issuing the WRKDEVBRM command.

```
Work with Devices                                RCHAS70

Position to . . . . .                          Starting characters

Type options, press Enter.
 1=Add  2=Change  4=Remove  5=Display
 8=Work with status

Opt  Device      Category  Type/Model  Text
-----
OPT01  *OPT          6331/003  Entry created by BRM configuration
OPT02  *OPT          6333/002  Entry created by BRM configuration
TAP01  *TAP          3580/004  *NONE
VOPT01 *VRTOPT       632B/001  *NONE
```

Figure 3-28 Work with Devices using BRM (WRKDEVBRM)

Figure 3-29 shows all options available for the WRKMEDBRM command.

```

                                Work With Media
                                System:  RCHAS70
Position to . . . . .           Starting characters

Type options, press Enter.
  1=Add  2=Change  4=Remove  5=Display  6=Work with serial set  7=Expire
  8=Move 10=Reinitialize ...

      Volume      Volume
Opt  Serial      Identifier
                                Optical
                                Media Type

      GEN001      GEN001
      GEN002      GEN002

                                                                Bottom

Parameters or command:
====>
F3=Exit  F4=Prompt  F5=Refresh  F11=Volume information  F12=Cancel
F17=Top  F18=Bottom F23=More options

```

Figure 3-29 Work with Media using BRM (WRKMEDBRM)

Media library status tool

This tool produces a report that displays the users of the resources in a tape library for all systems in a BRMS network that share that library. To find the tape library resource status for a specific system, issue:

```
QSYS/CALL QBRM/Q1A0LD PARM('DEVUSESTS ' 'SYSTEM1 SYSTEM2')
```

The second parameter can specify multiple systems, restricted only by the parameter length of 200 characters.

To find the tape library resource status for all systems in the BRMS network, issue:

```
QSYS/CALL QBRM/Q1A0LD PARM('DEVUSESTS ' '*NETGRP')
```

Both commands generate an output file, QTEMP/DEVICELOG, and a printer report, QP1ADEVUSE.

Note: All characters in the PARM field need to be in capital letters.

Note: The tool is available through PTFs for IBM i 5.4 (SI24433) and IBM i5/OS V5R3 (SI24427). In order to get status from a particular system, that system must have the appropriate PTF applied and must be in the BRMS network.

3.18.7 Network enhancements

IBM i 6.1 includes the enhancements described in the following topics.

Automatically choose the network communication method

If both TCP/IP and SNA have been configured between networked systems and the Communication method is set to *AVAIL, BRMS will attempt to use TCP/IP first, but if that is unsuccessful, BRMS will then attempt to use SNA for its network operations. Figure 3-30 shows an example of a network group configuration.

```
Change Network Group                                RCHAS70 APPN

Network group . . . . : *MEDINV                    Position to . . . . :
FlashCopy state . . . : *ENDPRC
Receive media info . . : *NONE                      *NONE, *LIB
Communication method : *AVAIL                      *AVAIL, *IP, *SNA, *NONE
Notify period . . . . : 30                          30-99999 seconds
Type options, press Enter.
  1=Add  4=Remove  8=Set time
      Remote Local  Remote  Receive  System      Network
Opt Location Name Network ID Media Info Status      Status

(No entries found)
```

Figure 3-30 Setting the communication method to *AVAIL

TCP/IP interfaces

The BRMS networking features now support TCP/IP IPv6 (Internet Protocol version 6) and VIPA (Virtual IP Address) restricted state support.

BRMS now also supports multiple TCP/IP interfaces, including virtual and physical interfaces. This new feature allows you to share BRMS control groups, policies, media information, and storage locations across the network. This enables you to manage save and restore operations across all of your systems in a consistent manner.

3.18.8 Install enhancements

With 6.1 and follow-on releases, the BRMS file conversion now includes the BRMS printer files. The conversion will save current attributes of the printer files that are in library QBRM, install the shipped ones, and then apply the saved attributes of old printer files to the newly installed ones. This preserves any attribute changes that you have made prior to upgrading.

3.18.9 Other enhancements

IBM i 6.1 also includes the following enhancement.

Volume identifier (VOL)

The Volume identifier (VOL) parameter was changed from a 6-character parameter to a 32-character parameter. This is to support the new Optical media that can have up to 32-character volume identifiers. If the volume identifier exceeds 6 characters, BRMS will use

a Volume Prefix found in the System Policy to generate a unique 6-character volume serial that will map to the 32-character volume identifier.

All commands that use the VOL parameter have been changed to use the new parameter length.

3.19 Journal management

The following sections describe new and enhanced journaling functions.

3.19.1 Journaling libraries and automatically starting journaling

You can now journal libraries by using the Start Journal Library (STRJRNLIB) command. You can record and replay changes to libraries by using the Apply Journalized Changes (APYJRNCHG) command. The library journal code is Y.

Journaling can start automatically for objects such as database files, data areas, and data queues that are created, moved into, or restored into a journaled library. The library's inherit rules determine which objects to start journaling automatically and with what journaling attributes. You set a library's inherit rules with the Start Journal Library (STRJRNLIB) command and change them with the Change Journalized Object (CHGJRNOBJ) command. You can view the library's inherit rules with the Display Library Description (DSPLIBD) command or retrieve the inherit rules with the Retrieve Library Description (RTVLIBD) API.

Note: If the specified library contains a data area named QDFTJRN, no objects being created, moved, or restored into the library will inherit the journaling state of the library. Instead, objects will automatically start journaling based on the contents of the QDFTJRN data area. When the QDFTJRN data area is deleted, objects being created, moved, and restored into the library will inherit the journaling state of the library based on the inherit rules.

For more information, refer to the “Journaling libraries and automatically starting journaling” section on the *What's new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.2 Apply Journalized Changes (APYJRNCHG) enhancements

The Apply Journalized Changes (APYJRNCHG) command now applies additional journal entries, including journal entries for creating objects such as database files, data areas, and data queues in journaled libraries. The command also applies changes to logical files. The APYJRNCHG command replays any operations that could formerly only be replayed by the Apply Journalized Changes Extend (APYJRNCHGX) command.

For more information, refer to the “Apply Journalized Changes (APYJRNCHG) enhancements” section on the *What's new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.3 Using data port services with remote journaling

Remote journaling can now transfer the contents of journal receivers to other systems using data port services. Data port services allow for up to four lines of communication, allowing for greater resiliency.

For more information, refer to the “Using data port services with remote journaling” section on the *What's new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.4 Additional remote journal options and information

New options for remote journals are Synchronous Sending Time-out (SYNCTIMO) and Validity Checking (VLDCHK). Additional information about remote journals can be retrieved with the Work with Journal Attributes (WRKJRNA) command and the Retrieve Journal Information (QjoRetrieveJournalInformation) API. The new information includes measurements of how far behind the source system is at sending data to the remote system and the numbers and sizes of transfer bundles.

For more information, refer to the “Additional remote journal options and information” section on the *What's new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.5 Exit point for detaching journal receivers

The system calls all user exit programs registered on the QIBM_QJO_CHG_JRNRCV exit point whenever a journal receiver is detached from a journal. See Change Journal Receiver Exit Program for more information.

For more information, refer to the *Change Journal Receiver Exit Program* page on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/apis/xchgrcv.htm>

3.19.6 Ending the journaling of opened files

Database files can have journaling ended even while they are open, unless they have uncommitted changes. In previous releases files had to be closed before ending journaling. See the End Journal Physical File (ENDJRNPF) command for more information.

For more information, refer to the *End Journal Physical File (ENDJRNPF)* page on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/cl/endjrnpf.htm>

3.19.7 Enhancements to start journaling and end journaling commands

The Start Journal Physical File (STRJRNPF), End Journal Physical File (ENDJRNPF), Start Journal Object (STRJRNOBJ), and End Journal Object (ENDJRNOBJ) commands can now

start or end journaling all of the objects of the specified type in a library, or start or end journaling based on a generic name.

For more information, refer to the “Enhancements to start journaling and end journaling commands” section on the *What’s new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.8 System-managed access path protection

A new display has been added to the Display Recovery for Access Paths (DSPRCYAP) and the Edit Recovery for Access Paths (EDTRCYAP) panels that shows which access paths are eligible for system-managed access path protection (SMAPP) but are not currently being protected.

For more information, refer to the “Display access paths that are eligible but not journaled by SMAPP” section on the *What’s new for V6R1* page on Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiwhatsnew.htm?tocNode=int_170023

3.19.9 Replay Journal Entry (QjoReplayJournalEntry) API

The new Replay Journal Entry (QjoReplayJournalEntry) API can be used to replay the creates of data areas and data queues. The recreated objects will have all the same attributes the original objects were created with.

For more information, refer to the *Replay Journal Entry (QjoReplayJournalEntry) API* page on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/apis/qjorplay.htm>

3.19.10 Changing the journal recovery count and the cache wait time

The Change Journal Attributes (CHGJRNA) command has been introduced and can be used to change system-wide attributes for all journal environments on the system.

When a system-wide journal attribute is changed, the changed attribute will go into effect immediately for all applicable journal environments on the system. Not every journal environment will be affected by a change to a system-wide journal attribute. For example, if caching is not enabled for a local journal then the cache wait time will have no effect on that journal environment. Also, if the journal recovery count for a local journal is not set to *SYSDFT then the system-wide journal recovery count will have no effect on that journal environment.

For more information, refer to the *Change Journal Attributes (CHGJRNA)* page on Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/cl/chgjrna.htm>

3.20 Restrictions for previous release support

In IBM i 6.1, the maximum size of a save file is approximately 2 TB (where TB equals 1 099 511 627 776 bytes). In IBM i 5.4 and i5/OS V5R3, the maximum size of a save file is approximately 1 TB. Take care when saving save files on a 6.1 system intended for a previous release because they may not restore if they are too large.

3.20.1 NetWare Enhanced Integration

IBM i no longer supports NetWare Enhanced Integration for backing up Novell data. Use IBM Tivoli Storage Manager to back up and restore your Novell data.



High availability

This chapter discusses the following topics:

- ▶ “PowerHA for i” on page 76:
- ▶ “Clustering” on page 80
- ▶ “IBM iCluster for i” on page 83
- ▶ “Independent auxiliary storage pool” on page 85
- ▶ “OptiConnect” on page 87

4.1 PowerHA for i

PowerHA for i was developed in response to customers' requests for an enterprise-level IBM high availability product. PowerHA for i, along with IBM i 6.1, offers:

- ▶ IBM i 6.1 enables integrated support that uses IBM i and PowerHA for i to enable an end-to-end high availability (HA) deployment and management solution. You can cluster IBM i partitions on a single system, multiple systems, or a combination of systems. PowerHA for i capabilities include switching of switchable devices such as iASPs, printers, tape and optical devices, various communication line descriptions, and network server descriptions.
- ▶ Cluster administrative domain covers the majority of system-based objects required by switchable applications. This means that in most cases, the application environments on the cluster nodes are synchronized. Application and data can be activated seamlessly on an alternate node.
- ▶ PowerHA for i includes integrated source and target side tracking for geographic mirroring. This means that when you detach a target system, the resynchronization operation, after reattaching, includes only the changed objects on the source and target systems.
- ▶ PowerHA enables you to perform role swap operations using Metro Mirror, a synchronous replication product for the DS8000® server. You can readily perform both planned and unplanned switching operations with minimal impact on operational procedures. You should use Metro Mirror for best-case recovery point objective (RPO) and recovery time objective (RTO).
- ▶ For disaster recovery solutions, you need to have a remote recovery center. In such cases the primary objective is not RTO but rather RPO. You can use Global Mirror within PowerHA to perform asynchronous data replication operations.

System i PowerHA for i helps protect critical business applications from outages. Combined with IBM i 6.1, PowerHA for i delivers tools for configuring, monitoring, and managing your high availability (HA) clustering solution.

PowerHA for i includes three types of interfaces for controlling and managing the HA environment:

- ▶ High Availability Solutions Manager (solution-based) GUI for small client environments (deploying on a single iASP topology, which can include a mirrored copy through geographic mirroring)
- ▶ Cluster Resource Services (task-based) GUI for all environments
- ▶ Traditional IBM i command-line interface

The GUI operations are deployed through IBM Systems Director Navigator for IBM i, which is a Web-based console.

These user interfaces invoke cross-site mirroring (XSM) to manage geographic mirroring (native IBM i replication), Metro Mirror (DS8000 synchronous replication), or Global Mirror (DS8000 asynchronous replication). You can also perform flash copy operations on the DS8000 server.

Support for Metro Mirror, Global Mirror and FlashCopy is included as part of the PowerHA for i licensed program. In addition to the PowerHA for i licensed program you will also need external storage units provided by IBM System Storage and Copy Services. Figure 4-1 on page 77 shows the relationship between the components.

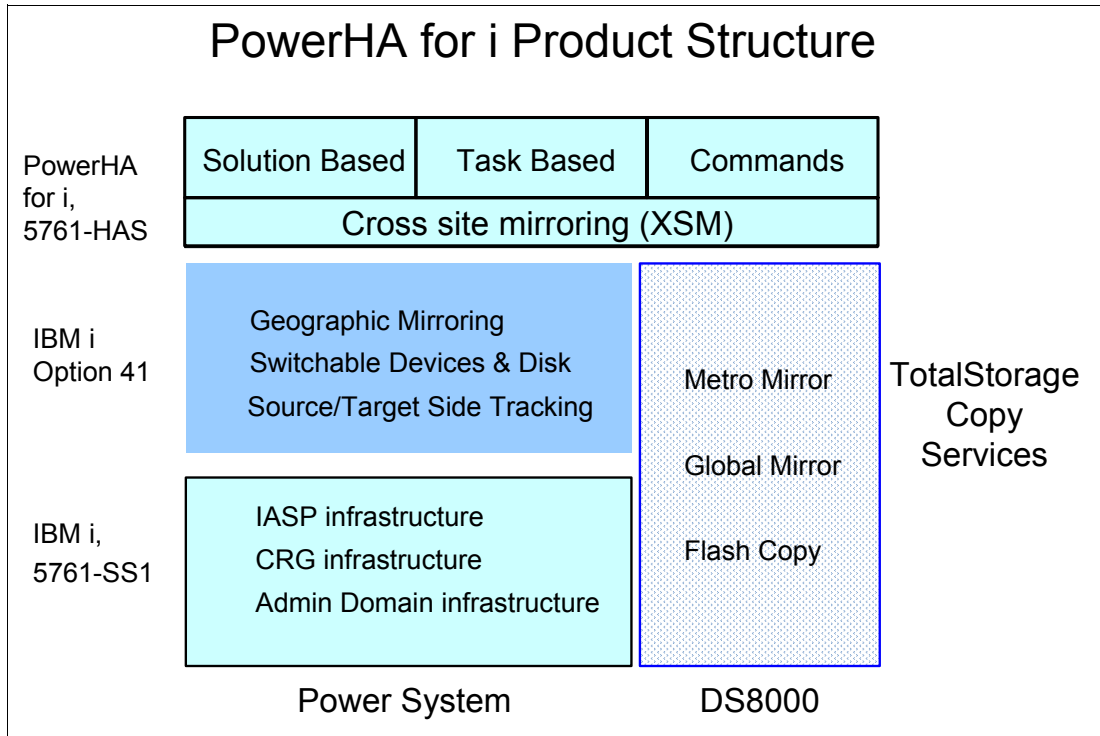


Figure 4-1 PowerHA for i product structure

IASP represents Independent Auxiliary Storage Pool support and CRG represents Cluster Resource Group support.

The following sections expand on the terminology used in this topic and provide additional support details.

PowerHA provides the cluster management, while the DS8000 provides the data resiliency with Metro Mirror and Flash Copy for nondisruptive off-line backups.

4.1.1 Geographic mirroring

Geographic mirroring, when used with i5/OS cluster technology, provides a high availability solution where a consistent copy of data stored in an independent disk pool at the production system is maintained on a mirrored copy. Geographic mirroring maintains a consistent backup copy of an independent disk pool using internal or external storage.

In the event of an outage on the production site, production can be switched to the mirrored copy of data, usually located at another location. Data is mirrored before write operations complete on the production system and is typically used for applications that cannot suffer any data loss in the event of a failure.

Geographic mirroring provides logical page level mirroring between independent disk pools through the use of data port services. Data port services manage connections for multiple IP addresses, which provides redundancy and greater bandwidth in geographic mirroring environments.

Geographic mirroring does not currently support an asynchronous mode. When implementing a geographic mirroring solution, the distance between production and mirrored independent disk pools will determine the application response time. Longer distances between the

production and mirrored copies may impact response times. Longer distances may require additional communication bandwidth to improve response times. Before implementing a high availability solution that uses geographic mirroring, you should understand your distance needs and their associated implications for communication requirements.

Geographic mirroring is a subfunction of cross-site mirroring (XSM), which is part of IBM i Option 41, High Availability Switchable Resources.

4.1.2 Metro Mirror

Metro Mirror is part of the optional Copy Services feature of IBM System Storage external storage unit support. Metro Mirror was known earlier as Synchronous Peer-to-Peer Copy (PPRC). Metro Mirror synchronously maintains a consistent copy of data between two IBM System Storage external storage units. When used with cluster technology, it provides high availability and disaster recovery.

The data mirroring occurs at the subsystem level, so your applications are not aware of the replication process.

Like geographic mirroring, this technology also mirrors data stored in independent disk pools. However, with Metro Mirror, disks are located on either IBM System Storage DS6000™ or DS8000 external storage units. Mirroring occurs from the source external storage units, which is usually located at the production site, to a set of target storage units, which are usually at the backup site. The data is copied between the external storage units to provide availability for both planned and unplanned outages.

Metro mirror is a function of the external storage unit where the target copy of a volume is constantly being updated to match changes made to a source volume. It is typically used for applications that cannot suffer any data loss in the event of a failure. The source and target volumes can be on the same external storage unit or on separate external storage units. In the case of separate units, the target storage unit can be located at another site up to 300 kilometers (186 miles) away. However, there may be performance implications when using synchronous communications over this distance and it may be more practical to consider a shorter distance to minimize the performance impact.

4.1.3 Global mirror

Global mirror maintains a consistent copy of data between two IBM System Storage external storage units. Global mirror was previously known as Asynchronous PPRC and PPRC Extended Distance (PPRC-XD) and is part of the optional Copy Services feature of IBM System Storage external storage unit support. It provides disk I/O subsystem level mirroring between two external storage units. This asynchronous solution provides good performance at unlimited distances, by allowing the target site to trail in currency a few seconds behind the source site.

Global mirror provides a remote, long-distance copy across two sites using asynchronous technology. It operates over high-speed, fibre channel communication links and is designed to maintain a complete and consistent remote mirror of data asynchronously at virtually unlimited distances with almost no impact to application response time.

Separating data centers by longer distances helps to provide protection from regional outages. This asynchronous technique provides better performance at unlimited distances. With global mirror data copied to the backup site is current with the production site in a matter of seconds. Global mirror creates a disaster recovery solution that provides high performance and a cost effective approach to data replication across global distances.

Since the replication is asynchronous there is potentially a significant time lag between the data changes made to the source disk repository and those same changes being applied on the target configuration. The latency depends on several factors including the amount of I/O activity on the source configuration and the bandwidth of the communication media linking to the target configuration.

4.1.4 FlashCopy

FlashCopy provides an almost instant, point-in-time copy of a logical volume (LUN) within independent disk pools (iASPs) on external storage, which can reduce the time it takes to complete daily backups. Point-in-time copy functions give you an instantaneous copy, or view, of what the original data looked like at a specific point-in-time.

The target copy is totally independent of the source independent disk pool and is available for both read and write access once the FlashCopy command has been processed.

FlashCopy uses a bitmap to represent the relationship between a source volume LUN and target volume LUN. This mechanism enables a FlashCopy operation to complete in a few seconds. The target volume is available for use immediately after the FlashCopy operation has completed.

4.1.5 Using PowerHA for i and IBM Total Storage Copy Services functions

Figure 4-2 shows an example of using the PowerHA for i product and IBM Total Storage products Copy Services facilities in a high availability configuration.

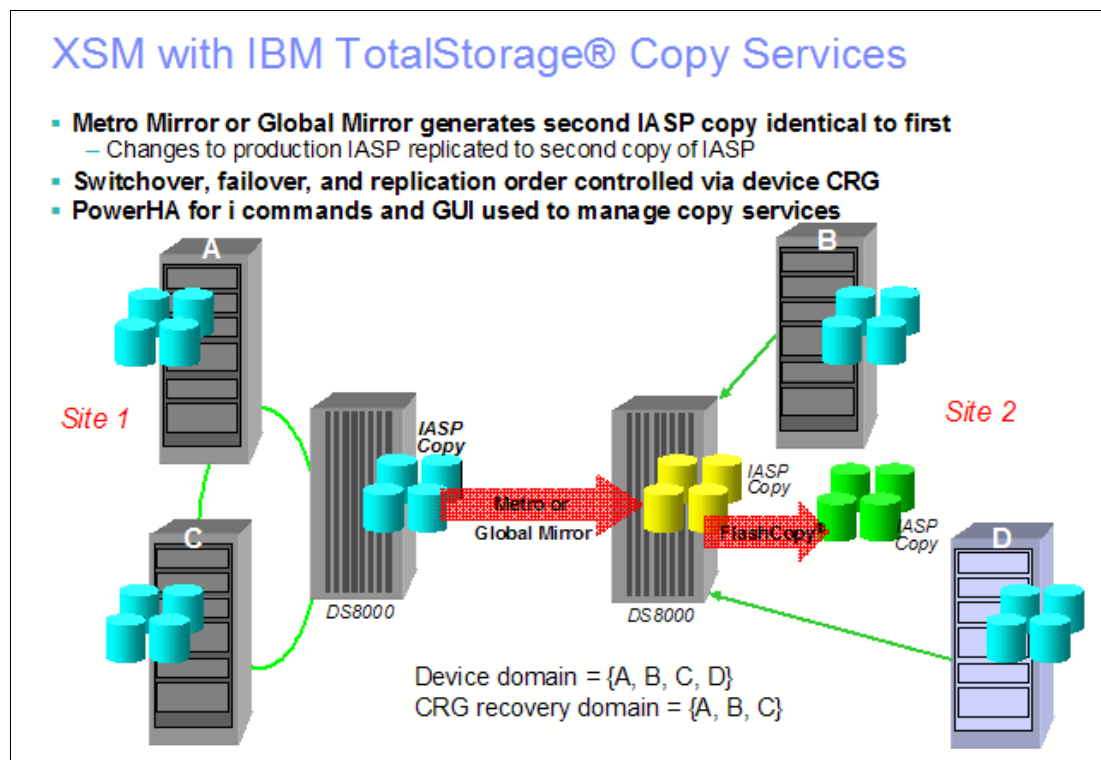


Figure 4-2 PowerHA for i, IBM i cross-site mirroring, Copy Services network example

PowerHA provides the cluster management, the DS8000 provides the data resiliency with Metro Mirror and Flash Copy for non-disruptive off-line backups.

4.1.6 Further information

For further information about PowerHA for i, refer to *Implementing PowerHA for IBM i*, SG24-7405, available at:

<http://www.redbooks.ibm.com/abstracts/sg247405.html>

Go to Infocenter, starting with the topic *Implementing high availability* at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaig/rzaigimplmentkickoff.htm?tocNode=int_169338.

Note: PowerHA for i was previously named High Availability Solutions Manager (HASM). At the time of writing, Infocenter had not been updated to reflect the name change.

4.2 Clustering

The following sections describe the clustering function, including cluster management jobs, switchable devices, cluster-level and CRG-level messages, and cluster administrative domain.

Note: With IBM i 6.1, new and enhanced clustering interfaces have been moved from IBM i to PowerHA for i. IBM i 5.4 clustering command support has been moved to library QUSRTOOL. The only cluster-related commands that remain in QSYS are:

DLTCRG	Delete Cluster Resource Group
DMPCLUTRC	Dump Cluster Trace
CHGCLURCY	Change Cluster Group Recovery
STRCHTSVR	Start Clustered Hash Table Server
ENDCHTSVR	End Clustered Hash Table Server

4.2.1 Cluster management jobs

In IBM i 6.1 cluster management jobs (such as canceling jobs and critical cluster jobs) have moved to “system jobs.” This minimizes accidental termination of these critical jobs. These jobs include:

- ▶ Cluster control job consists of one job that is named QCSTCTL
- ▶ Cluster resource group manager consists of one job that is named QCSTCRGM

Note: The QCSTCTL and QCSTCRGM jobs are cluster critical jobs. That is, the jobs must be running in order for the node to be active in the cluster.

- ▶ Each cluster resource group consists of one job per cluster resource group object. The job name is the same as the cluster resource group name.
- ▶ Cluster administrative domain jobs consist of a single system job running on every node in the cluster. The name of the system job is the name of the cluster administrative domain.

4.2.2 Switchable devices

In addition to independent disk pool devices, cluster resource service now supports the ability to switch other devices, such as Ethernet lines and network servers, which are defined in a device cluster resource group (CRG).

The new device types that are supported in i 6.1 are:

- ▶ Network server host adapter
- ▶ Optical device
- ▶ Tape device
- ▶ Ethernet line
- ▶ Token-ring line
- ▶ Network server

A device cluster resource group (CRG) is made up of a pool of hardware resources that can be switched as an entity.

To create switchable devices within a high-availability solution, the nodes that use these devices need to be a part of a device CRG. Prior to creating a device CRG, add all nodes that will share a switchable resource to a device domain.

To create a device CRG, use the IBM Systems Director Navigator for i interface as an example. Complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select Cluster Resource Services from the IBM Systems Director Navigator for i window. Recall that 5761-HAS and 5761-SS1 option 41 must be installed to see this link on the IBM Systems Director Navigator for i window.
4. On the Cluster Resource Services page, select **Work with Cluster Resource Groups** to show a list of cluster resource groups in the cluster.
5. On the Cluster Resource Group page, click the **Select Action** menu.
6. Select **New Device CRG** and click **Go**. The New Device CRG wizard is shown. The New Device CRG task is only available if all the nodes in the recovery domain are started.
7. Follow the instructions in the New Device CRG wizard to create the new device CRG. While running this wizard, you can create a new device CRG. You can also create either a new independent disk pool or specify an existing disk pool to use.

The device CRG keeps the hardware resource information identical on all recovery domain nodes and verifies that the resource names are identical. You can also configure a cluster administrative domain to keep the enrolled attributes of the configuration objects, which might include resource names, identical across the cluster administrative domain. If you are using cross-site mirroring, you should create separate device CRGs for independent disk pools and other types of switchable devices at each site.

4.2.3 Cluster-level and cluster resource groups-level messages

Cluster resource services supports the creation of cluster-level messages and cluster resource groups (CRG)-level messages to help in managing and controlling failover situations in a high-availability environment.

4.2.4 Cluster administrative domain

Cluster administrative domain now supports additional monitored resource types and enables single and compound attributes for monitored resource types. In addition, cluster administrative domain provides more detailed status messages for both monitored resources and the cluster administrative domain and synchronizes information regarding switchable devices.

- ▶ Enhanced monitored resources
 - User profiles (*USRPRF)
 - Class (*CLS)
 - Job description (*JOBDD)
 - ASP device description (*ASPDEV)
 - System values (*SYSVAL)
 - Network attributes (*NETA)
 - Environment variables (*ENVVAR)
 - TCP/IP Attributes (*TCPA)
- ▶ New monitored resources
 - Subsystem Descriptions (*SBSD)
 - Network Server Descriptions (*NWSD) of types *WINDOWSNT, *IXSVR, and *ISCSI.
 - NWS Configurations (*NWSCFG)
 - NWSH Device Descriptions (*NWSHDEV)
 - NWS Storage Spaces (*NWSSTG)
 - Tape Device Descriptions (*TAPDEV)
 - Optical Device Descriptions (*OPTDEV)
 - Ethernet Line Descriptions (*ETHLIN)
 - Token-ring Line Descriptions (*TRNLIN)

Other changes to Cluster Administrative Domain support include:

- ▶ New commands and interfaces

CHGCAD	Change Cluster Admin Domain
ENDCAD	End Cluster Admin Domain
STRCAD	Start Cluster Admin Domain
ADDCADMRE	Add Admin Domain Monitored Resource Entry
RMVCADMRE	Remove Cluster Admin Domain Monitored Resource Entry
ADDCADNODE	Add Cluster Admin Domain Node Entry
RMVCADNODE	Remove Cluster Admin Domain Node Entry
- ▶ Changed commands and interfaces

WRKCLU	Display admin domain monitored resource entries
CHGCLU	Change Cluster - The Change Cluster Configuration (renamed from CHGCLUCFG)
CRTCAD	Create Cluster Admin Domain (renamed from CRTADMDMN)
DLTCAD	Delete Cluster Admin Domain (renamed from DLTADMDMN)

Changes have also been made to increase automation and customization:

- ▶ Resource synchronization occurs at different times, for example:
 - When change is made to object and active nodes in Admin Domain
 - When node rejoins Admin Domain
 - When partitioned environment is corrected
- ▶ Synchronization options
 - Last change (default)
 - The last change that was made before the node joined the cluster administrative domain is processed by all nodes in the active domain.
 - The last change could have been made in the active domain or on the joining node while it was inactive.
 - Active domain
 - Only changes made on active nodes in an active cluster administrative domain are processed.
 - Changes made on a node while it was inactive are not passed to the active domain.
 - When a node joins the cluster administrative domain, it will be synchronized with the values from the active domain.

Note that additional coverage of the PowerHA for i product, 5761-HAS, and its capabilities is beyond the scope of this publication. You are referred to *Implementing PowerHA for IBM i*, SG24-7405, for additional information.

4.3 IBM iCluster for i

iCluster® is the logical replication-based solution offered by IBM. Note that there are other logical replication solutions also provided by IBM HA ISV partners.

IBM iCluster for i helps protect critical business applications and data during both planned and unplanned outages. Whether you run a small to medium business or a large enterprise with multiple mission-critical applications, the iCluster solution can meet your needs.

The iCluster product consists of the following:

- ▶ iCluster, a logical replication solution based on IBM i journaling
- ▶ Complete high-availability solution based on logical software replication and IBM i journaling and clustering
- ▶ Java-based GUI and green screen command and menu-based operations
- ▶ Monitoring, synchronization checking, role swapping, alarms and alerts, and automatic resynchronization
- ▶ Support and integration with DS6000 and DS8000 servers and Enterprise Storage Servers as well as third-party SAN hardware disk-based replication tools
- ▶ iBalance, a chargeable feature of iCluster, that performs bidirectional replication for workload balancing

Figure 4-3 demonstrates the iCluster product in a high-availability environment (using IBM i synchronous journal data transmission to a secondary system) along with a disaster recovery configuration (asynchronous data transmission to a tertiary system).

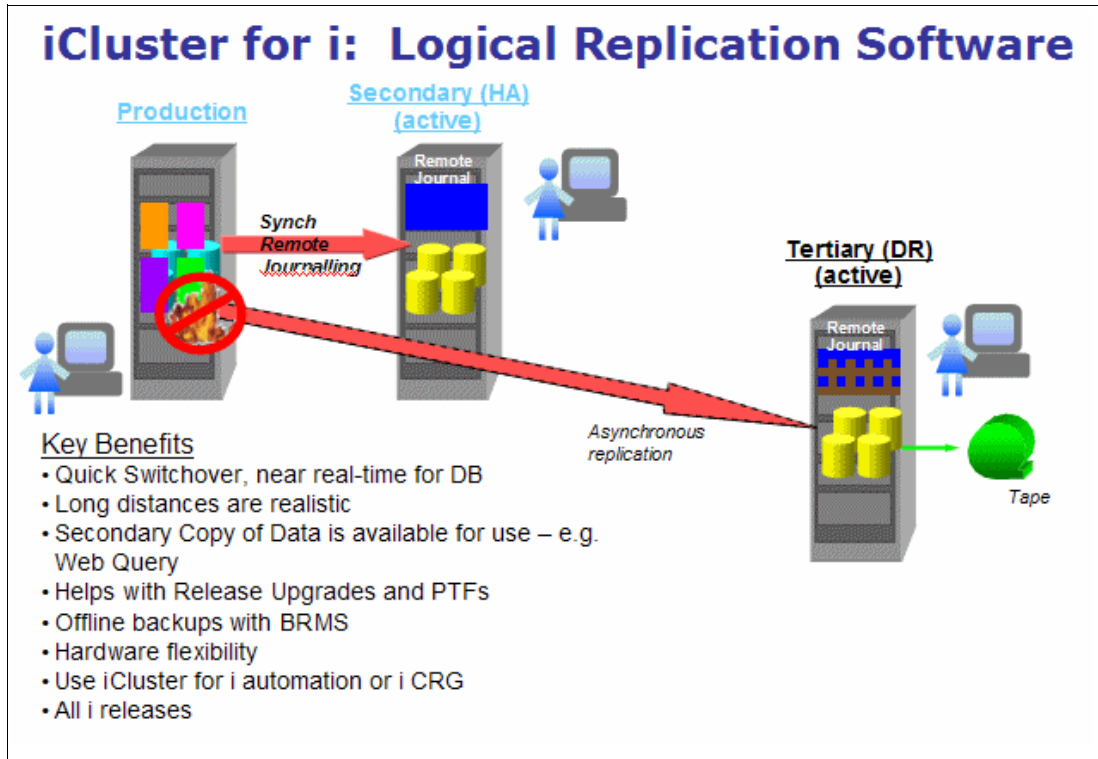


Figure 4-3 iCluster example showing possible HA and DR configurations

Note: iCluster was announced prior to IBM i 6.1 and is supported on both IBM i 5.4 and 6.1. iCluster enhancements were announced in October 2008 and have been included in this book because the information may be relevant when selecting a high-availability solution.

4.3.1 Complete integration with IBM i

iCluster is now installed as a licensed program (5733-ICL) with licensing by standard IBM i license keys. Fixes will now be delivered as PTFs.

Support has been streamlined with PMRs and APARs to track issues across both iCluster and IBM i.

There are two IBM iCluster offerings as of October 2008:

- ▶ IBM iCluster for i (5733-ICL) V5R1, which offers:
 - Full-featured IBM i HA offering
 - Remote journaling, sync checking, monitoring, alarms and alerts, and switchover
 - Auto-configuration command for ease of implementation
 - Scalable from SMB to large enterprises

IBM iBalance for i:

- Chargeable feature of iCluster
- Provides bi-directional HA for workload balancing on CBU Edition
- Conflict detection and resolution built in

4.3.2 Additional iCluster enhancements

Through October 2008, the following are key areas of enhancements in iCluster:

- ▶ Usability
 - Consolidation of multiple sync check types into easy-to-use groupings
 - Consolidated, simplified reports
 - Continuous Sync Check option that always runs as part of mirroring
 - Examining attributes of non-journaled objects
 - Improved integration between Sync Check status and auto-recovery (suspension) status
- ▶ Remote journaling performance and recoverability
 - Faster startup when starting from manually set initialization points (set position, mark position)
 - Enhanced recovery option for source system failure
- ▶ Integrated clustering support
 - Ability to configure and interoperate iCluster among pre-existing cluster configurations
 - Easier control of recovery domain upon node failure

For more information on iCluster go to:

<http://www.ibm.com/systems/power/software/availability/i5os.html#icluster>

4.4 Independent auxiliary storage pool

In this section we describe independent auxiliary storage pools (iASP) enhancements available under IBM i 6.1

4.4.1 Quiesce I/O for an auxiliary storage pool

IBM i 6.1 introduces the new Change ASP Activity (CHGASPACT) command, which enables you to control the activity of an auxiliary storage pool at a storage pool level, either:

- ▶ *SYSBAS: System pool 1 and all other existing dependent (basic) storage pools 2 to 32
- ▶ A named independent auxiliary storage pool

With this command you can:

- ▶ Suspend database transactions and database and Integrated File System file change operations for the system and configured basic auxiliary storage pools (ASPs) or an independent ASP.
- ▶ Resume suspended transactions and file operations for the system and configured basic auxiliary storage pools (ASPs) or independent ASPs that are currently in a suspended state.
- ▶ Force to disk changed pages associated with the partition and configured basic auxiliary storage pools (ASPs) or independent ASPs.

You can suspend an ASP in order to facilitate vary on after performing a flash copy or geographic mirror detach followed by a vary on. For more information about the new quiesce function, refer to 18.8, “Quiesce I/O for an Auxiliary Storage Pool” on page 455.

4.4.2 Target side tracking for geographic mirroring

The tracking space was introduced with IBM i 5.4. It enables geographic mirroring to track changed pages while in suspended status. With tracked changes we can avoid full resynchronization after a resume in many cases, therefore minimizing exposed times where you do not have a valid mirror copy. Tracking space gets configured when you configure geographic mirroring or change geographic mirroring attributes.

When migrating IBM i 5.3, tracking space is automatically added to your existing IASPs. If a tracking space was defined for that IASP during configuration (which is the default behavior) then tracking automatically occurs.

Tracking space is allocated inside the independent ASPs. The more tracking space you specify, the more changes the system can track. You can define the amount of space for tracking up to 2% of the total independent ASP's capacity. Be aware that this tracking space does not contain any changed data—it just holds information on what pages in the IASP have been changed.

In a geographic mirroring implementation, if the target side iASP is detached with tracking, a full synchronization is no longer required when the iASP is reattached.

When the mirror copy is reattached and geographic mirroring is resumed, only tracked changes will need to be synchronized. Changes made on the production copy (since the detach was done) are sent to the mirror copy and, in addition, any change made on the mirror copy will be overlaid with the original production data coming from the production copy of the IASP. Logically, any changes made on the detached mirror copy are undone, and any tracked changes from the production copy are applied.

Tracking of changes to the detached mirror copy enables you to achieve a reasonable resynchronization time when the detached mirror copy needs to be used for some reason (for example, backup).

4.4.3 System base commitment control transactions

In IBM i 6.1, support has been added for *SYSBAS commitment control transactions when the current name space of the job is an independent ASP.

The current name space of the job determines which independent ASP the commitment definition is created in. You can open a file under commitment control for the commitment definition under the following conditions:

- ▶ When the file is in the same independent ASP as the commitment definition.
- ▶ When the file is in the *SYSBAS ASP and the commitment definition has no resources registered in the independent ASP.

If you open a file that is in the *SYSBAS ASP while the commitment definition is in an independent ASP, the commitment definition is moved to the *SYSBAS ASP and no resources in the independent ASP can be registered.

4.4.4 Job queues

Job queues are now supported in an iASP, enabling an application to run in an iASP with fewer changes.

Important:

- ▶ Jobs on job queues in an iASP will not persist across a vary off and vary on.
- ▶ Jobs on job queues in an iASP will not be available on a backup system after switchover or failover.

4.5 OptiConnect

OptiConnect provides high-speed system-to-system interconnection and is useful in high-availability environments and high-speed access of data on multiple systems. OptiConnect includes a set of application programming interfaces (APIs) that allow application programs to access the high-speed bus transport. These APIs are intended for use by IBM Business Partners in the development of high-availability solutions. Applications that provide database mirroring using the OptiConnect connections are available from IBM Business Partners.

The mechanism used by OptiConnect software to access database files on another system is modeled after the mechanism used by distributed data management (DDM). DDM uses a DDM file and Advanced Program-to-Program Communication (APPC) communications protocols to redirect file access operations to another system. Similarly, OptiConnect uses DDM files and a special transport service to redirect file access operations to another system in an OptiConnect network. Thus, OptiConnect can achieve transport efficiencies that are not possible with a more general purpose, wide-area communications protocol.

When the OptiConnect feature is ordered and installed, an embedded IBM i driver streamlines the DDM implementation between nodes or systems or partitions. Data warehouse, Distributed Relational Database Architecture™ (DRDA®), and data propagation functions, such as journaling, can use this technology.

The underlying communications “fabric” includes either:

- ▶ Using the RIO-2 loop architecture to connect to one or more systems
- ▶ Virtual OptiConnect to connect among partitions within the same system

The older hardware SPD OptiConnect is no longer supported.

The following topics cover what is new in IBM 6.1.

4.5.1 Logical name support for DDM files

In a distributed data management (DDM) file, you can now use a logical name to identify the target system where data requests are sent over OptiConnect. Previously, you could use only the real name of the target system.

A logical name of the target system is the name of a data area in the QSOC library that contains the real system name. By using a logical name, you can create only one DDM file while being able to switch real system names.

The fast path method uses a special device description of QYCTSOC. When an Advanced Program-to-Program Communication (APPC) conversation is directed at this device, the device driver redirects the conversation through the OptiConnect bus, bypassing most of the standard distributed database management (DDM), Distributed Relational Database Architecture (DRDA), and APPC code.

Note that the QYCTSOC device description is created during the OptiConnect software installation (5761-SS1, option 23). For more information about logical name support, refer to the “Configuring fast path routing” topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/topic/rzau3/rzau3route1.htm?resultof=%22%43%6f%6e%66%69%67%75%72%69%6e%67%22%20%22%63%6f%6e%66%69%67%75%72%20%22%66%61%73%74%22%20%22%70%61%74%68%22%20>

4.5.2 Work with OptiConnect Activity (WRKOPCACT) command

A new parameter (OUTPUT) has been added to the Work with OptiConnect Activity (WRKOPCACT) command so that you can print the output of the command to a spooled file or to a database file.



Database

This chapter discusses the following topics:

- ▶ “Distributed database programming” on page 96
- ▶ “OnDemand and availability” on page 99
- ▶ “Commitment control and independent disk pools” on page 99
- ▶ “<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzakj/rzakjoptimize.htm> Miscellaneous database enhancements” on page 100
- ▶ “IBM OmniFind Text Search Server for DB2 for i5/OS” on page 102
- ▶ “IBM Systems Director Navigator for i” on page 403

5.1 Database performance and query optimization

In IBM i 6.1, support has been extended as described in the following topics:

5.1.1 SQL Query Engine

IBM i 6.1 eliminates two major SQL characteristics that prevented the usage of the SQL Query Engine (SQE):

- ▶ National Language Sort Sequences (NLSS)
- ▶ Functions relying on low-level translation functions

The low-level translation capability is required by commonly used SQL functions such as Upper and Lower. The NLSS support in IBM i is frequently used in overseas markets where the application requires the character data be sorted in a manner that matches the local language instead of the default *HEX ordering, which closely mirrors the sorting of the English alphabet. With these key functions now supported by SQE in 6.1, the most common items that prevent the usage of SQE are:

- ▶ ICU 2.6.1 sort sequences
- ▶ Non-SQL interfaces (OPNQRYF, Query/400, QQQQRY API)
- ▶ Logical File reference on FROM clause
- ▶ Select/Omit logical files defined in tables

While the SQE optimizer is unable to utilize Select/Omit logical files when building query access plans, IBM did change the default value for the Ignore_Derived_Index QAQQINI parameter to enable more SQE usage.

Prior to 6.1, the default value for this parameter was *NO, causing the SQE optimizer to reroute processing of any SQL request to the Classic Query Engine (CQE) anytime a derived logical file was encountered during the query optimization process. The default value in 6.1 for the Ignore_Derived_Index QAQQINI parameter has been changed to *YES.

This value allows the SQE query optimizer to ignore any keyed logical files that contain select/omit criteria during optimization instead of rerouting processing to CQE.

Note: An *access plan* is an internal IBM i control structure that describes the actions necessary to satisfy each query request. It contains information about the data and how to extract it. For any query, whenever optimization occurs, the query optimizer develops an optimized plan for how to access the requested data.

The SQE optimizer can also improve SQL performance with new indexing technologies. The SQE query optimizer can dramatically speed up the processing of the Sum and Average aggregate functions with Encoded Vector Indexes (EVI) whenever an EVI exists over the columns referenced on the function call.

While eliminating barriers to SQE usage is important, the most interesting 6.1 enhancements are those new capabilities that leverage the extensible architecture of the SQL Query Engine.

A prime example is DB2 for IBM i, the first product that uses self-learning query optimization. A self-learning query optimizer (initially introduced in i5/OS V5R3 and enhanced in IBM i 5.4) is able to analyze poor-performing query plans and dynamically adjust its internal algorithms based on feedback to select a better query plan for future executions. In 6.1, the SQE

optimizer will automatically analyze poor performing query plans to examine the I/O characteristics, as well as record retrieval patterns and compare them to the values used during optimization of the query. If the query optimizer detects significant mismatches, then the optimizer modifies its assumptions and algorithms during the next execution of that SQL request to generate a better plan.

To complement the learning optimizer, the DB2 runtime engine is also equipped with adaptive technologies that allow the plan of a currently running query to be modified on the fly to improve the efficiency of the query for the remainder of the run.

Self-learning only applies when an SQL statement has been run at least seven times and the estimated runtime is greater than 1 second. Examples of self-learning include:

- ▶ ColdIO to WarmIO Mitigation

The optimizer assumes a certain number of page faults for tables referenced in a query, if the optimizer is noticing that the plan is not causing page faults during execution (that is, the table is staying in memory). The optimizer reoptimizes the query with the assumption that the table will not incur page faults. The access plan may not be changed, but the optimizer at least tried to change its default behavior. WarmIO assumptions may cause table scans to look more attractive to the optimizer because a table scan method would not need to worry about I/O.

- ▶ FirstIO to AllIO Mitigation

For access plans built with a FirstIO optimization goal, the optimizer attempts reoptimization if it sees the access plan continually reading all of the rows in the table to return a result set. Because this is an ALLIO behavior, the optimizer reoptimizes the query with the ALLIO optimization goal to determine if a better performing plan can be built.

New access plan rebuild reasons were added for these two new optimizer-initiated plan rebuilds.

Note: These self-learning checks will only be done once by the optimizer (there is a flag in the plan cache entry recording whether or not the optimizer has performed mitigation on the plan's statement).

In addition to the self-learning enhancements, self-adjusting query execution has also been introduced to the SQE. Examples of self-adjusting query execution include:

- ▶ The optimizer builds an index over a hash table to minimize page fault issues associated with a poorly sized hash table.
- ▶ The statistics manager updates stale stats in small tables in real time.

5.1.2 SQL runtime performance improvements

In IBM i 6.1, support for SQL runtime performance has been improved.

Optimization time

Two changes have been made to improve the optimization time:

- ▶ The code path has been shortened where an SQL statement is first run.
- ▶ The optimizer's algorithms have been enhanced to enable the final access plan to be selected quicker.

Materialized query table refresh

The maintenance of indexes during the repopulation of a Materialized query table (MQT) has been changed to rebuild the index after the MQT is refreshed, instead of maintaining indexes for each new row.

Host Server variable length compression

Variable length columns in results sets are now compressed.

SQL CALL cache

Improvements have been made to the caching of repeated stored procedure calls.

SQL_STMT_REUSE INI option

This new INI option can be used to specify the number of times the statement must be reused before the statement is stored in the SQL extended dynamic package. If the number of times the statement has been reused is less than the specified INI option, a temporary copy of the statement is used for the query and any other job preparing the statement will also do a complete prepare.

The default value is 0 but any number from 1 to 255 can be specified.

5.1.3 Temporary distinct sorted list

A temporary distinct sorted list combines the features of the temporary hash table and the temporary sorted list.

Like the hash table, this temporary list allows the optimizer to collate the rows based on a column or set of columns. Like the sorted list, this temporary list also allows the optimizer to sequence the rows.

A temporary distinct sorted list contains a hash table data structure set up for efficient access to aggregate rows during population. In addition, a binary tree data structure is maintained over the hash table data structure so that the data can be accessed in sequence. The sorted aspect of the data structure allows for the efficient computation of super-aggregate rows in SQL statements that contain GROUP BY ROLLUP.

A temporary sorted aggregate hash table is an internal data structure and can only be created by the IBM i database manager.

5.1.4 Temporary values list

The temporary values list is a temporary object that allows the optimizer to store rows of data specified by the user in a VALUES clause of a SELECT statement or a CREATE VIEW statement.

The list is an unsorted data structure that is used to simplify the operation of the query. Since the list does not have any keys, the rows in the list can only be retrieved by a sequential scan operation.

A temporary values list is an internal data structure and can only be created by the database manager.

5.1.5 Optimizing performance using the plan cache

The plan cache is a repository that contains the access plans for queries that were optimized by SQE. The purposes of the plan cache are to:

- ▶ Facilitate the reuse of a query access plan when the same query is re-executed.
- ▶ Store runtime information for subsequent use in future query optimizations.
- ▶ Provide performance information for analysis and tuning.

Various new functions are available with SQL plan cache to assist in optimizing performance. These include:

- ▶ SQL details for jobs when working with active jobs
- ▶ SQL plan cache size threshold
- ▶ Additional actions available from the SQL plan cache statements window
- ▶ SQL plan cache event monitors

For more information about these new functions, refer to topics starting with 5.7.12, “Change the SQL plan cache size threshold” on page 118.

5.1.6 QAQQINI updates

Previous to IBM 6.1, the query options file QAQQINI was introduced to provide the ability to dynamically modify or override the environment in which queries are executed through the Change Query Attributes (CHGQRYA) command and the QAQQINI file. The query options file QAQQINI is used to set some attributes used by the database manager.

For each query that is run the query option values are retrieved from the QAQQINI file in the schema specified on the QRYOPTLIB parameter of the CHGQRYA CL command and used to optimize or implement the query.

IBM i 6.1 offers control of four additional environmental attributes through QAQQINI:

- ▶ **SQL_DECFLOAT_WARNINGS (*DEFAULT|*YES|*NO)**
This parameter provides the ability to get warnings for SQL DECFLOAT computations and conversions involving division by 0, overflow, underflow, an invalid operand, an inexact result, or a subnormal number.
- ▶ **SQL_FLAGGER (*DEFAULT|*NONE|*STDIDB2)**
Specifies whether non-standard statements are flagged. This parameter allows you to flag SQL statements to verify whether they conform to the Core level of the ISO/IEC 9075-2003 standards.
- ▶ **SQL_PSEUDO_CLOSE (*DEFAULT|<Integer value>)**
Prior to 6.1, SQL cursor open processing checks for the presence of a data area named QSQPSCLS1 in the job's library list. If the data area is found, all reusable cursors are marked as candidates for reuse and are pseudo closed the first time rather than the second time they are closed by the application. Without this data area, the default is that a cursor will not become reusable until the second close.

Pseudo closing the first time results in leaving some cursors open that may not be reused. This can increase the amount of auxiliary and main storage required for the application. This can be monitored using the WRKSYSSTS command. For the amount of auxiliary storage used, look at the *% system ASP used*. For the amount of main storage, examine the faulting rates on the WRKSYSSTS display. The format and the contents of the data

area are not important. The data area can be deleted using the following command:
DLTDTAARA DTAARA(QGPL/QSQPSCLS1).

The existence of the data area is checked during the first SQL open operation for each job. It is checked only once and the processing mode remains the same for the life of the job. Because the library list is used for the search, the change can be isolated to specific jobs by creating the data area in a library that is included only in the library lists for those jobs. The new SQL_PSUEDO_CLOSE attribute can now be used instead of the data area.

► SQL_STMT_REUSE (*DEFAULT|<Integer value>)

Specifies the number of times the statement must be reused before it is stored in the SQL extended dynamic package. If the number of times the statement has been reused is less than the specified INI option, a temporary copy of the statement is used for the query and any other job preparing the statement will also do a complete prepare.

For further information about QAQQINI options, refer to the “QAQQINI query options” topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzajq/qryopt.htm>

5.1.7 Derived key index

You can use the SQL CREATE INDEX statement to create a derived key index using an SQL expression.

Traditionally, an index could only specify column names in the key of the index over the table it was based on. With this support, an index can have an expression in place of a column name that can use built-in functions, user-defined functions, or some other valid expression. Additionally, you can use the SQL CREATE INDEX statement to create a sparse index using a WHERE condition.

The Upper function is typically one of the more popular functions used in SQL-derived indexes. The Upper function is often used by application developers to implement case-insensitive searches.

While this type of case-insensitive search functions in IBM i 5.4, performance is improved with 6.1. This is because prior to 6.1 the UPPER function prevented the query optimizer from using an index to speed up this search. The new SQL-derived index support allows the index to be created with the Upper function in the key definition and the SQE query optimizer to use this index to quickly retrieve data because the key expression on the index exactly matches the column expression in the case-insensitive query.

Some customers have successfully used shared-weight sort sequences to improve performance issues associated with case-insensitive searches. However, that approach does require a more complex configuration when creating the database and applications. In addition, sort sequences cannot be used with more complex expressions such as:

```
WHERE UPPER(CONCAT(FirstName,LastName)) = 'ALBERTYOUNG'.
```

A large majority of SQL built-in functions can be included in the key definition of an SQL-derived index. For a complete list of the functions, refer to the *DB2 for i SQL Reference* PDF in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/db2/rbafz.pdf>

For examples and information about restrictions on derived indexes, refer to the “Using derived indexes” topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzajq/rzajqusederived.htm?resultof=%22%74%6f%74%61%6c%69%78%22%20>

5.1.8 Full outer join

Full outer join is supported by the SQE optimizer. Just as a right outer or right exception join is rewritten by the optimizer to reflect the supported join types of inner, left outer, or left exception, a full outer join is rewritten in 6.1.

A full outer join of A FULL OUTER JOIN B is equivalent to a (A LEFT OUTER JOIN B) UNION ALL (B LEFT EXCEPTION JOIN A). The following example illustrates the rewrite.

```
SELECT EMPNO, LASTNAME, DEPTNAME FROM CORPDATA.EMPLOYEE XXX FULL OUTER JOIN
CORPDATA.DEPARTMENT YYY ON XXX.WORKDEPT = YYY.DEPTNO
```

This query is rewritten as follows:

```
SELECT EMPNO, LASTNAME, DEPTNAME FROM CORPDATA.EMPLOYEE XXX LEFT OUTER JOIN
CORPDATA.DEPARTMENT YYY ON XXX.WORKDEPT = YYY.DEPTNO
UNION ALL
SELECT EMPNO, LASTNAME, DEPTNAME FROM CORPDATA.DEPARTMENT YYY LEFT EXCEPTION
JOIN CORPDATA.EMPLOYEE XXX ON XXX.WORKDEPT = YYY.DEPTNO
```

A query with multiple FULL OUTER JOIN requests, such as A FULL OUTER JOIN B FULL OUTER JOIN C can quickly become complicated in this rewritten state. This is illustrated in the example below. Provided your query does not run in live data mode, the optimizer may choose to facilitate performance both during optimization and at runtime by encapsulating intermediate results in a temporary data object that can be optimized once and plugged in to both the scanned and probed side of the rewrite. These shared temporaries eliminate the need to make multiple passes through the specific tables to satisfy the request. In this example the results of the (A FULL OUTER JOIN B) would be a candidate for encapsulation during its FULL OUTER JOIN with C.

```
A FULL OUTER JOIN B FULL OUTER JOIN C
```

This query is rewritten as follows:

```
((A LEFT OUTER JOIN B) UNION ALL (B LEFT EXCEPTION JOIN A)) LEFT OUTER JOIN C )
UNION ALL (C LEFT EXCEPTION JOIN ((A LEFT OUTER JOIN B) UNION ALL (B LEFT
EXCEPTION JOIN A))
```

As FULL OUTER implies that both sides of the full outer join request can generate NULL values in the resulting answer set, local selection provided in the WHERE clause of the query could result in the appropriate downgrade of the FULL OUTER to a LEFT OUTER or INNER JOIN. If you want FULL OUTER behavior and want some local selection applied, you should specify the local selection in the ON clause of the FULL OUTER JOIN or use common table expressions. For example:

```
WITH TEMPEMP AS (SELECT * FROM CORPDATA.EMPLOYEE XXX WHERE SALARY > 10000)
SELECT EMPNO, LASTNAME, DEPTNAME
FROM TEMPEMP XXX FULL OUTER JOIN CORPDATA.DEPARTMENT YYY
ON XXX.WORKDEPT = YYY.DEPTNO
```

5.1.9 New database monitor format

The performance statistics table, QSYS/QAQQDBMN, shipped with IBM i, has a new format in IBM i 6.1.

For further information, refer to the “Database monitor SQL table format” topic in Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzajq/phfdds.htm?tocNode=int_170370

5.2 Distributed database programming

The following topics describe distributed database programming including Distributed Relational Database Architecture (DRDA) and Distributed Data Management (DDM).

5.2.1 Distributed Relational Database Architecture

In IBM i 6.1, Distributed Relational Database Architecture (DRDA) support has been extended to include the following:

Decimal floating-point data type

IBM i 6.1 introduced support for a new data type, decimal floating point.

A decimal floating-point number is an IEEE 754R number with a decimal point. The position of the decimal point is stored in each decimal floating-point value. The maximum precision is 34 digits. The range of a decimal floating-point number is either 16 or 34 digits of precision, and an exponent range of 10-383 to 10384 or 10-6143 to 106144, respectively.

The minimum exponent, E_{min} , for DECFLOAT values is -383 for DECFLOAT(16) and -6143 for DECFLOAT(34). The maximum exponent, E_{max} , for DECFLOAT values is 384 for DECFLOAT(16) and 6144 for DECFLOAT(34).

In addition to the finite numbers, decimal floating-point numbers can also represent the following three special values:

- ▶ Infinity - A value that represents a number whose magnitude is infinitely large.
- ▶ Quiet NaN (Not a Number) - A value that represents undefined results which does not cause an invalid number warning.
- ▶ Signaling NaN - A value that represents undefined results which will cause an invalid number warning if used in any numerical operation.

When a number has one of these special values, its coefficient and exponent are undefined. The sign of an infinity is significant (that is, it is possible to have both positive and negative infinity). The sign of a NaN has no meaning for arithmetic operations.

Note: Decimal floating-point variables can only be used in C or Java.

Decimal floating-point is now also supported in DRDA.

IBM i 6.1 introduces a more granular approach to encrypted user IDs and passwords. In previous releases it was only possible to specify that a password be encrypted (or not) at connection time. You can now specify whether encryption is required on either the user ID or password, or both.

New special values on the Remote authentication method (RMTAUTMTH) parameter on the Change RDB Directory Entry (CHGRDBDIRE) command and the Lowest authentication method (PWDRQD) parameter on Change DDM TCP/IP Attributes (CHGDDMTCPA) can be used to change the authentication level. The new values are:

*USRIDPWD	Refuse the connection if a DDM connection request does not contain a password associated with the user ID.
*USRENCPWD	Require the sending of an encrypted password along with the user ID by a remote client system attempting to connect.
*ENCUSRPWD	Require the sending of an encrypted user ID and encrypted password by a remote client system attempting to connect.

For more information on DRDA authentication, refer to the *Distributed database programming* PDF in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/ddp/rba11.pdf>

SQL statement interruption

SQLCancel() is used to end the processing of an SQL statement operation that is running synchronously. To cancel the function, the application calls SQLCancel() with the same statement handle that is used by the target function, but on a different thread. How the function is canceled depends on the operating system. DRDA in IBM 6.1 supports the SQL Cancel() statement.

Note: It is now recommended to use the value *USRENCPWD in place of *ENCRYPTED.

SQL limits

DB2 for i imposes some SQL and database limits. Some of these limits have changed in IBM i 6.1. In order to get further information about which limits have been changed, refer to the *SQL Limits* table in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/db2/rbafz1limitabs.htm>

Heterogeneous scrollable cursors

To enable scrollable cursors on the DB2 for Linux, UNIX, and Windows V8 clients, you must be using DB2 FixPak 4, or later. You can also use a client that runs a later version of DB2 for LUW (Logical Unit of Work) that provides scrollable cursor support to the IBM i platform. If you are using FixPak 4 on a client that runs DB2 for Linux, UNIX, and Windows V8, you must perform one of the following actions:

Note: If you are using a later version of the DB2 for LUW client that provides the scrollable cursor support to the System i platform, you do not need to perform these actions.

1. Issue this command:

```
UPDATE CLI CFG FOR SECTION i_dbname USING CURSORTYPES 1
```

Where *i_dbname* is the name of your IBM i database.

2. Edit the `db2cli.ini` file using this syntax:

```
CURTORTYPES = 1
```

Set Client Information (SQLESETI) API

The SQLESETI() API permits an application to set client information associated with a specific connection, provided that a connection already exists.

By using this API, the application server can pass the client's user ID, workstation information, program information, and other accounting information to the DB2 server. The default value for the special registers is the empty string.

The application can choose not to specify a relational database name. If a relational database name is not specified, any client information will be set for all existing, as well as future, connections.

This API can be used to establish values before a connection to a database, or it can be used to set or modify the values after a connection has been established. If a database connection does not exist when this API is called, the information will be cached globally and applied to all new connections established.

DRDA now supports the Set Client Information (SQLESETI) API.

Internet Protocol version 6 (IPv6)

IPv6 is the next evolution in Internet Protocol. Most of the Internet uses IPv4, which has been used for over 20 years. However, IPv4 has limitations that might cause problems as the Internet expands. IPv6 is the updated version of IPv4 and is gradually replacing IPv4 as the Internet standard.

Extended indicators

As in previous releases, you can use an indicator variable to set a null value for a column in an INSERT or UPDATE statement.

In IBM i 6.1 there are now two forms of indicators for INSERT and UPDATE statements, *normal indicators* and *extended indicators*. When you use normal indicators, an indicator set to any negative value is interpreted as the null value. When you use extended indicators, the negative values have several different meanings.

When processing UPDATE or INSERT statements using normal indicators, SQL checks the indicator variable (if it exists). If it contains a negative value, the column value is set to null. If it contains a value greater than -1, the column is set from the associated host variable value.

For more information and examples, refer to the "Indicator variables used to assign special values" topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzajp/rzajpnu11var.htm>

Optimistic locking

Using commitment control to ensure database integrity incurs overhead. The overhead that is associated with locking not only increases utilization, but also reduces concurrency.

A new alternative is to use *optimistic locking*, which involves issuing explicit UPDATES by using a WHERE clause that uniquely determines a particular record. Optimistic locking ensures that the record does not change after it is retrieved. DRDA now supports optimistic locking.

5.2.2 Distributed Data Management

In IBM i 6.1, Distributed Data Management (DDM) support has been extended to include the following.

Decimal floating-point data type

IBM i 6.1 introduced support for a new data type, decimal floating point.

A decimal floating-point number is an IEEE 754R number with a decimal point. The position of the decimal point is stored in each decimal floating-point value. The maximum precision is 34 digits. The range of a decimal floating-point number is either 16 or 34 digits of precision, and an exponent range of 10-383 to 10384 or 10-6143 to 106144, respectively.

See “Decimal floating-point data type” on page 96, under 5.2.1, “Distributed Relational Database Architecture” on page 96, for more details about this topic.

Note: Decimal floating-point variables can only be used in C.

Decimal floating point is now also supported in DDM.

5.3 OnDemand and availability

You can now journal libraries by using the Start Journal Library (STRJRNLIB) command. For more information, refer to 3.19.1, “Journaling libraries and automatically starting journaling” on page 70.

5.4 Commitment control and independent disk pools

IBM 6.1 includes enhancements to commitment control when it is used in combination with independent disk pools (iASPs).

5.4.1 System base commitment control transactions

In IBM i 6.1, support has been added for *SYSBAS commitment control transactions when the current name space of the job is an independent ASP.

The current name space of the job determines which independent ASP the commitment definition is created in. You can open a file under commitment control for the commitment definition under the following conditions:

- ▶ When the file is in the same independent ASP as the commitment definition
- ▶ When the file is in the *SYSBAS ASP and the commitment definition has no resources registered in the independent ASP

If you open a file that is in the *SYSBAS ASP while the commitment definition is in an independent ASP, the commitment definition is moved to the *SYSBAS ASP and no resources in the independent ASP can be registered.

5.4.2 Quiesce processing

IBM i 6.1 introduced support for quiescing a base or independent disk pool. Quiescing is the process of suspending database transactions and database and Integrated File System file change operations for the system and configured basic auxiliary storage pools (ASPs) or independent ASPs. For more information on quiescing, refer to 18.8, “Quiesce I/O for an Auxiliary Storage Pool” on page 455.

Quiescing an independent disk pool while commitment control is active can have a detrimental effect on performance. For further information, refer to the “Optimizing performance for commitment control” topic in Infocenter at:

5.5 <http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzakj/rzakjoptimize.htm>

Miscellaneous database enhancements

The following sections describe the Run SQL Statements (RUNSQLSTM) and DB2 for i catalog views enhancements.

5.5.1 Run SQL Statements (RUNSQLSTM)

The Run SQL Statements (RUNSQLSTM) command has proven to be a popular tool with IBM i developers who need to execute SQL scripts from a CL program or command line interface.

With 6.1, the RUNSQLSTM command now has the ability to process and execute SQL statements stored in a stream file within the integrated file system. Support prior to 6.1 required the SQL statements to be stored in a source physical file member. The new stream file support has been implemented by adding the new source stream file (SRCSTMF) parameter, as shown in Figure 5-1 on page 101.


```

                                Run SQL Statements (RUNSQLSTM)

Type choices, press Enter.

Source file . . . . . Name
  Library . . . . . *LIBL Name, *LIBL, *CURLIB
Source member . . . . . Name
Source stream file . . . . .
.....
Commitment co : Source stream file (SRCSTMF) - Help :
Naming . . . : :
: Specifies the path name to the file that contains the SQL :
: statements. The path name can be either absolute or :
: relative. :
: : Bottom :
: F2=Extended help F3=Exit help F10=Move to top :
: F12=Cancel F13=Information Assistant F14=Print help :
: :
:.....

```

Figure 5-1 Run SQL statements (RUNSQLSTM)

This new stream file support can assist developers in several ways:

- ▶ SQL statements no longer have to be limited to 80 characters or less when stored in a stream file. This 80-character limit for statements stored in a source file member made SQL coding quite difficult and tedious for IBM i developers.
- ▶ A stream file can accommodate larger, more complex SQL scripts since they have a size limit of 1 TB—much larger than the 16 MB limit of a source file member.
- ▶ Your IBM i SQL scripts can be managed by a stream file-based change management system.

5.5.2 DB2 for i catalog views

The database manager maintains a set of tables containing information about the data in each relational database. These tables are collectively known as the *catalog*. The catalog tables contain information about tables, user-defined functions, distinct types, parameters, procedures, packages, views, indexes, aliases, sequences, constraints, triggers, and languages supported by DB2 for i. The catalog also contains information about all relational databases that are accessible from this system. There are 10 new views in 6.1:

- ▶ **SYSCOLUMNSTAT**
Contains one row for every column in a table partition or table member and one row for every column statistics collection.
- ▶ **SYSPARTITIONINDEXSTAT**
Contains one row for every index built over a table partition or table member.
- ▶ **SYSPARTITIONSTAT**
Contains one row for every table partition or table member.
- ▶ **SYSTABLEINDEXSTAT**

Contains one row for every index that has at least one partition or member built over a table. If the index is over more than one partition or member, the statistics include all those partitions and members.

▶ SYSTABLESTAT

Contains one row for every table that has at least one partition or member. If the table has more than one partition or member, the statistics include all partitions and members.

▶ SYSINDEXSTAT

Contains one row for every SQL index. Use this view when you want to see information for a specific SQL index or set of SQL indexes.

▶ SYSMQTSTAT

Contains one row for every materialized table. Use this view when you want to see information about a specified materialized query table or set of materialized query tables.

▶ SYSPACKAGESTAT

Contains one row for each SQL package in the SQL schema.

▶ SYSPARTITIONMQTSTAT

Contains one row for every index built over a table partition or table member. Use this view when you want to see index information for indexes built on a specified table or set of tables.

▶ SYSPROGRAMSTAT

Contains one row for each program, service program, and module that contains SQL statements.

The information returned in SYSINDEXSTAT, SYSMQTSTAT, SYSPARTITIONINGINDEXES, and SYSPARTITIONMQTS is similar to that returned through Show Indexes in System i Navigator.

Note: The first five views shown were PTFed back to IBM i 5.4.

For more information about views, refer to the “DB2 for i catalog views” topic in Infocenter at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/db2/rbafzcatalog.htm?tocNode=int_170523

5.6 IBM OmniFind Text Search Server for DB2 for i5/OS

IBM Text Search for DB2 for i5/OS (5733-OMF) allows a DB2 for i5/OS user to issue SQL statements to satisfy familiar text search queries on data that is stored in a DB2 database.

IBM Text Search for DB2 for i5/OS provides a set of administrative stored procedures and two built-in functions (CONTAINS and SCORE) that are used to search text indexes created from documents stored in a DB2 table. The administrative stored procedures are used to enable and disable text searching and to create, update, and drop text indexes.

A text index can be created over any CHAR, VARCHAR, CLOB, BLOB, DBCLOB, GRAPHIC, VARGRAPHIC, BINARY, or VARBINARY column that contains plain text, HTML, XML, or many rich document types (such as PDF files). The data is read from the text column and is converted to Unicode (CCSID 1208) before it is indexed.

Text indexes are not normal DB2 indexes. They are not maintained automatically, cannot be journaled, and cannot be backed up using normal backup and restore methods. Text indexes are created and stored on a text search server, which can be on a server running i5/OS, Linux, UNIX, AIX, or Windows. The text search server contains a collection of significant terms extracted from each row of the column. A TCP/IP connection is used to communicate with the text search server.

The CONTAINS and SCORE functions are built-in functions that are integrated into DB2.

For further information, refer to the *OmniFind® Text Search Server for DB2 for i5/OS V1.1* PDF in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzash/rzash.pdf>

5.7 System i Navigator database task

Various System i Navigator database tasks have been added or updated in this release. For more information and examples, see the following topics.

5.7.1 Create multiple indexes with Run SQL Scripts

To create multiple indexes with Run SQL Scripts, select the indexes you want to create, right-click the advised indexes in the Index Advisor window and select Show SQL. This launches the Run SQL Scripts window that contains the CREATE INDEX statements for each index selected.

Figure 5-2 shows the first step to create multiple indexes, selecting **Index Advisor**.

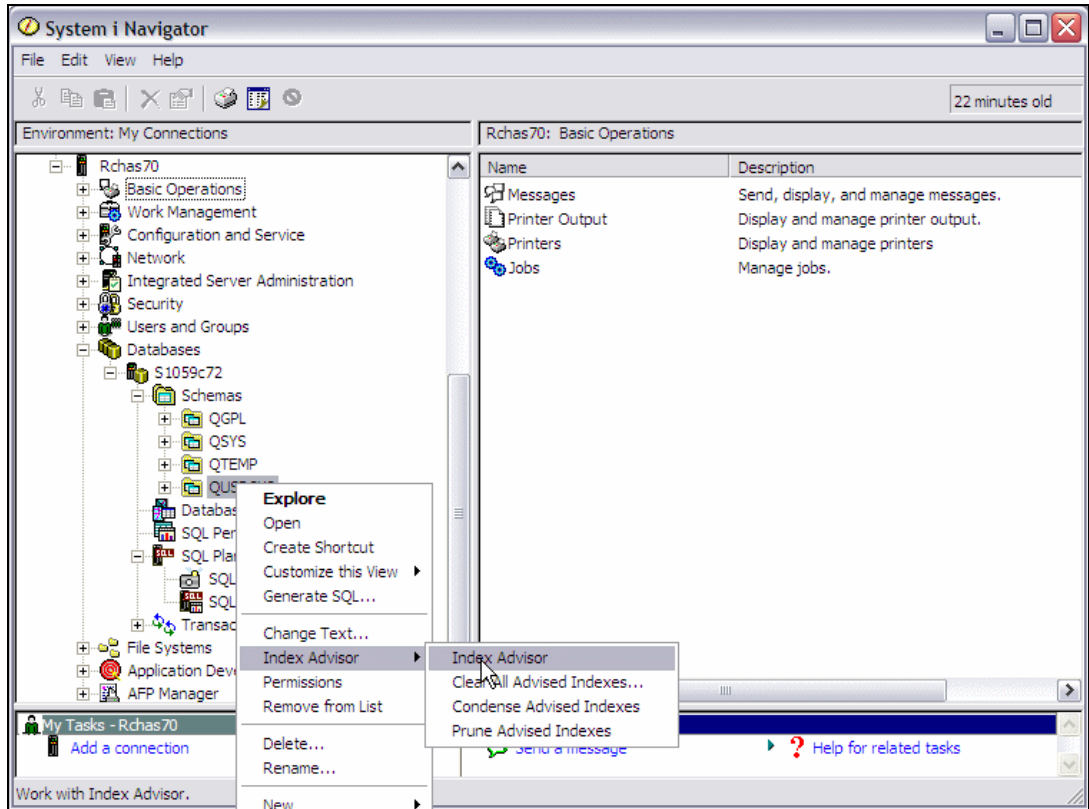


Figure 5-2 Selecting Index Advisor

Figure 5-3 shows the selected indexes.

Table for Which Index was Advised	Schema	Short Name	Partition	Keys Advised	Leading Keys Order Independent	Index Type Advised
QASZRAIRD	QUSRSYS	QASZRAIRD		COMPID, FULLP...	COMPID, FULLP...	Binary Radix
QATOCTCPIP	QUSRSYS	QATOCTCPIP		KEYWORD, DATA	KEYWORD	Binary Radix
QATOCTCPIP	QUSRSYS	QATOCTCPIP		DATA		Encoded vector ..
QATOCTCPIP	QUSRSYS	QATOCTCPIP		DATA		Binary Radix
QASZRAIRC	QUSRSYS	QASZRAIRC		PACKPROD		Binary Radix
QATOCRTE	QUSRSYS	QATOCRTE		NEXTHOP		Binary Radix
QATOCRTE	QUSRSYS	QATOCRTE		NEXTHOP		Encoded vector ..
QATOCRTE	QUSRSYS	QATOCRTE		RTEDEST, NEXT...	RTEDEST	Binary Radix
QASZRAIRC	QUSRSYS	QASZRAIRC		VENDOR, VERSI...	VENDOR, VERSI...	Binary Radix
QASZRAIRC	QUSRSYS	QASZRAIRC		VENDOR, VERSI...	VENDOR, VERSI...	Binary Radix
QASZRAIRC	QUSRSYS	QASZRAIRC		PRODNAME, VE...	PRODNAME, VE...	Binary Radix
QASZRAIRC	QUSRSYS	QASZRAIRC		PRODNAME, CO...	PRODNAME, CO...	Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Memory Handle...	"Memory Handle...	Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Created by", "S...	"Created by", "S...	Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Free Storage M...	"Free Storage M...	Binary Radix
QAUGDBPMD	QUSRSYS	QAUGDBPMD		"Memory Handle"		Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Memory Handle"		Binary Radix
QAEZDISK	QUSRSYS	QAEZDISK		DIOBTP, DIOBNM	DIOBTP	Binary Radix
QAEZDISK	QUSRSYS	QAEZDISK		DIOBLI		Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Name", "Create...	"Name", "Create...	Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Status", "Free S...	"Status", "Free S...	Binary Radix
QAUGDBPMD2	QUSRSYS	QAUGDBPMD2		"Status", "Name"	"Status"	Binary Radix
QAUGDBPMD	QUSRSYS	QAUGDBPMD		"Status", "Name"	"Status"	Binary Radix
QAUGDBNAV	QUSRSYS	QAUGDBNAV		DTACPRAMNAME		Binary Radix

Figure 5-3 Selecting indexes

Figure 5-4 shows the result of an SQL statement.

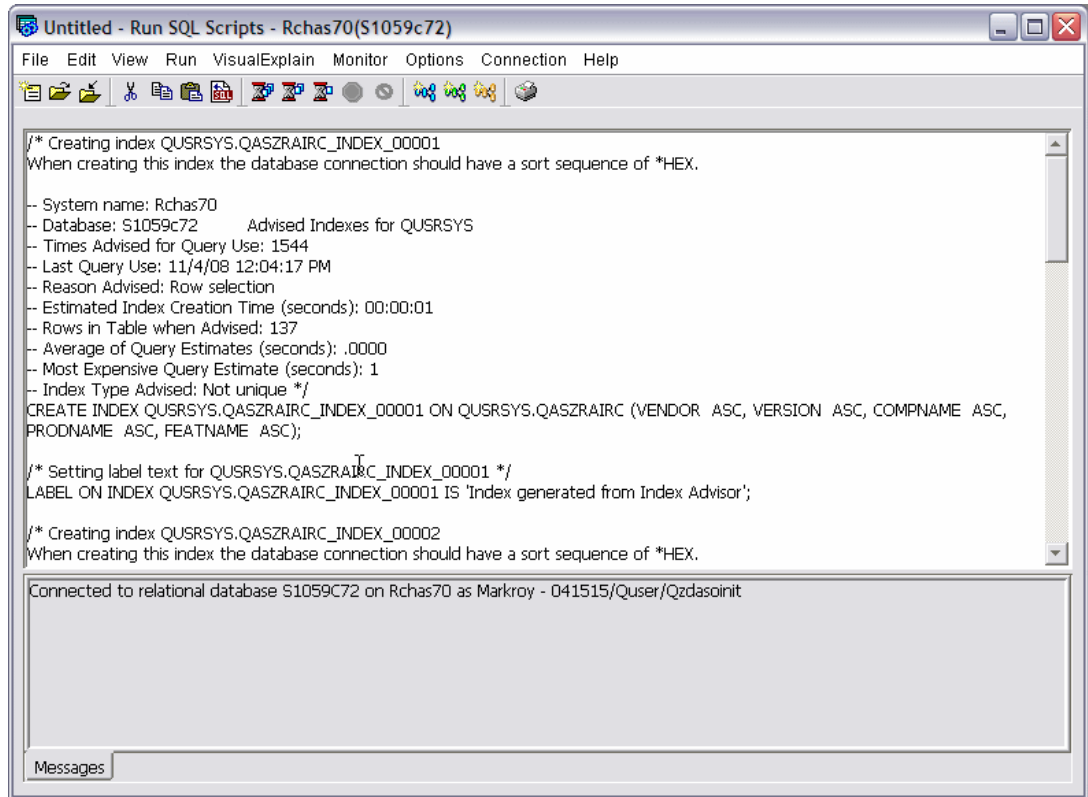


Figure 5-4 Running multiple SQL scripts

5.7.2 Add a comment to a database object

To add a comment to a database object, right-click the object in the right pane and select Comments as shown in Figure 5-5.

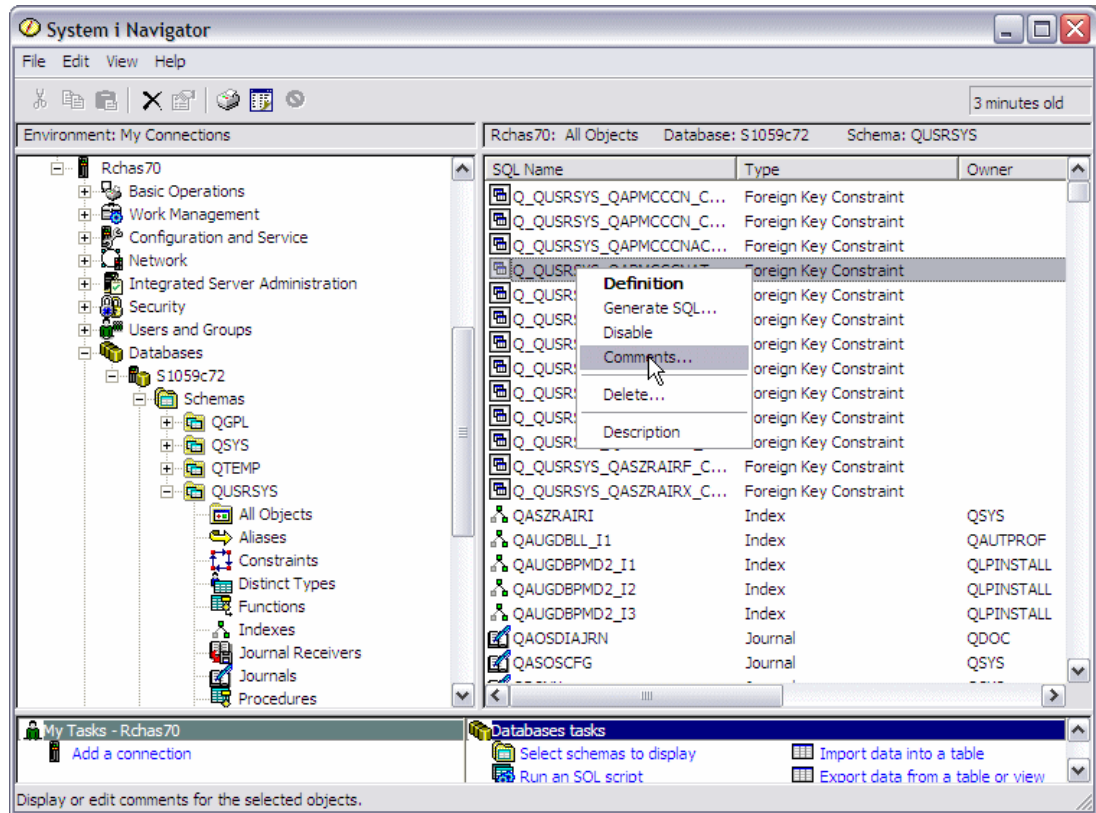


Figure 5-5 Add a comment to a database object

5.7.3 Reset usage counts for indexes and materialized query tables

Resetting the usage counts for a table allows you to determine how the changes that you have made to your indexing strategy are affecting the indexes and constraints that are created on that table. For example, if your new strategy causes an index to never be used, you may be able to delete that index. Resetting usage counts on a table affects all indexes and constraints that are created on that object.

Note: Resetting usage counts for a keyed physical file or a constraint in the Show Indexes window resets the counts of all constraints and keyed access for that file or table.

To reset usage counts for a specific index, right-click the table in the Tables window and select **Reset Usage Counts**. This launches the Confirm Reset Usage Counts dialogue box, to confirm which action you want to take. Refer to the example shown in Figure 5-6.

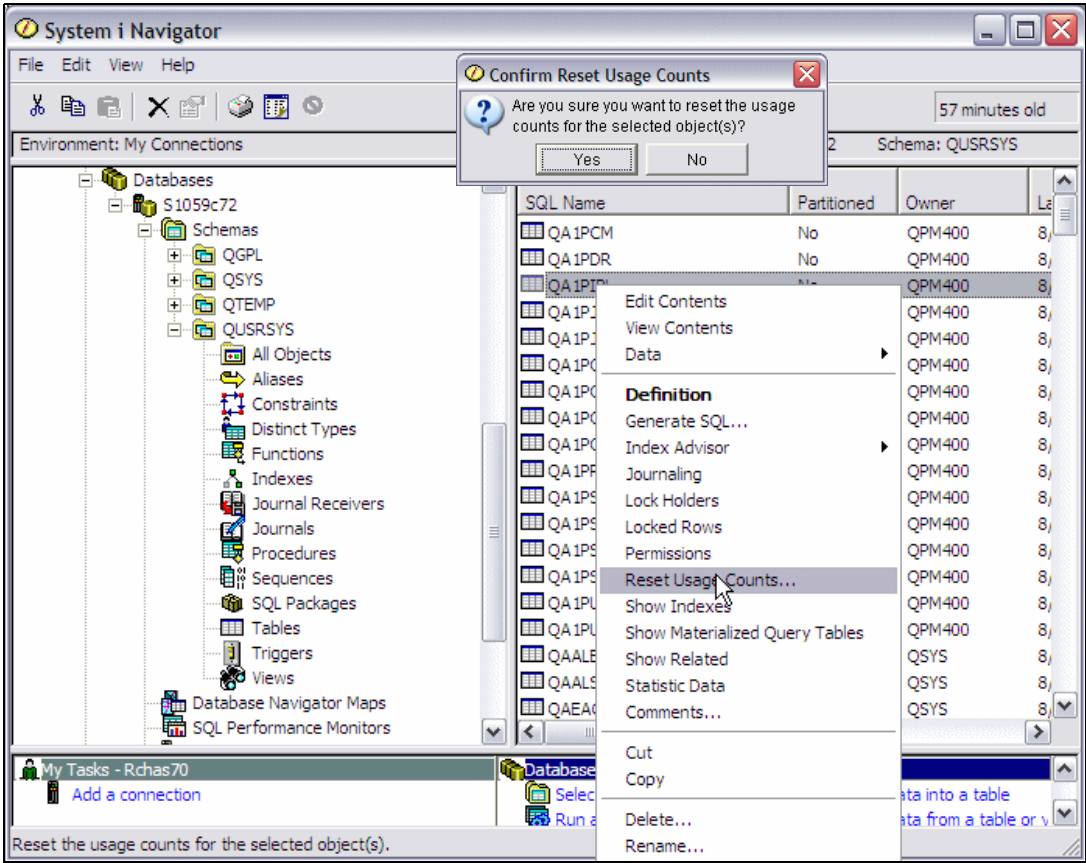


Figure 5-6 Reset usage counts for indexes

Similarly, to reset usage counts for a specific materialized query table, right-click the materialized query table in the Indexes window and select Reset Usage Counts. This launches the Confirm Reset Usage Counts dialogue box, to confirm which action you want to take.

5.7.4 Show indexes for all tables in a schema

To view indexes for all tables in a schema, right-click the Tables in the left pane and select **Show Indexes**. This launches the Indexes window that contains the indexes. Figure 5-7 shows how to list indexes for all tables in a schema.

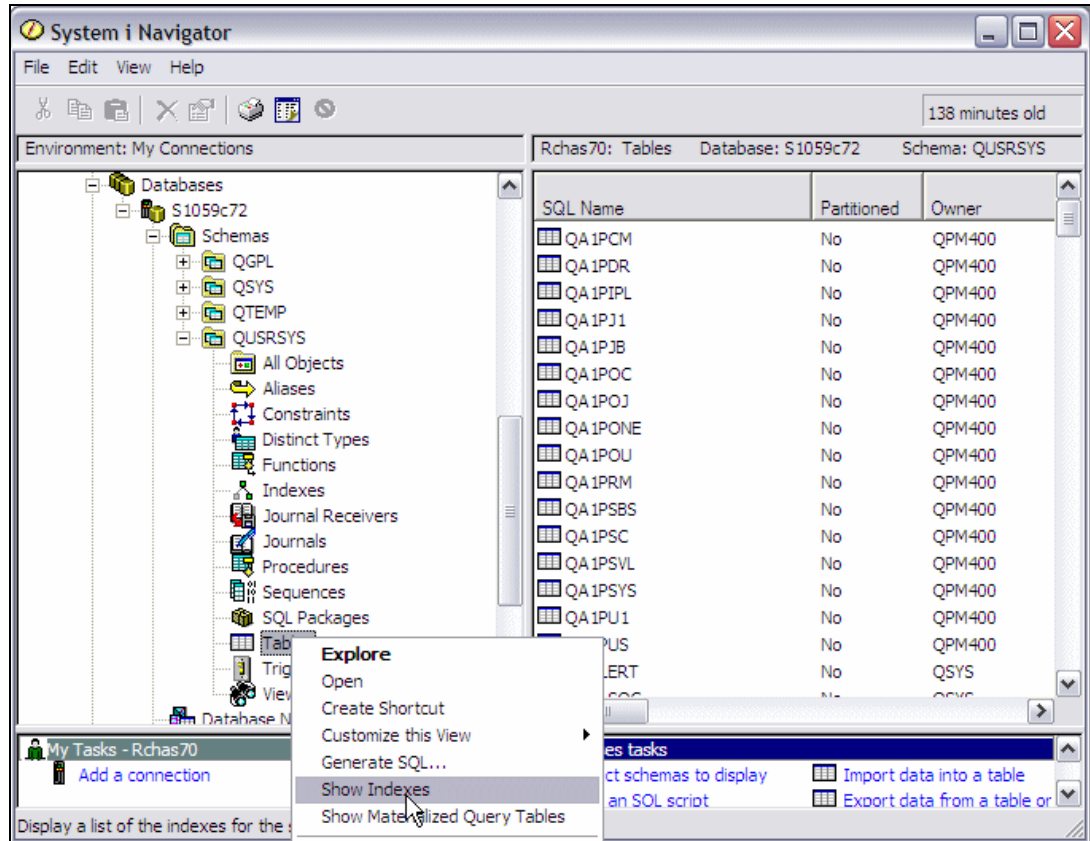


Figure 5-7 Show indexes for all tables in a schema

Figure 5-8 shows a list of indexes for all tables in a schema.

SQL Name	Type	Schema	Owner	Short Name	Text
*1 MAINTAINED TEMPORARY...	Temporary index			1 MAINTAINED T...	1 MAINTA
Q_QUSRSYS_QAPMCCCLS_C...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCN_C...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCN_C...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCN_C...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNAC...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNAC...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNAT...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNAT...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNDF...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNDF...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNTB...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNTB...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNTB...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCNTB...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCR_C...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCR_C...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCRAC...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCCRAC...	Foreign Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCORG...	Primary Key Con...	QUSRSYS	QSYS		
Q_QUSRSYS_QAPMCCORG...	Primary Key Con...	QUSRSYS	QSYS		

Figure 5-8 Indexes for all tables in a schema

5.7.5 Show indexes for a specific table

To view indexes for a specific table, right-click the table in the Tables window and select **Show Indexes**. This launches the Indexes window that contains the indexes. Figure 5-9 shows how to list indexes for a specific table.

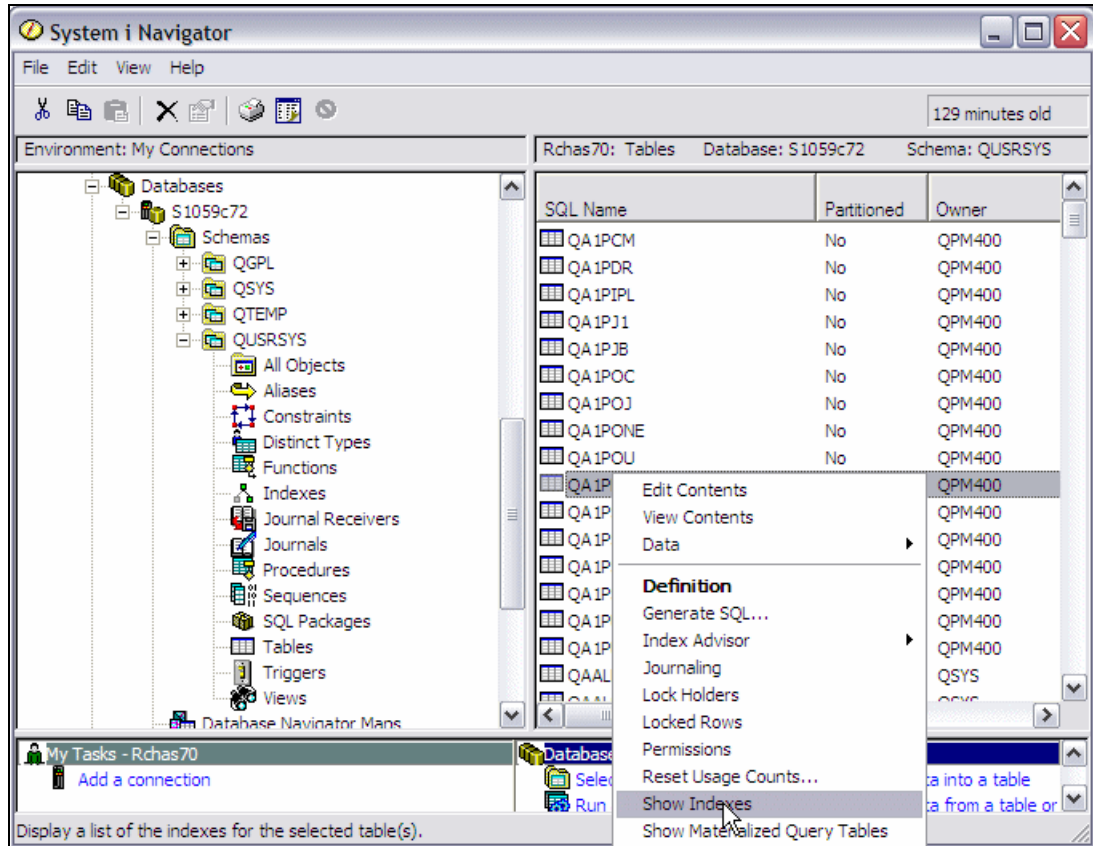


Figure 5-9 Show indexes for a specific table

Figure 5-10 shows a list of indexes for a specific table.

SQL Name	Type	Schema	Owner	Short Name	Text
QA1PRM1	Keyed Logical File	QUSRSYS	QPM400	QA1PRM1	
QA1PRM2	Keyed Logical File	QUSRSYS	QPM400	QA1PRM2	
QA1PRM3	Keyed Logical File	QUSRSYS	QPM400	QA1PRM3	

Figure 5-10 Indexes for a specific table

5.7.6 Show materialized query tables for all tables in a schema

To view materialized query tables for all tables in a schema, right-click Tables in the left pane and select **Show Materialized Query Tables**. This launches the MQT window that contains the indexes.

Figure 5-11 shows materialized query tables for all tables in a schema.

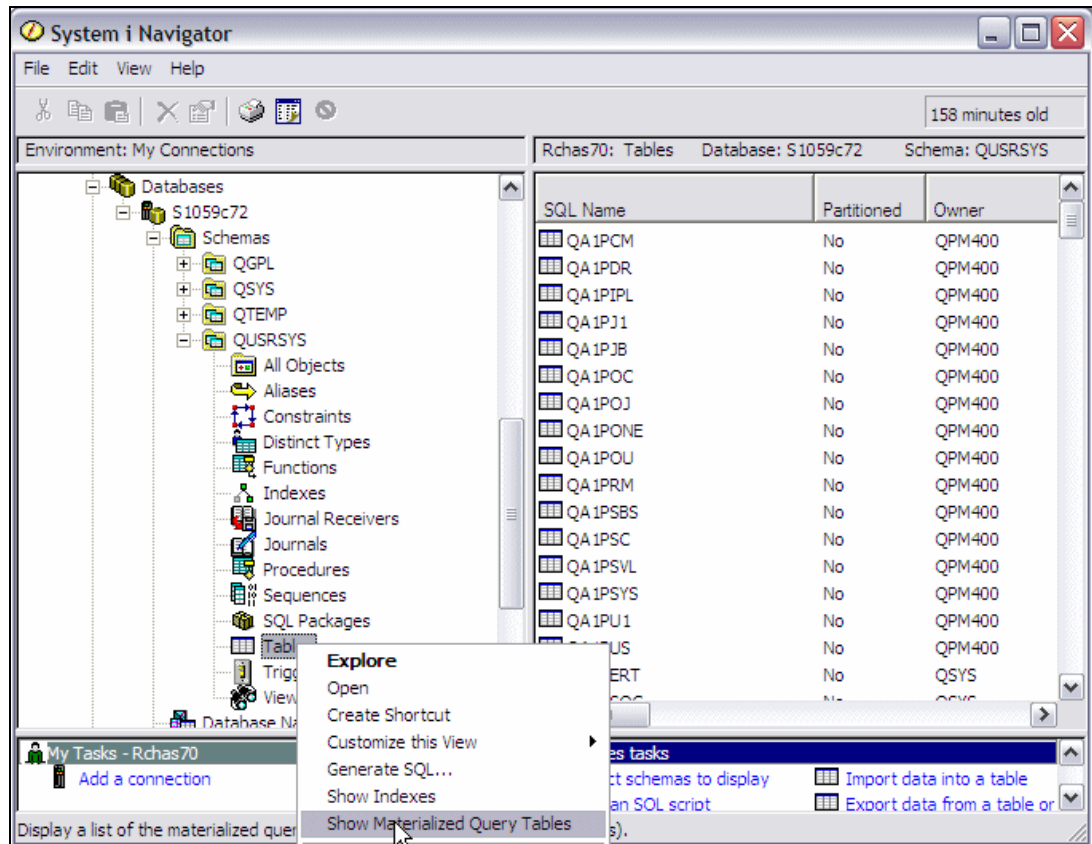


Figure 5-11 Show materialized query tables for all tables in a schema

5.7.7 Show materialized query tables for all views in a schema

To view materialized query tables for all views in a schema, right-click Views in the left pane and select **Show Materialized Query Tables**. This launches the MQT window that contains the indexes.

Figure 5-12 shows materialized query tables for all views in a schema.

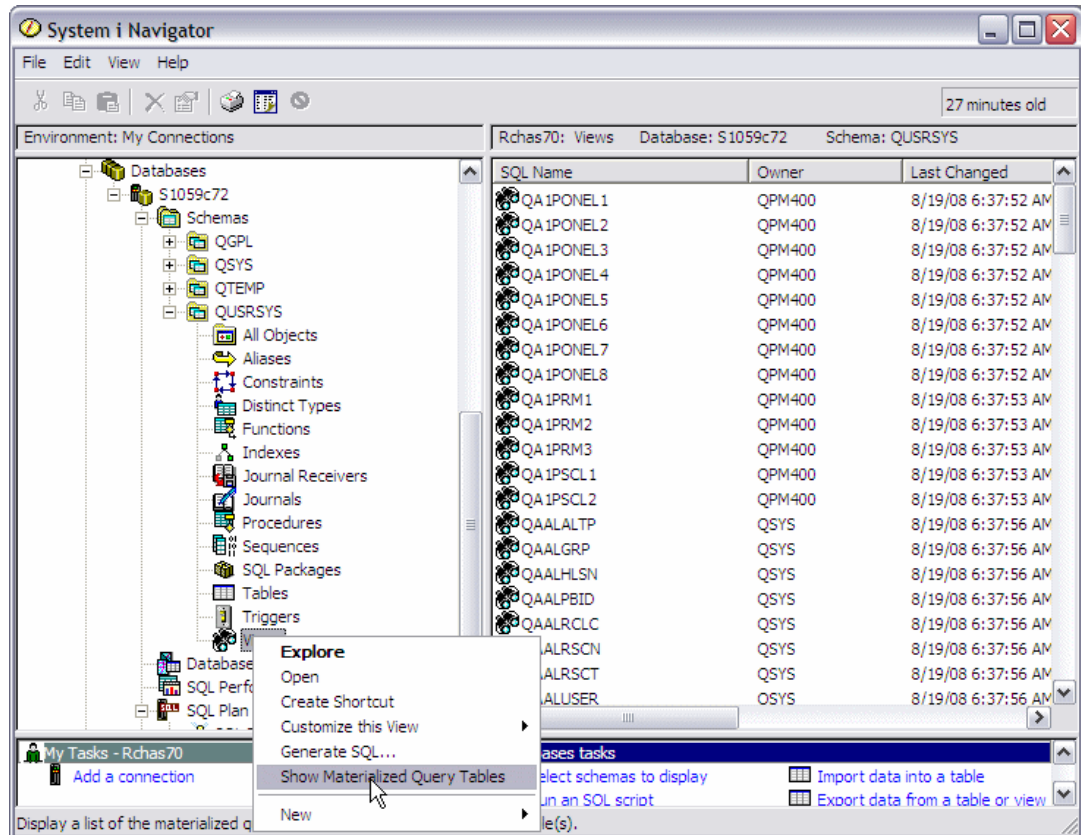


Figure 5-12 Show materialized query tables for all views in a schema

5.7.8 Show materialized query tables for a specific table

To view materialized query tables for a specific table, right-click the table in the Tables window and select **Show Materialized Query Tables**. This launches the MQT window that contains the indexes.

Figure 5-13 shows materialized query tables for a specific table.

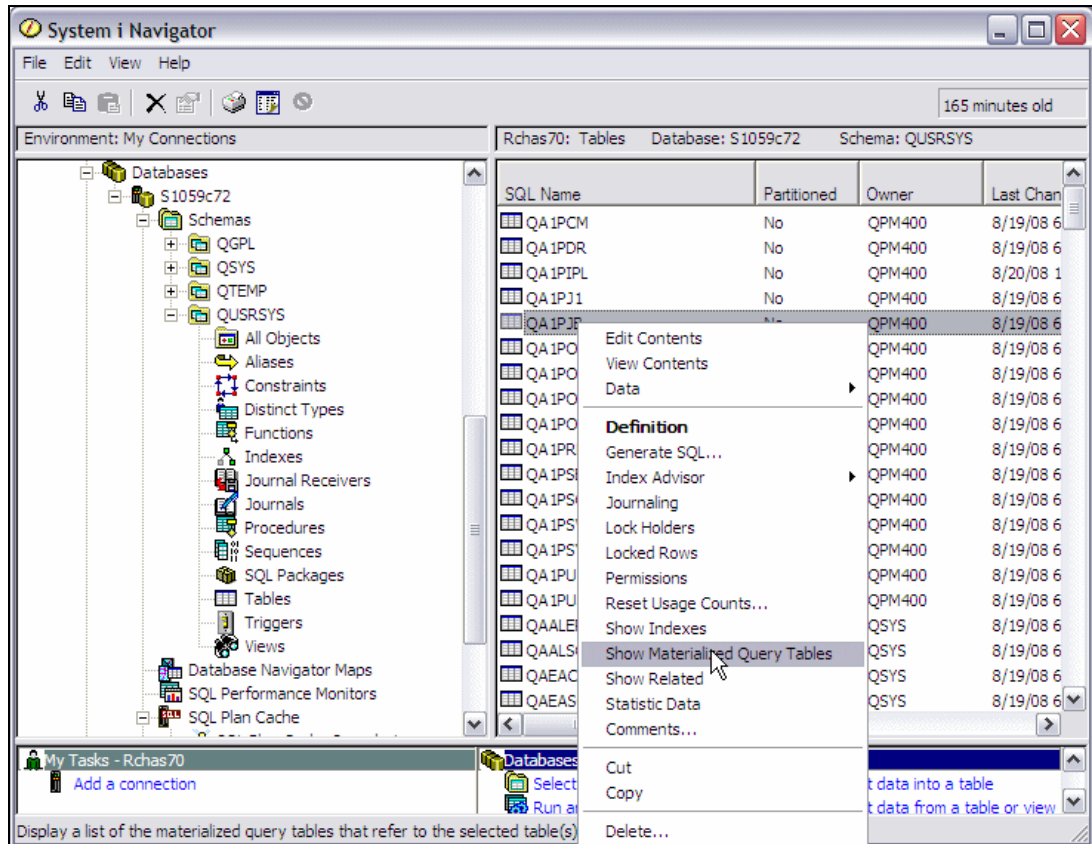


Figure 5-13 Show materialized query tables for a specific table

5.7.9 Show materialized query tables for a specific view

To view materialized query tables for a specific view, right-click the view in the Views window and select **Show Materialized Query Tables**. This launches the MQT window that contains the indexes.

Figure 5-14 shows the options to be selected to show materialized query tables for a specific view.

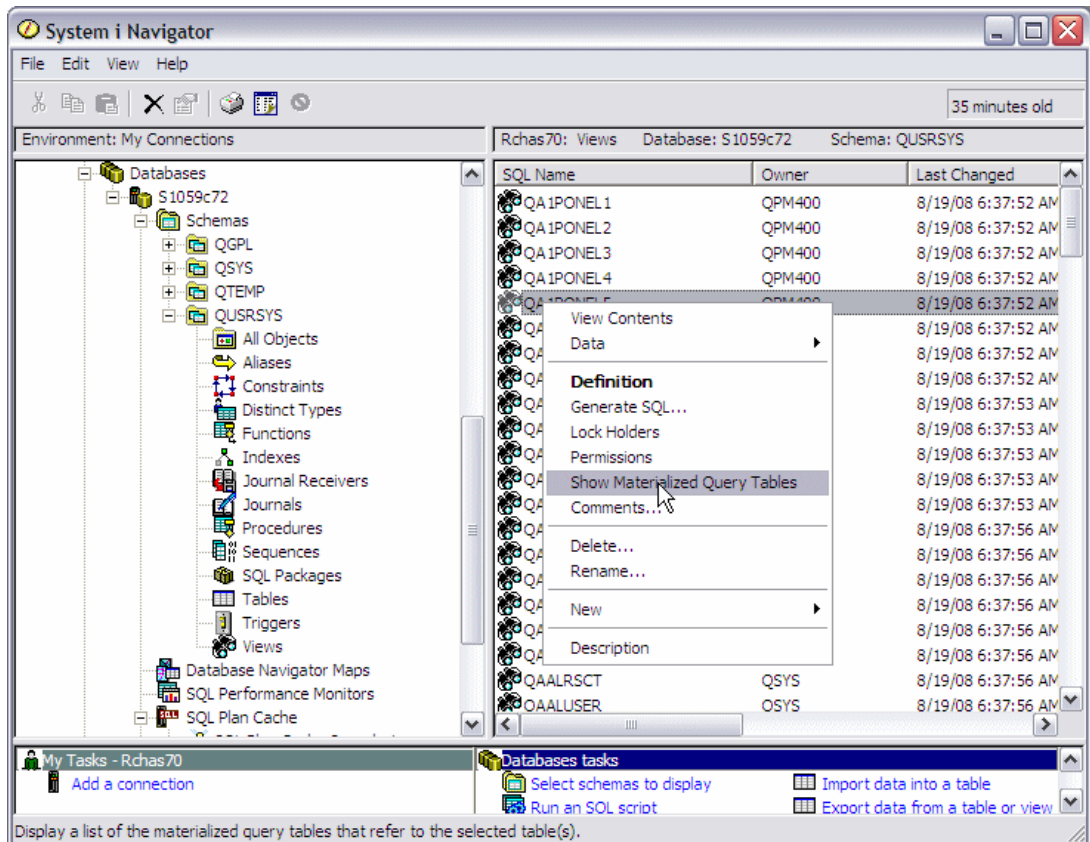


Figure 5-14 Show materialized query tables for a specific view

5.7.11 SQL details for jobs

When working with active jobs, IBM i 6.1 provides additional information about the last SQL statement executed by the job. Figure 5-16 shows how to get SQL details for jobs.

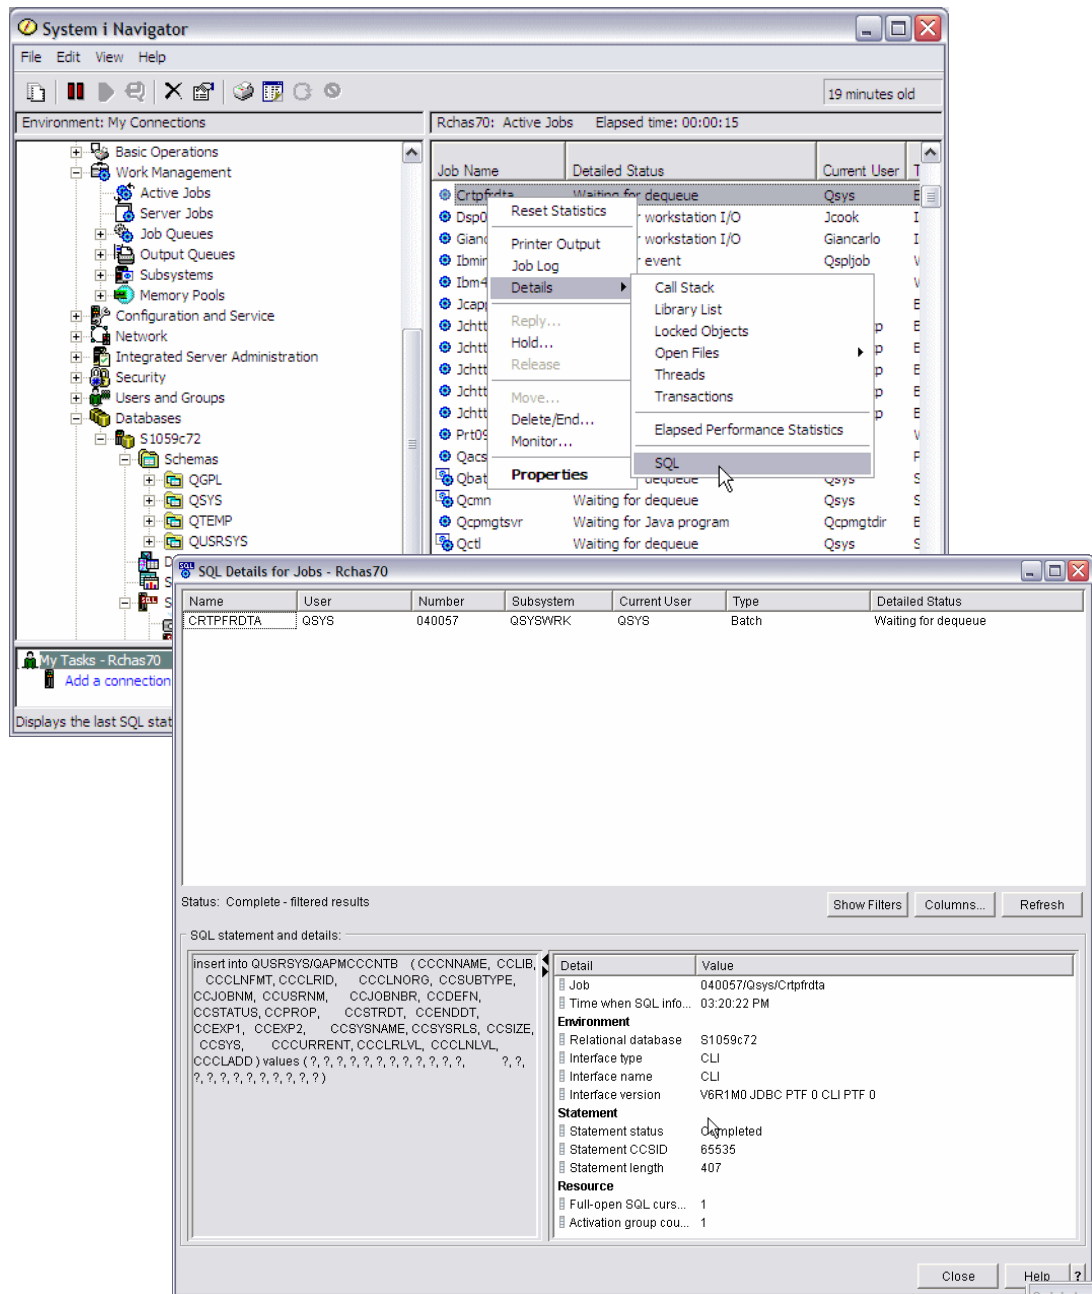


Figure 5-16 SQL details for jobs

5.7.12 Change the SQL plan cache size threshold

The SQL Plan Cache Properties option shows high-level information about the cache, including, for example, cache size, number of plans, and the number of full opens and pseudo opens that have occurred.

This information can be used to view overall database activity. If tracked over time, it provides trends to help you better understand the database utilization peaks and valleys throughout the day and week.

You may edit the Plan Cache Size Threshold property of your plan cache by right-clicking a property and selecting **Edit Value**. Under normal circumstances this properties value is sufficient and altering is not necessary. If it is altered, care should first be taken to assess the performance consequences of the change. Note that the value is restored to its default on an IPL.

Figure 5-17 shows how to change the SQL plan cache.

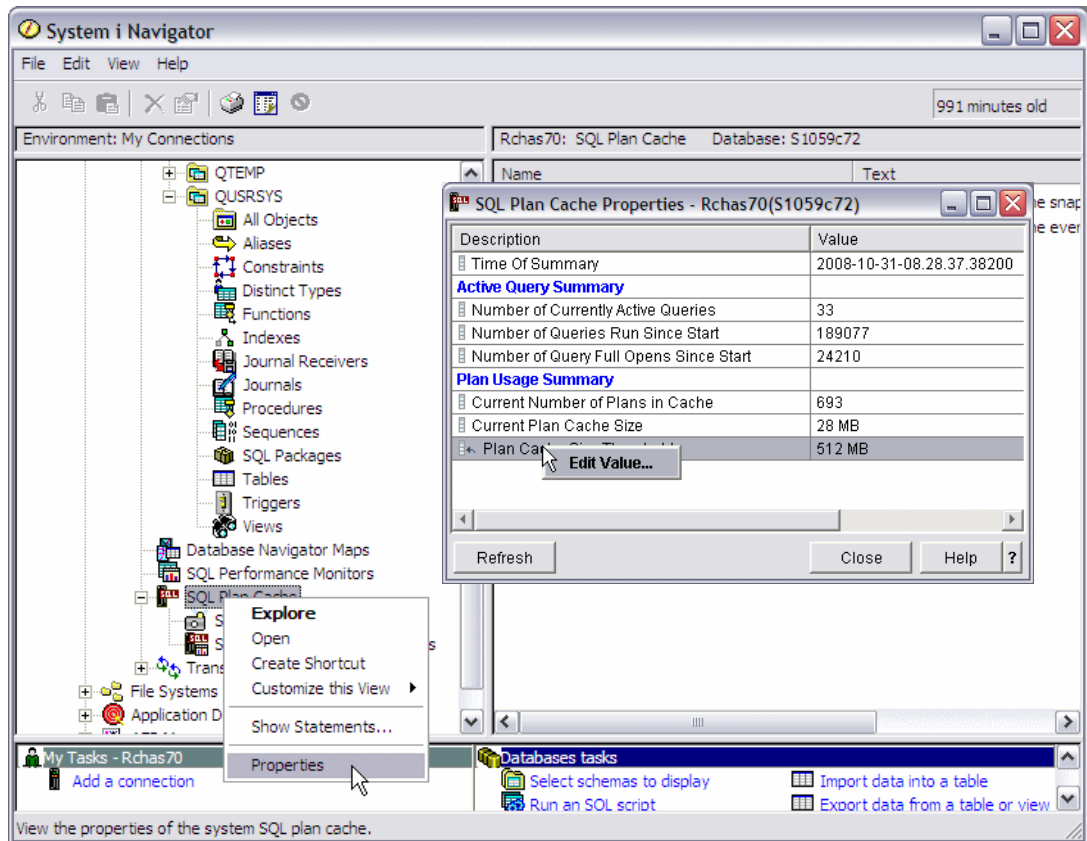


Figure 5-17 Change the SQL plan cache size threshold

5.7.13 SQL plan cache statements

Right-clicking an SQL statement in the Statements frame on the SQL Plan Cache Statements window now provides options to do the following:

- Show Longest Runs

Shows details of up to ten of the longest running instances of that query. In the Longest Run list you can right-click an entry and either visually explain it, save it to a snapshot, or remove it. These are useful if you want to capture the information for that specific run in visual explain. It also provides a way to remove old or superfluous runs to make room to capture future runs. Note that only one plan can be highlighted to perform this action. Also note that any runs removed only affect what is shown in the list, the total time, total number of runs and others, because the plan still includes all the runs, including the runs removed from the list.

- ▶ Show Active Jobs
Displays a list of jobs on the system that are currently using that plan.
- ▶ Show User History
Gives a list of all user IDs that have run that plan along with the last time they ran it.
- ▶ Work with SQL Statement
Brings up a scripting window with the SQL statement in it. This is useful when working with and tuning the statement directly or for just viewing the statement in its own window. Only one plan can be highlighted to perform this action.
- ▶ Save to New
Enables you to create a snapshot of the selected entries.
- ▶ Plan
Provides more options where certain modifying options on the plan can be performed:
- ▶ Change Plan Score
Allows you to set the plan's score to a specific value. The plan score is used to determine when a plan should be removed from the cache. A lower score plan will be removed before a higher score plan. By setting the plan score high, the plan will remain in the cache for a longer period of time. Setting the plan score to a low value will cause the plan to be pruned sooner than might otherwise have occurred.
- ▶ Delete
Allows you to remove the plan immediately from the cache. Note under normal circumstances there should not be a need to modify a plan's attributes. Normal database processing will age and prune plans appropriately. These modifying options are provided mostly as tools for minute analysis and for general interest.

Note: The user and job name information given for each plan entry on the Show Statements panel is the user and job that initially caused the creation of the plan (the user where full optimization took place). This is not necessarily the same as the last user to run that query. The Longest Runs panel, however, does show the particular user and job for that individual run.

Figure 5-18 shows the SQL plan cache statement actions.

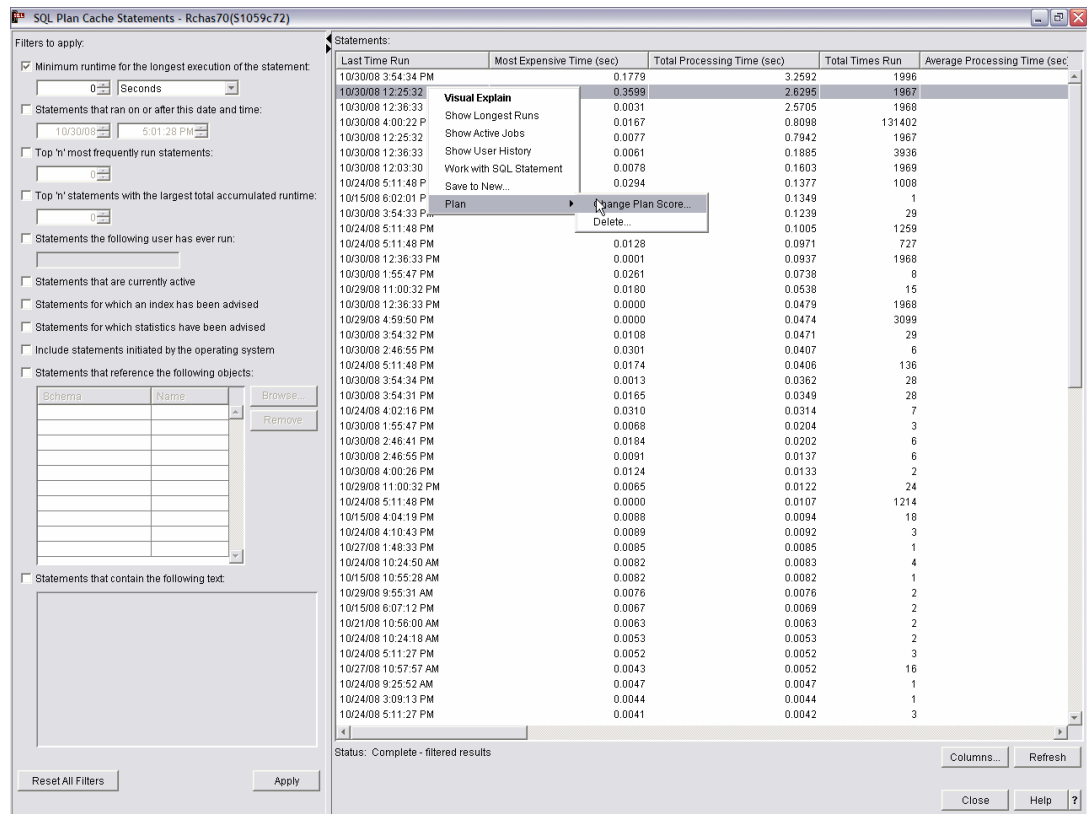


Figure 5-18 SQL plan cache statement actions

5.7.14 SQL plan cache event monitor

The SQL plan cache event monitor records changes in your plan cache.

The SQL plan cache event captures monitor records of plans as they are removed from the plan cache. The event monitor is useful for ensuring that all performance information potentially available in the cache is captured even if plans are removed from the cache. Combining the event monitor output with a plan cache snapshot provides a composite view of the cache from when the event was started until the snapshot was taken.

The event monitor allows the same filtering options as described for show statements. It creates a snapshot with the exception of the Top n most frequently run statements and the Top n statements with the largest accumulated runtime. Once a statement is removed from the cache, it no longer participates in comparison to other plans, therefore these two 'Top n' filters do not make sense for these pruned plans.

To start an SQL plan cache event monitor, right-click SQL Plan Cache Event Monitors in the left pane and select **New** → **Event Monitor**. This invokes the SQL Plan Cache Event Monitor Wizard, where you can define the attributes of the monitor as shown in Figure 5-19 on page 122.

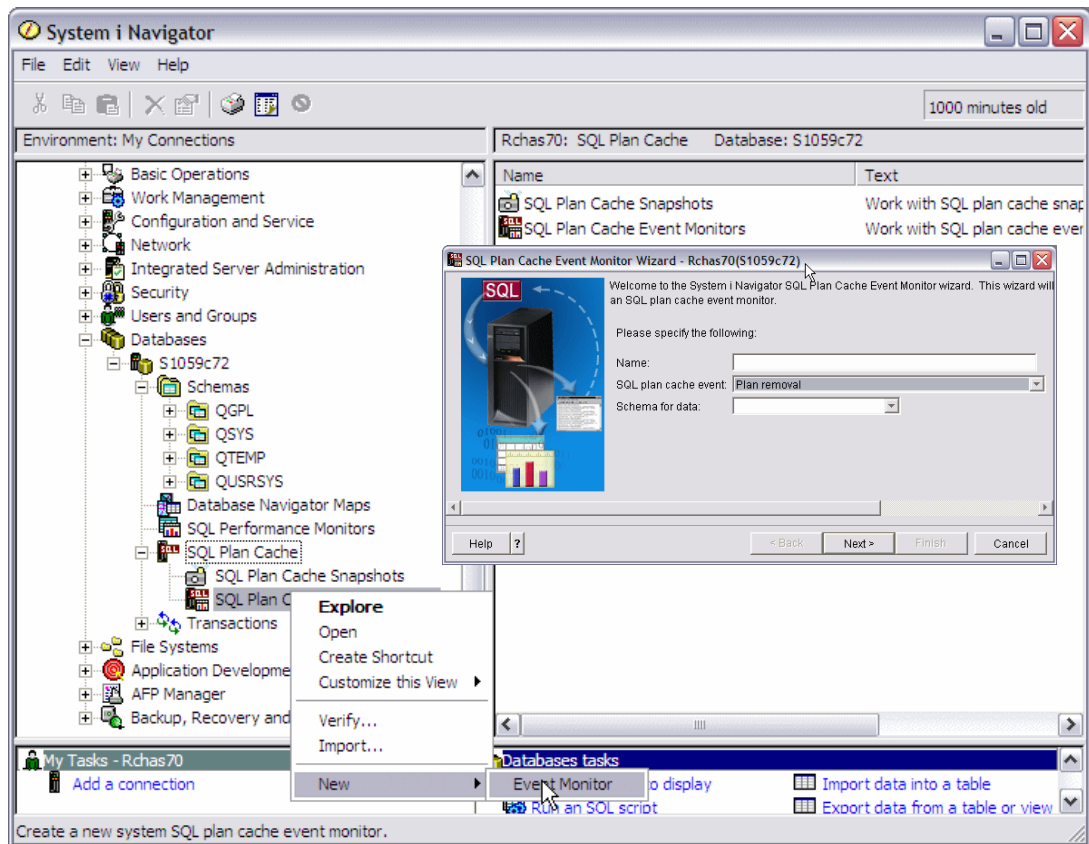


Figure 5-19 Start an SQL plan cache event monitor

5.7.15 Working with an SQL plan cache event monitor

This section describes the SQL plan cache event monitor.

Right-clicking an event monitor in the SQL Plan Cache Event Monitor frame provides options to do the following:

- ▶ Analyze

Launches the SQL Performance Data Analysis window, which provides a summary of the collected data. Various analysis options are available from this window either at the summary or statement level.

- ▶ Show Statements

Launches the Statements window where you can filter the statements or drill down on statement details.

- ▶ Compare

Launches the SQL Performance Data Comparison window, which summarizes the data from each collection and displays the results in a table-like format. To compare event monitors, select two or more monitors, then right-click and select **Compare**.

Figure 5-20 shows a view of working with an SQL plan cache event monitor.

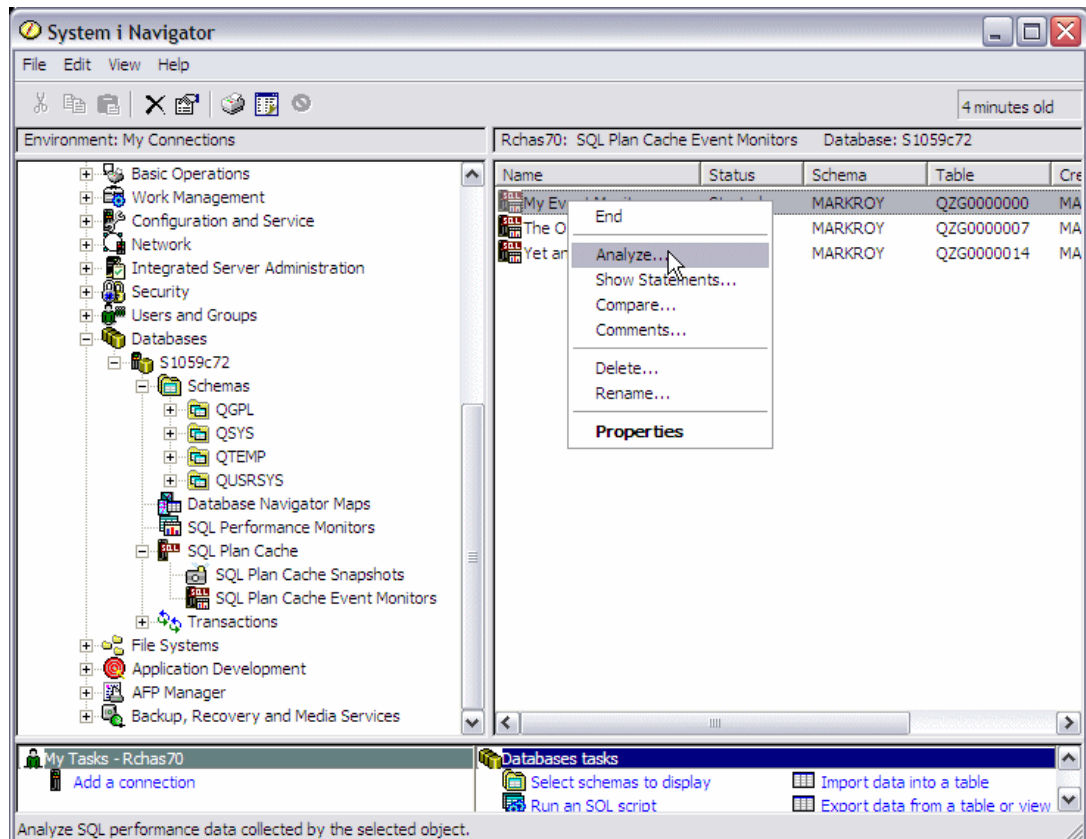


Figure 5-20 Working with an SQL plan cache event monitor



Performance

This chapter discusses the following topics:

- ▶ “Java and WebSphere performance” on page 126
- ▶ “Java Performance Tools” on page 129
- ▶ “IBM Performance Tools for i5/OS (PT1) packaging” on page 134
- ▶ “IBM Performance Tools for i5/OS enhancements” on page 138
- ▶ “IBM iDoctor for IBM i” on page 138
- ▶ “IBM Systems Director Navigator for i” on page 403
- ▶ “Work with System Activity (WRKSYSACT)” on page 180
- ▶ “Performance Explorer” on page 180
- ▶ “Collection Services” on page 182
- ▶ “Performance Management Agent” on page 184
- ▶ “System i Navigator On Demand Performance Center enhancements” on page 184
- ▶ “Workload Estimator” on page 185

6.1 Java and WebSphere performance

In IBM i 6.1 the performance of the 32-bit IBM Technology for Java virtual machine has significantly improved. The same improvements also benefit the new 64-bit IBM Technology for Java virtual machine.

IBM has tested the Java performance improvements with a sample stock trading application called Trade 6.1 running on a POWER5+™ system with IBM i 5.4, IBM Technology for Java 32-bit JVM, DB2 for i, WebSphere 6.1 and compared it with the same Java and WebSphere Application Server levels running under IBM i 6.1. Both JDBC drivers (“native” and Toolbox) were tested.

The results show a significant performance improvement for both JDBC drivers being used. The “native” JDBC driver shows an improvement of 68% over IBM i 5.4, the Toolbox JDBC driver a 78% improvement. One change was made in IBM i 6.1 to support large 64 KB pages which all POWER5+ systems support, but is only available on 6.1 and later. Figure 6-1 shows a comparison of Java performance, release to release.




Figure 6-1 Java performance release to release comparison

6.1.1 IBM Technology for Java, J9

J9 is the IBM Java virtual machine and classlibrary technology. It is used by both the 32-bit and 64-bit IBM Technology for Java JVMs but not the Classic JVM.

The IBM Technology for Java provides a significant reduction in memory requirements compared to the “classic” 64-bit JVM due to more effective use of memory heap. However, it should be noted that the 32-bit JVM effectively supports a maximum heap size of about 3 GB. This could be further reduced when native PASE programs are being used in conjunction with Java applications.

Applications with larger heap requirements may need to use one of the 64-bit JVMs—either IBM Technology for Java or the Classic VM.

Because IBM Technology for Java shares a common implementation with all IBM JVMs, the available tuning parameters are essentially the same on IBM i as on other platforms.

6.1.2 Classic JVM (64-bit)

The performance of the Classic JVM (Java Virtual Machine) in IBM i 6.1 is approximately equivalent to Classic JVM on IBM i 5.4. However, performance improvements in the JDBC driver and database may improve performance of some applications.

Also, Direct Execution (DE) is not supported in 6.1. Applications that used DE in previous releases may have improved performance in 6.1 due to use of the Just In Time (JIT) compiler.

In 6.1 the default JVM is the IBM Technology for Java 32-bit JVM and Java 5, while in IBM i 5.4 the 64-bit Classic JVM and Java 1.4 was the default. After 6.1 is installed, Java applications may show significant performance improvements because the 32-bit JVM and Java 5 are used by default (unless the application overrides the default), without any modifications to the application environment.

Note: IBM i 6.1 is the last release supporting the “classic” JVM, and application developers and providers should be moving to the 32-bit or new 64-bit IBM Technology for Java JVM.

6.1.3 IBM Technology for Java (32-bit)

On IBM i 6.1 performance of IBM Technology for Java (32-bit) is 10-15% improved over current IBM i 5.4.

Additional improvements of 10% or more are possible for some applications by using 64-KB pages. This can be achieved by specifying the `-X1p64k` flag on the Java command line prompt to enable large page support during start-up of the application. Without this new flag, applications will use 4-KB pages. 64-KB pages have the most performance impact when the heap size is greater than 1.5 GB.

In IBM i 6.1, IBM Technology for Java (32-bit) generally offers a 30% or more performance improvement and typically a 40% or more reduction in main memory footprint, relative to Classic JVM.

For more about Java support functional enhancements supported by IBM i 6.1 in this publication, see:

- ▶ 13.3, “WebSphere Application Server Express V7.0” on page 305
- ▶ 6.1, “Java and WebSphere performance” on page 126

6.1.4 IBM Technology for Java (64-bit)

In addition to the 32-bit JVM provided in IBM i 5.4, a new 64-bit IBM Technology for Java JVM is provided in 6.1.

This JVM provides the improved performance known in the 32-bit JVM and the scalability of the “classic” 64-bit JVM. This scalability comes at a cost and causes about a 10% performance overhead relative to the 32-bit JVM, but allows the application memory heap to grow to the size of the physical memory available on the system.

Due to the large memory heap being supported in the 64-bit JVM, the Java applications can benefit from the large address space and large caches now available.

The new JVM has the same underlying implementation as 32-bit, but with no practical limit to the heap size. The main difference is that object references are 64 bits instead of 32 bits. The same command line parameters work for both JVMs.

Typically, if an application is changed from IBM Technology for Java 32-bit to IBM Technology for Java 64-bit, the heap is 70% larger. Therefore, heap sizes should be increased for 64-bit JVM. Increasing heap size may improve 64-bit performance at the cost of a larger memory footprint. Larger address space may allow for larger caches, which increases chances for improved performance.

The new footprint is generally smaller than the Classic JVM (reductions up to 10% may be seen).

6.1.5 Choosing which JVM to use

In general, the 32-bit IBM Technology for Java JVM should primarily be used. However, functional differences may influence the decision. These functional differences include:

- ▶ Only Classic supports JDK 1.4.2 or WebSphere Application Server 6.0.
- ▶ IBM Technology for Java (32-bit) cannot call 64-bit PASE Native Methods.
- ▶ IBM Technology for Java (64-bit) cannot call 32-bit PASE Native Methods.

If the application has dependencies to prior levels of JDK, for example Java 1.4 (or earlier) or use of certain releases of WebSphere, the use of the classic 64-bit JVM might be necessary. However, as mentioned previously, IBM 6.1 is the last release to support the classic JVM.

If functional differences are not a concern, use IBM Technology for Java (32-bit) if the application will “fit” in a 32-bit environment. If the current heap uses up to about 5 GB using the “classic” JVM, it would be possible to consider the 32-bit JVM. Otherwise, use IBM Technology for Java (64-bit).

6.1.6 Comparing Java 5.0 to Java 6

The differences between Java 5 and 6 are minimal using the “classic” JVM. However, using the IBM Technology for Java JVM, 10-20% could be experienced in some applications. Some garbage collection parameters might have to be tuned in order to achieve those results.

Note: WebSphere 6.1 does not support Java 6. WebSphere 7.0 does require Java 6.

6.1.7 Determining JVM heap size for IBM Technology for Java JVMs

The following topics describe how to determine JVM heap size for IBM Technology for Java JVMs.

Moving an application from the Classic VM to IBM Technology for Java

You can use a tool such as Dump Java Virtual Machine (DMPJVM) or verbose garbage collection (GC) to determine how large the heap grows when running your application. This value can be used as the maximum heap size for 64-bit IBM Technology for Java.

In 32-bit IBM Technology for Java, about 75% of this value is a reasonable starting point. For example, if your application’s heap grows to 256 MB when running in the Classic VM, try setting the maximum heap size to 192 MB when running in the 32-bit VM. The initial heap size can be set to about half of this value—96 MB in our example.

These settings are unlikely to provide the best possible performance or the smallest memory footprint, but the application should run reasonably well. Additional performance tests and tuning could result in better settings.

New Java applications using the IBM Technology for Java

Using about 25% of the maximum heap often provides reasonable performance. The default is 4 MB minimum and 2 GB maximum heap size. Again, additional performance tests and tuning could result in better settings.

6.2 Java Performance Tools

IBM i 6.1 introduces two main performance tooling changes for IBM Technology for Java, Collection Services integration and new commands such as New Work with JVM Jobs (WRKJVMJOB). The following topics describe these changes.

6.2.1 Collection Services integration

A new QAPMJVM performance database file in 6.1 lists information about each IBM Technology for Java VM active during the collection. The data includes:

- ▶ Which JVM (32-bit or 64-bit) and version was used
- ▶ Garbage Collection policy, cycle count, cycle time
- ▶ Various heap and non-heap memory statistics

For further information about the data contained in QAPMJVM, refer to the “Collection Services data files: QAPMJVM” topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahx/rzahxqapmjvm.htm?resultof=%22%51%41%50%4d%4a%56%4d%22%20>

6.2.2 The Work with JVM Jobs (WRKJVMJOB) command

This command allows the user to monitor Java Virtual Machines (JVMs) running in active jobs. If the values specified for the Job name (JOB) parameter identify a single job, a menu is displayed showing the options available for the specific JVM job. Otherwise, a list of jobs with active JVMs is displayed.

When a specific job is selected, several options are available to retrieve configuration and environment variables of a particular JVM. The options are:

- ▶ Arguments and options used when the JVM was started
- ▶ Environment variables for both ILE and PASE
- ▶ Java lock requests outstanding for the JVM job
- ▶ Garbage collection information
- ▶ Properties when the JVM was started
- ▶ Properties of the JVM currently running
- ▶ List of threads associated with the JVM
- ▶ Partially completed job log for the JVM job
- ▶ Work with spooled input and output files for the JVM job

- ▶ Generate JVM (system, heap, Java) dumps from a panel option
- ▶ Enable and disable verbose garbage collection

Some of the most important pieces of information would be the Java lock requests, the garbage collection, the list of threads, and the current set properties for the JVM.

Note: This command only works for IBM Technology for Java (not Classic).

In Figure 6-2 we use the command to determine the accumulated time spent in garbage collection. The accumulated time, relative to the length of time a JVM has been running, may indicate a poorly tuned JVM. In our example, we know the name of the JVM but not the job number, so we start with the command without specifying any parameters. Figure 6-2 shows an example of the Work with JVM Jobs (WRKJVMJOB) command.

```

Work with JVM Jobs (WRKJVMJOB)

Type choices, press Enter.

Job name . . . . . *ALL      Name, generic*, *ALL
User . . . . . *ALL      Name, generic*, *ALL
Number . . . . . *ALL      000000-999999, *ALL
Option . . . . . *SELECT   *SELECT, *CRTOPT, *ENVVAR...

```

Figure 6-2 Work with JVM Jobs (WRKJVMJOB)

Figure 6-3 on page 131 shows the results options available on the Work with JVM Jobs (WRKJVMJOB) command display. From here, we select the option to work with the job we are interested in.

```

Work with JVM Jobs
RCHAS65
11/07/08 16:43:20

Active JVMs on system: 9

Type options, press Enter.
 5=Work with 7=Display job log 8=Work with spooled files
 9=Display GC information 11=Display threads 12=Dump 13=Print

Opt Job Name User Number Function Status
QSRVMON QSYS 245000 JVM-ServiceMon DEQW
QJVACMSRV QIBMHELP 245088 JVM-org.eclips THDW
WSERVICE1 QWSERVICE 245474 JVM-com.ibm.lw THDW
JSCOM3C QWEBQRYADM 249267 PGM-jvmStartPa TIMW
WQLWI7 QWEBQRYADM 249295 JVM-com.ibm.lw THDW
SERVER1 QEJBSVR 252206 PGM-jvmStartPa THDW
ADMIN3 QLWISVR 266587 JVM-com.ibm.lw THDW
5 ADMIN QLWISVR 266591 JVM-com.ibm.lw THDW
ADMIN2 QLWISVR 266595 JVM-com.ibm.lw THDW

Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Print F9=Retrieve
F11=Display subsystem information F12=Cancel F16=Resequance

Bottom

```

Figure 6-3 Work with JVM Jobs

The option to view garbage collection information is on the first panel presented, as shown in Figure 6-4. If other options are needed, they can be displayed with the Page Down key.

```

Work with Java Virtual Machine
System: RCHAS65

Job . . . . . : ADMIN PID . . . . . : 18205
User . . . . . : QLWISVR JDK . . . . . : 1.5.0
Number . . . . . : 266591 Bits . . . . . : 32

Select One of the following:

1. Display JVM arguments
2. Display environment variables
3. Display PASE environment variables
4. Display Java lock information
5. Display garbage collection information
6. Display initial Java system properties
7. Display current Java system properties
8. Display Java threads
9. Display job log

20. Work with spooled files

Selection or command
====> 5
F3=Exit F4=Prompt F9=Retrieve F12=Cancel

More...

```

Figure 6-4 Work with Java Virtual Machine

The initial Display Garbage Collection Information display shows information about the heap size. Figure 6-5 shows the Display Garbage Collection Information.

```

Display Garbage Collection Information
System: RCHAS65
Job . . . . . : ADMIN          PID . . . . . : 18205
User . . . . . : QLWISVR       JDK . . . . . : 1.5.0
Number . . . . . : 266591      Bits . . . . . : 32

Garbage collected heap:
Initial heap size . . . . . : 64.000M
Maximum heap size . . . . . : 2048.000M
Current heap size . . . . . : 64.489M
Heap in use . . . . . : 48.007M

Other memory:
Internal memory size . . . . . : 164.644M
JIT memory size . . . . . : 16.000M
Shared classes memory size . . . . . : 0.000M

More...
F3=Exit  F5=Refresh  F6=Print  F7=Print GC tables  F12=Cancel

```

Figure 6-5 Display Garbage Collection Information - part 1

We use the Page Down key to display the GC information we are looking for, as shown in Figure 6-6.

```

Display Garbage Collection Information
System: RCHAS65
Job . . . . . : ADMIN          PID . . . . . : 18205
User . . . . . : QLWISVR       JDK . . . . . : 1.5.0
Number . . . . . : 266591      Bits . . . . . : 32

General GC information:
Current GC cycle . . . . . : 36
GC policy type . . . . . : OPTTHRUPUT
Current GC cycle time . . . . . : 304
Accumulated GC time . . . . . : 2418

Bottom
F3=Exit  F5=Refresh  F6=Print  F7=Print GC tables  F12=Cancel

```

Figure 6-6 Display Garbage Collection Information - part 2

6.2.3 The Print JVM Job (PRTJVMJOB) command

The Print JVM Job (PRTJVMJOB) command allows you to print Java Virtual Machines (JVMs) running in active jobs.

The Print JVM Job command prints most of the parameters and variables displayed by the WRKJVMJOB command. However, since the command invokes a list function, more history information can be provided through the generated spool file.

The additional information shown (compared to WRKJVMJOB) includes:

- ▶ Garbage collection cycle reason, time spent, and more
- ▶ Sweep and compact times
- ▶ Tenured object space used information
- ▶ Finalizer, Soft, Phantom object information

A good example of this is the garbage collection information, where the display only lists the current GC cycle, the printed list provides detailed information on each of the GC cycle since the start of the JVM. This is particularly useful in determining when the application reaches a level point. It does not place additional memory requirements on the heap memory. Alternatively, in a few cases, it could show an ever-increasing demand for more memory in the heap, because of a potential Java object leak.

Note: This command only works for IBM Technology for Java (not Classic).

6.2.4 The Generate JVM Dump (GENJVMDMP) command

The Generate JVM Dump (GENJVMDMP) command generates Java Virtual Machine (JVM) dumps upon request. Three types of dumps are available:

Java dump (*JAVA) Generates multiple files that contain diagnostic information for the JVM and the Java applications running within the JVM. The IBM tool, IBM Thread, and Monitor Dump Analyzer for Java Technology can analyze javacore and diagnose monitor locks and thread activities in order to identify the root cause of hangs, deadlocks, and resource contention or monitor bottlenecks.

System dump (*SYSTEM)

Generates a binary format raw memory image of the job that was running when the dump was initiated. This dump is primarily used by service personnel to debug functional issues, but sometimes can give clues on performance problems as well, for example running out of memory.

Heap dump (*HEAP) Generates a dump of all the heap space allocations that have not yet been freed. The IBM tool Memory Dump Diagnostic for Java (MDD4J) can be used to further analyze the collected data.

The easiest and recommended way to find the location of the dumps is to look at the joblog using Work with JVM Jobs (WRKJVMJOB).

Notes:

- ▶ The dump request typically stalls the application while dump data is generated. The application continues to run after the dump data collection has completed.
- ▶ This command only works for IBM Technology for Java (not Classic).

6.3 IBM Performance Tools for i5/OS (PT1) packaging

In IBM i 6.1 the Performance Tools licensed program (5761-PT1) is packaged differently from previous versions of IBM i.

- ▶ BASE - IBM Performance Tools for i5/OS
- ▶ Option 1 - Manager Feature
Full-function and now includes Disk Watcher.
- ▶ Option 2 - Agent Feature
Subset of Manager Feature excludes:
 - Performance and trace reports
 - Performance utilities (job traces and the select file utilities)
 - Performance graphics (5250 display station graphs)
- ▶ Option 3 - Job Watcher
 - English only.
 - This is not equivalent to the Job Watcher feature of the iDoctor service offering.

Note: The Job Watcher and Disk Watcher functions in IBM Performance Tools for i5/OS collect and present data similar to the IBM iDoctor for the IBM i product and service offering.

IBM iDoctor for IBM i has more features than IBM Performance Tools for i5/OS and is a more complex product to use. The Job Watcher and Disk Watcher features in IBM Performance Tools for i5/OS are for those new to performance analysis, whereas the similar features in IBM iDoctor for i are intended for performance analysis specialists.

For example, here are some orderable combinations that will work:

- ▶ *Base + Option 1 (includes Disk Watcher)
- ▶ *Base + Option 1 + Option 3 (includes Disk Watcher and Job Watcher)
- ▶ *Base + Option 2
- ▶ *Base + Option 2 + Option 3 (includes Job Watcher)
- ▶ *Base + Option 3 (includes Job Watcher)

The main differences in the packaging structure are:

- ▶ The Disk Watcher data viewing options are now included in the Performance Tools Manager Option (Option 1).
- ▶ A new option, Job Watcher (Option 3).

The full suite of Performance Tools is only available when both the Performance Tools Manager Option (Option 1) and Performance Tools Job Watcher Option (Option 3) are installed (in addition to the base functions).

Customers using IBM i 6.1 without 5761-PT1 are still able to collect and save performance collections and have that data investigated on another system that does have 5761-PT1. To enable the collection of the specific Job Watcher and Disk Watcher performance data on

partitions or systems that have no Performance Tools licensed product installed, the following commands are now delivered with IBM i 6.1:

- ▶ Job Watcher commands:
 - Add Job Watcher Definition (ADDJWDFN)
 - End Job Watcher (ENDJW)
 - Remove Job Watcher Definition (RMVJWDFN)
 - Start Job Watcher (STRJW)
- ▶ Disk Watcher commands:
 - Add Disk Watcher Definition (ADDDWDFN)
 - End Disk Watcher (ENDDW)
 - Remove Disk Watcher Definition (RMVDWDFN)
 - Start Disk Watcher (STRDW)

Table 6-1 shows the functions and the components required for each.

Table 6-1 Performance data collection and analysis tools

Functional area	Function	Tool	Component required
Data collection	Collection services	STRPFCOL	Base IBM i
	Disk Watcher	STRDW	Base IBM i
	Job Watcher	STRJW	Base IBM i
	Performance Explorer	STRPEX	Base IBM i
Data analysis - reporting	Collection services	PRTnnnRPT	PT1 Opt 1 Manager
	Disk Watcher	N/A	N/A
	Job Watcher	N/A	N/A
	Performance Explorer	PRTPEXRPT	Base IBM i
Data analysis - graphical	Collection services	Investigate data	Base IBM i
		Collection Services Investigator	iDoctor
	Disk Watcher	Display performance data	System i navigator
		Investigate data	PT1 Opt 1 Manager
	Job Watcher	Disk Watcher	iDoctor
		Investigate data	PT1 Opt 3 Job watcher
	Performance Explorer	Job Watcher	iDoctor
		PEX Analyzer	iDoctor
		PTDV	alphaWorks® download

Note: In all cases, data analysis can be conducted using user-defined queries using products such as Query and SQL.

6.3.1 Disk Watcher

Disk Watcher provides you with the ability to collect either a statistical summary of disk performance data or a trace of all disk I/O events that occur on a system. The trace mode is recommended because it provides more options for analyzing the data and determining potential disk problems.

The Disk Watcher GUI in iDoctor and the browser interface in IBM Systems Director Navigator for i provide many graphs with drill-downs for each mode of collection (statistical or trace). Using Disk Watcher you can take a trace, then summarize the trace data into an interval size desired for the purpose of easily graphing the statistics at either broad or detailed level.

The Disk Watcher GUI in the iDoctor product is available for i5/OS V5R3 and higher. The Disk Watcher GUI in the IBM i 6.1 Investigate Data feature will support Disk Watcher data from the previous releases; however, the Disk Watcher data must include the hardware information file, and also be moved to reside on an IBM i 6.1 partition for analysis. At IBM i5/OS V5R3 and IBM i 5.4, the required Disk Watcher PTFs must be installed to add the Disk Watcher data collection commands to IBM i. At IBM i 6.1 the Disk Watcher data collection commands are included.

6.3.2 Job Watcher

Job Watcher returns real-time information about all jobs, threads, and/or LIC tasks running on a system (or on a selected set of jobs/threads or tasks). The data is collected by a server job, stored in database files, and displayed on a client through the iDoctor GUI or through the browser interface in IBM Systems Director Navigator for i. Job Watcher is similar in sampling function to the system commands WRKACTJOB and WRKSYSACT in that each *refresh* computes delta information for the ending snapshot interval. Refreshes can be set to occur automatically, as frequently as every 100 milliseconds. The data harvested from the jobs/threads/tasks being watched is done so in a nonintrusive manner (similar to WRKSYSACT).

Using the iDoctor GUI or the browser interface in IBM Systems Director Navigator for i, the Job Watcher data is summarized to show high-level overviews of system performance over time. From these overview charts you can select a time period of interest and drill down. The drill-down graphs from the overview charts into rankings graphs to show the job/thread experiencing the highest amount of work for the desired statistic. From the rankings graphs, you can select one or more job/threads to show how they performed over time.

Job Watcher, like Collection Services in IBM i 6.1, collects information about *waits* and summarizes them in *wait buckets*. These buckets consist of waits that are generally considered good or bad, and seeing the bad ones on a graph makes it easy to identify problem areas for further investigation.

Additional information harvested by Job Watcher includes:

- ▶ Standard WRKSYSACT type info such as CPU, DASD I/O breakdown, DASD space consumption, and others
- ▶ Some data previously only seen in Collection Services: *real* user name, seize time, breakdown of what types of waits (all waits) that occurred

- ▶ Some data not available anywhere else in real time: details on the current wait (duration, wait object, conflicting job info, specific LIC block point ID), 1000 level deep invocation stack including LIC stack frames
- ▶ SQL statements, host variables, communications data, and activation group statistics

6.3.3 Commands moved to IBM i

In IBM i 6.1, the following commands have moved from IBM Performance Tools for IBM i to IBM i:

Performance Review

The Work with System Activity (WRKSYSACT) command allows you to interactively work with the jobs and tasks currently running in the system. Besides having the capacity to view this data on the display station, you may also direct the data to be stored in a database file for future use.

Note: For information about the function enhancements made to WRKSYSACT, refer to 6.7, “Work with System Activity (WRKSYSACT)” on page 180.

Performance Data

The following performance functions have been moved to IBM i:

▶ Collection Services

Copy Performance Data (CPYPFRDTA), Delete Performance Data (DLTPFRDTA), and the Performance Collection (*nnn*PFRCOL) commands. These commands perform common collection management activities.

Convert Performance Collection (CVTPFRCOL). This command converts performance data from a previous release to the format needed for processing by the current release of IBM Performance Tools for i5/OS.

Note: For information about the function enhancements made to CVTPFRCOL, refer to 6.8.10, “Convert Performance Collection (CVTPFRCOL)” on page 182.

Create Performance Summary (CRTPFRSUM). This command creates additional performance database files containing summary information for an existing collection services performance database file collection. It is also required to “Investigate Data” with IBM Systems Director Navigator for i.

▶ Performance Explorer

Start Performance Explorer (STRPEX), End Performance Explorer (ENDPEX), the PEX Definition (*nnn*PEXDFN), and the Print PEX Report (PRTPEXRPT) commands. These commands perform common performance explorer management and reporting activities.

▶ Disk Watcher

Disk Watcher (*nnn*DW) and Disk Watcher Definition (*nnn*DDWDFN) commands. These commands perform common disk watcher management activities.

6.4 IBM Performance Tools for i5/OS enhancements

The following sections describe IBM Performance Tools for IBM i (5761-PT1) enhancements, including Print System Report (PRTSYSRPT) and wait buckets.

6.4.1 Print System Report (PRTSYSRPT)

A system report generated in 6.1 will include a new section covering the basic performance metrics of the entire system (that is, across all logical partitions).

6.4.2 Wait buckets

Wait buckets are used to account for the time a job (this means a task, primary thread, or secondary thread) spends waiting and to provide some indication as to the type of wait. Since the reasons for a wait are too numerous to handle individually, they are grouped into sets of functionally related waits.

In IBM i 6.1, 16 additional wait buckets have been added, bringing the total to 32.

Note: Wait buckets from previous releases cannot be mapped to 6.1 wait buckets. To analyze wait bucket data from previous releases you can use the corresponding release level of the iDoctor GUI.

For more information about wait buckets, refer to the “Collection Services data files: QAPMJOBWT” topic in Infocenter at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahx/rzahxqapmjobwt.htm?resultof=%22%77%61%69%74%22%20%22%62%75%63%6b%65%74%73%22%20>

6.5 IBM iDoctor for IBM i

IBM iDoctor for IBM i is a suite of performance tools used by IBM and customers to productively collect and analyze performance data in order to quickly solve performance problems on IBM i. The tools may be used to monitor overall system health at a high level or for analyzing performance details within jobs, disk units, and programs collected. When analyzing performance data at a high level there are many drill-down options to assist you with the most logical next step listed first.

IBM iDoctor for IBM i has been used for many years by groups within IBM: the IBM Rochester Support Center, the System i Benchmark Center, as well as IBM Lab Services (and others) for performance consulting work. Through the use of the tools by these groups and customer experiences, iDoctor has grown to become one of the top tools relied on for solving difficult performance issues on IBM i.

IBM iDoctor for IBM i consists of the following components:

- ▶ IBM iDoctor for IBM i Job Watcher
- ▶ IBM iDoctor for IBM i Job Watcher-Collection Services Investigator (or Collection Services Investigator)
- ▶ IBM iDoctor for IBM i Job Watcher-Disk Watcher (or Disk Watcher)
- ▶ IBM iDoctor for IBM i PEX-Analyzer

- ▶ IBM iDoctor for IBM i Heap Analyzer (or Heap Analyzer)

Note: The Job Watcher and Disk Watcher functions in IBM Performance Tools for i5/OS collect and present data similar to the IBM iDoctor for the IBM i product and service offering.

IBM iDoctor for IBM i has more features than IBM Performance Tools for i5/OS and is a more complex product to use. The Job Watcher and Disk Watcher features in IBM Performance Tools for i5/OS are for those new to performance analysis, whereas the similar features in IBM iDoctor for i are intended for performance analysis specialists.

All components require IBM i 6.1 with the required PTFs for each installed. The required PTFs are listed on the iDoctor Web site at:

https://www-912.ibm.com/i_dir/iDoctor.nsf

The Web site also contains prerequisite information for earlier releases of IBM i.

The following section describes the changes to iDoctor for the IBM i 6.1 release.

6.5.1 Name changes

The following product names have changed:

- ▶ iDoctor for iSeries is now IBM iDoctor for IBM i.
- ▶ iDoctor for iSeries Job Watcher is now IBM iDoctor for IBM i Job Watcher.
- ▶ iDoctor for iSeries PEX-Analyzer is now IBM iDoctor for IBM i PEX-Analyzer.
- ▶ iDoctor for iSeries Heap Analysis Tools for Java is now IBM iDoctor for IBM i Heap Analyzer.

6.5.2 New components

The following components are now available:

- ▶ IBM iDoctor for IBM i Job Watcher-Collection Services Investigator is now available at IBM i5/OS V5R3, IBM i 5.4 and 6.1.
- ▶ IBM iDoctor for IBM i Job Watcher-Disk Watcher is now available at IBM i5/OS V5R3, IBM i 5.4 and 6.1.

6.5.3 Graphical user interface changes in IBM iDoctor

The following topics describe graphical user interface (GUI) changes in IBM iDoctor:

General

The following topics describe general GUI changes for IBM iDoctor:

- ▶ In multithreading support in iDoctor, you can now open up multiple reports and while the queries are running on the server, perform other tasks on the client.
- ▶ A new menu option called *Filter libraries...* from the component icon to filter the list of libraries shown. This setting applies to all components. This library filter is shown next to the component's name in the tree when active.
 - iDoctor can now handle SQL statements that contain SQL (long) table names.

- ▶ The toolbars in the main window and in the Data Viewer are now dockable (can move them to wherever you want).

Install

The following topics describe installation changes:

- ▶ Added check in the install up front for system value QALWOBJRST.
- ▶ Add/Remove programs in the Windows control panel now properly recognizes the iDoctor GUI. You can use this option to uninstall the client.

SQL Editor

The following topics describe changes for the SQL Editor:

- ▶ In the SQL Editor, added the menu item *Launch SQL in Run SQL Scripts...* that will take the current SQL statement, place it in a temporary file and launch it in Run SQL Scripts for the purpose of either running it within that interface or using Visual Explain to analyze it. Note that iDoctor uses System naming convention (library/file) and System i Navigator uses SQL naming convention (library.file). You will have to change the default connection properties (Connection → JDBC Setup... Format tab, Naming Convention) in order for the queries created in iDoctor to be executable.
- ▶ When opening a report from the list portion of a component tree or list window using the popup menu, added an option called *Edit* below the normal *Open table* or *Open graph* menu. This option allows you to quickly open a user-defined or iDoctor-supplied table or graph such that the SQL Editor is initially open without actually attempting to run the query. If the query will take a long time to run, this can be used to open and modify the SQL statement first. This technique can also be used in graphs to modify the graph settings before running the query behind the graph.
- ▶ The popup menu in the SQL Editor has a new option, *Execute in batch*. This will run the SQL statement in a separate thread and show the status of execution in the remote SQL statement status view. If an SQL select statement is provided, you will be optionally prompted for an SQL table or library name in order to create a table using the SQL statement given.

Remote SQL Status View

Added a new view, which processes database SQL statements in a separate thread. These requests can be queued up and are processed one at a time. This view is called the Remote SQL Statement Status View and is used by the summary programs in iDoctor as well as the *Execute in batch* option from the SQL Editor.

Property Sheet Views

Property sheets that are not modal (meaning you can go back to other panels while they are open) are now displayed inside of a view. This provides the following changes that users have been asking for:

- ▶ These property sheets are no longer always on top. You can click another view and it will move in front of the property sheet view.
- ▶ Property sheet views can be resized to any size.
- ▶ Drill-down actions from a property sheet view will now open on top of the property sheet view, like any other drill-down in the Data Viewer.

Preferences

The following topics describe changes for the Preferences setup:

- ▶ Added two new Preferences to the Graph views section of the Display tab. *Use normalize option*, when checked will cause the normalize option to be used when a graph is first opened (if it is available.) *Use variable width bar option*, when checked will cause the variable width bar option to be used when a graph is first opened (if it is available).
- ▶ Added a preference on the miscellaneous tab called *Show non-modal property pages in a view*. If you do not like the new property sheet views, you can uncheck this and go back to the way it was before.
- ▶ Added a scheduling preferences page that allows you to change the default scheduled date and time, from 1 week in the future, to whatever you prefer (in +days/hours from current date/time).

Graphing

Changes in Graphing are as follows:

- ▶ Added the capability to select time ranges through bar selection. Hold down Shift and click the first and last bar of the desired time range. To remove a selection, click anything except a bar. If a selection is made, it will be applied to the next drill-down made (if appropriate).
- ▶ Added the capability to select one or more objects in rankings charts. Use Shift to select a range of entries, and use Ctrl to select or deselect a single bar.
- ▶ Added a button on the toolbar in the Data Viewer to redisplay the current graph with variable width bars such that wider bars represent longer intervals. When displaying a graph in this manner the x-axis changes to date/time scaling.
- ▶ Added a button on the toolbar in the Data Viewer to normalize time interval-based graphs. This divides each value of the graph by the interval duration, which makes bars more consistent in height. This button is available on several Job Watcher graphs.
- ▶ Since graph tool tips can sometimes get in the way of what you are doing, added the ability to turn them on/off by pressing Ctrl+t while in a graph.
- ▶ The option to save a graph as a jpeg (Save As... menu when right-clicking on a graph) has been removed from the GUI.
- ▶ Added a button on the toolbar in the Data Viewer to show a synchronized table view beneath the graph. The table remains in sync with the graph as far as all selections and scrolling.

6.5.4 IBM iDoctor for IBM i Heap Analyzer changes

Changes to IBM iDoctor for IBM i Heap Analyzer are as follows:

- ▶ The compare function is now created by a stored procedure and the files generated are accessible from the server-side output files folder.
- ▶ Added a function that summarizes all object table dumps in a library into a single report. Right-click a Heap Analyzer library and choose the **Create Object Table Summary** menu to perform this option.
- ▶ The root finder function has been removed.

6.5.5 IBM iDoctor for IBM i Job Watcher changes

This section describes the changes to IBM iDoctor for IBM i Job Watcher.

General

Added a Tip of the day interface to the startup of Job Watcher that offers several short videos.

New Summary Program

Changes to the New Summary Program are:

- ▶ The summary program has been rewritten. Use the Summarize option under a collection to use this option. The primary files are single member SQL tables with the names QAIDRJWSUM_<collection name> and QAIDRJWGAP_<collection name>. These files are currently available under the server side output files folder and are utilized by the graphs. To display the SQL table names in the server side output files view, use the Select fields menu found by right-clicking the server side output files folder under a collection.
- ▶ Added a prompter for the summary program that allows you to pick and choose which summary functions to perform.
- ▶ Added a step to the summary program called Situational Analysis that looks for possible performance situations and displays them on the graphs using different background colors and flyovers.

Reporting

Changes to the Reporting setup are:

- ▶ All time interval graphs can now be toggled at analysis time to use one of six interval sizes.
- ▶ The graph folders, Wait graphs -> by time ranges and Wait graphs -> by interval have been combined into a single folder Wait graphs -> by time interval.
- ▶ The graph folders, CPU graphs -> by time ranges and CPU graphs -> by interval have been combined into a single folder CPU graphs -> by time interval.
- ▶ For summarized collections the following folders have been removed: IFS graphs -> by interval, JVM graphs -> by interval and Other graphs -> by interval. They have been replaced by the time interval graphs that allow configuration of the interval size.
- ▶ Several new job groupings drill-down graphs are now available for summarized collections from the collection overview time signature graph. The options include: Thread rankings, Job rankings, Job user rankings, Current user profile rankings, Generic Job rankings, Pool rankings, Priority rankings and Subsystem rankings.
- ▶ From rankings charts in summarized collections, if multiple bars are selected you are given additional drill-down options. Selected Threads (separate) produces a different graph for each thread selected. Selected Threads (combine) produces a single chart combining the values of the selected threads. Note that "Threads" could be any of the job grouping types: Pool, Job name, and others.
- ▶ Some JW drill-down options have changed such that the default next step is a level higher with sublevel menus beneath them in the same list to make them easier to use.
- ▶ Added an option from the Holder menu called "Holder chase" which produces a report of up to five levels deep for the current interval and thread, showing the holder of the holder of the holder, along with other useful information such as the wait object and wait information.
- ▶ Added several Classic JVM graphs. These graphs can be drilled down to rankings graphs showing the JVM jobs detected for an interval. From there a specific JVM job can be picked and the statistics graphed over time.

- ▶ Added a graph called “Current wait duration time signature” under the Wait graphs -> by interval folder. This graph displays only the *bad* current waits that were detected while snapshots were being taken. This has the effect of showing where these bad waits occurred even if they are relatively small.
- ▶ Added a graph for a single job/thread called “Job wait time signature with current wait contributions” to display the job wait time signature graph such that each interval that used CPU contains two components, the wait bucket times before the current wait occurred, and the wait time for the current wait. Intervals that did not use any CPU just have one component, which is the idle wait time. This graph should be viewed with the variable width bar option turned on. Using the normalize option as well will give the bars relatively fixed heights.
- ▶ The wait graphs -> by job folder has been renamed to wait graphs -> by thread because of the creation of the new job rankings graphs.
- ▶ A drill-down action from a rankings chart or a table will result in more than five reports being produced; added a warning message to make sure the user really wanted to do this.
- ▶ Added the “I/O pending page faults per second” field to the Thread rankings - > Page faults graph. Also added this field to the collection level I/O graphs by thread, page faults graph.
- ▶ Removed all normalized graphs defined in Job Watcher. Now any of the time interval graphs in Job Watcher can be normalized by you when desired by clicking **Normalize** in the toolbar of the Data Viewer.

Interval Details

The following topics describe changes in Interval Details:

- ▶ On the Interval Details pages, on the SQL tab, added a button to launch the SQL statement collected by Job Watcher into the Run SQL Scripts window.
- ▶ On the Interval Details pages, all three job priority values are now listed—Current XPF, Current LIC and Original LIC.
- ▶ On the Interval Details property pages, if looking at an interval for a job where CPU was not used, many fields shown are now retrieved from the last active CPU record (like current user profile, pool number, priority, and others). Previously these fields would be blank.

Collection Functions

Changes to the Collection Functions are:

- ▶ The Start Job Watcher Wizard has been renamed to Start Job Watcher Collection Wizard.
- ▶ In the Start Job Watcher Collection Wizard, replaced the ending options panel with the three termination options provided by the Job Watcher OS commands. As part of this change, the max data value on the startup options panel has been moved to the termination options panel. The default termination settings are 1 GB max size, or 100 intervals (whichever comes first).
- ▶ In the Start Job Watcher Collection Wizard, renamed the startup options panel to basic options. Also added the ability to automatically generate the collection name when using this interface from the current date or time. The library will also default to QJWDATA instead of blank.

Collection Search

Added key fields such as LDIO, IFS, and Transaction counters from the QAPYJWPRC file to the search function output. This affects the job/task name, subsystem, pool, and current user profile and task count search option.

Preferences

Changed the ship default preferences for Job Watcher summary graphs showing threads. Previously the default was only to show interval 1 to 10 or the last 10 intervals. Changed this to 1 to 999999 and the last 999999 intervals since most users look at the entire collection anyway.

6.5.6 IBM iDoctor for IBM i PEX-Analyzer changes

The following topic describes changes for IBM iDoctor for IBM i PEX-Analyzer:

Reporting

Changes for Reporting are:

- ▶ The PEX-Analyzer stats flat analysis for all jobs/threads is now available from PEX Stats Hier collections at IBM i5/OS V5R3 or higher. Also added the default menu for this analysis to IBM i5/OS V5R3+ Stats Hier collections. The latest PEX-Analyzer server builds are required for this fix to work.
- ▶ When working with PEX-Analyzer graphs and opening a graph that is too complex and get an out of memory error, the error will be shown inside the graph view itself rather than in a popup box.

6.6 IBM Systems Director Navigator for i

New in IBM i 6.1, IBM Systems Director Navigator for i is a Web-based console that consolidates all of the Navigator for i functions available on the Web. IBM Systems Director Navigator for i is the strategic tool for IBM i management and works in concert with IBM Systems Director products that enable heterogeneous and cross-system management. For further information about IBM Systems Director Navigator for i, refer to Chapter 16, “IBM Systems Director Navigator for i” on page 403.

In this chapter we focus on the functions in IBM Systems Director Navigator for i that relate to performance.

IBM Systems Director Navigator for i provides the following Performance interface:

- ▶ Investigate Data
 - Allows you to investigate previously collected performance data on your system.
- ▶ Collections
 - Allows you to manage the performance data of your system.

6.6.1 Collections

To access the pages to manage collected performance data, first we log on to the Integrated Solution Console in IBM Systems Director Navigator for i, as shown in Figure 6-7.

To access the logon page, use a URL similar to `https://<system-name>:2001/ibm/console/logon.jsp`, where `<system-name>` is your IBM i system name or IP address.

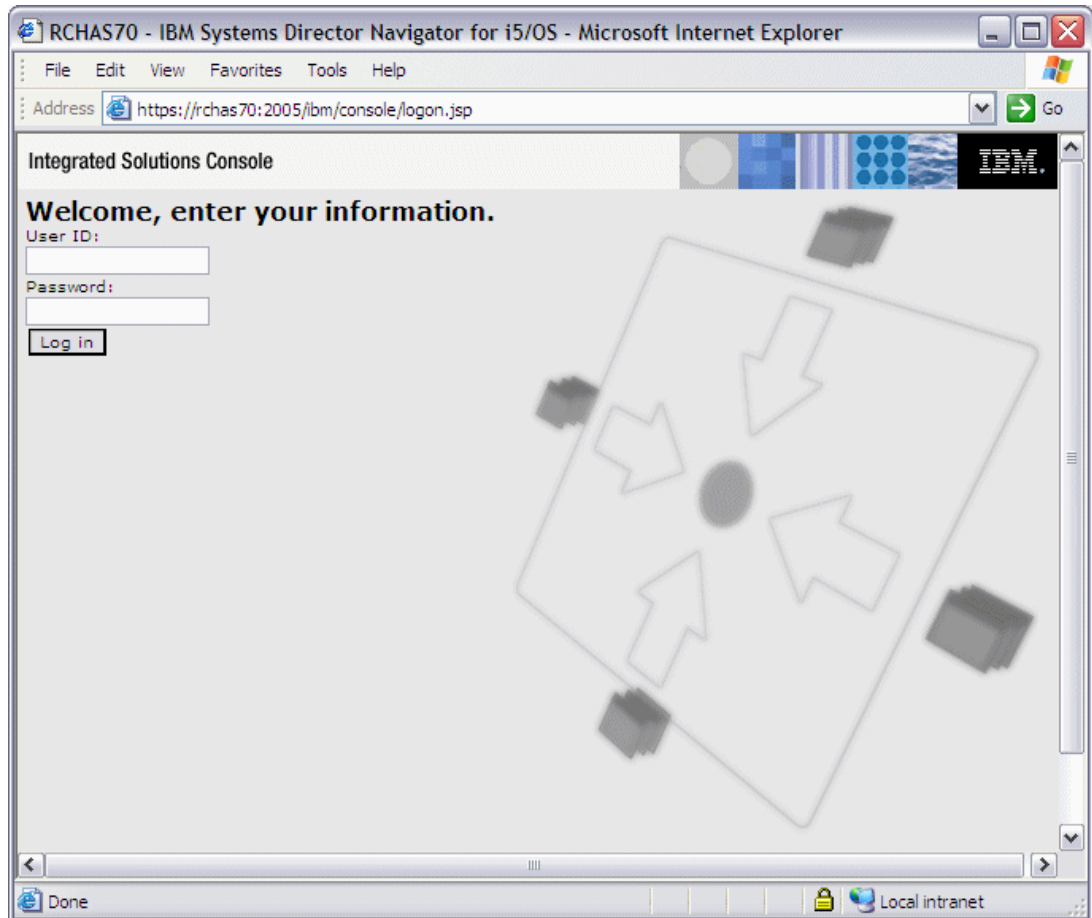


Figure 6-7 Log on to the Integrated Solution Console

After logging on, the Welcome page is displayed, shown in Figure 6-8. On the left pane, expand i5/OS Management to see the associated links.

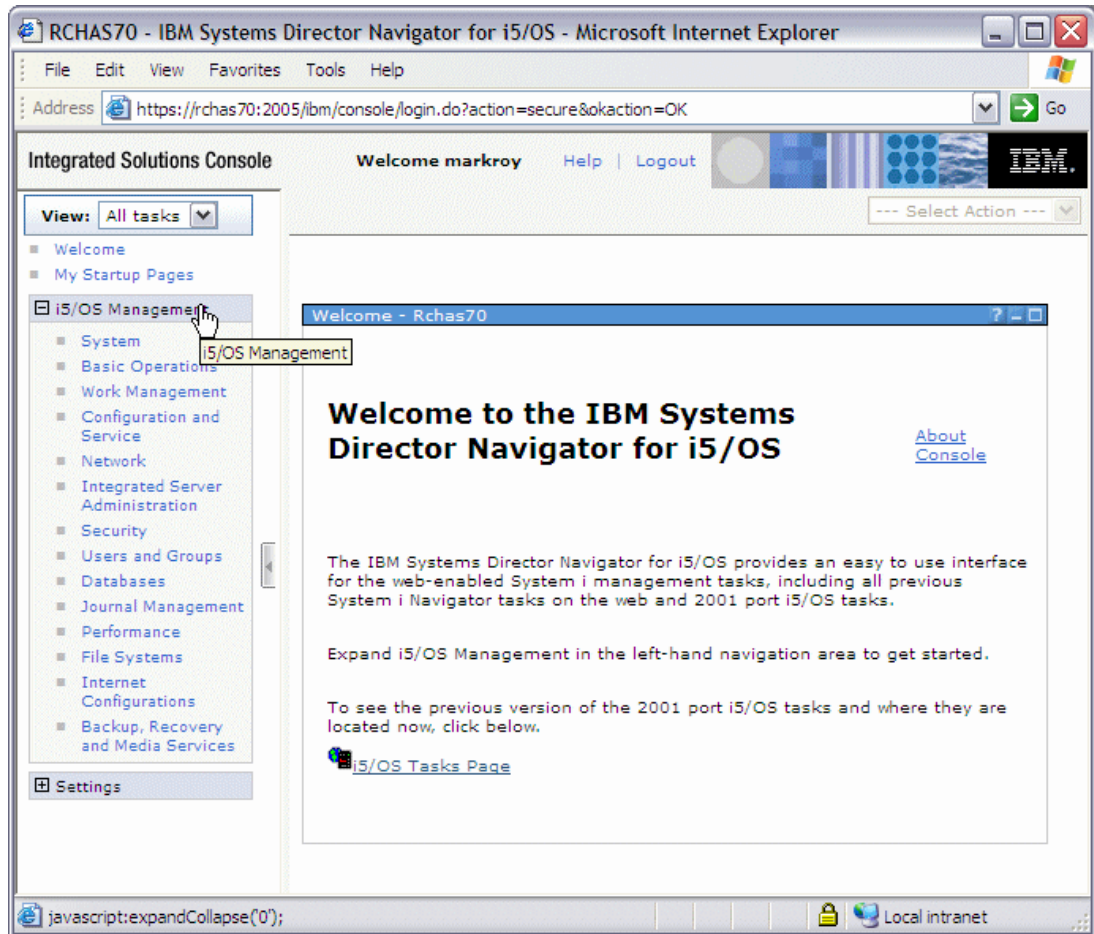


Figure 6-8 Welcome page

On the expanded list in the left pane, select **Performance** to open the Performance pane. On the newly displayed right pane, select **Collections** to display the Collections pane, as shown in Figure 6-9.

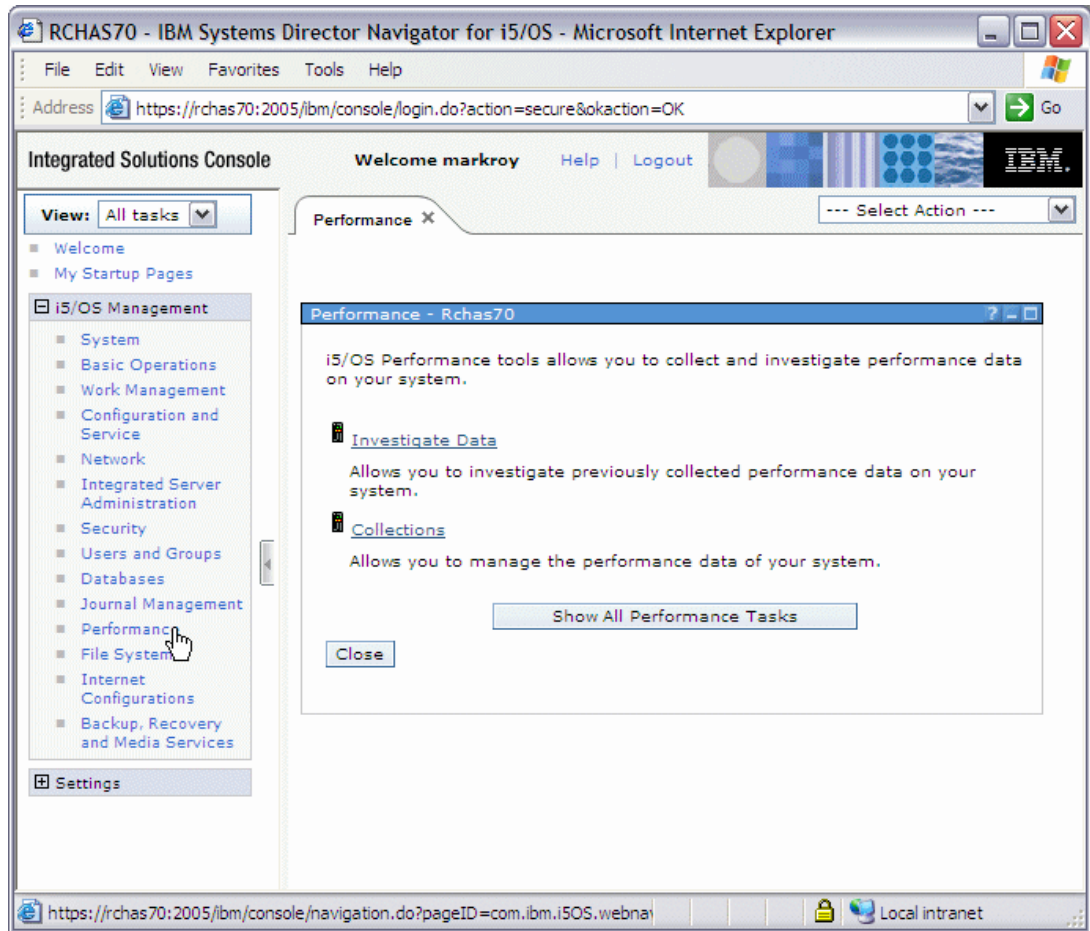


Figure 6-9 Managing collections

On the Collections pane, IBM Systems Director Navigator for i displays a list of all available collections and their attributes (status, size, and others). Near the upper left corner of the pane there are six buttons to help manage the data shown in the pane. The two left-most buttons allow selection or deselection of the entries in the list, the middle two allow filtering, and the right-most buttons allow sorting. These buttons are useful if you are managing a large list of collections. Figure 6-10 shows an example of a collections list.

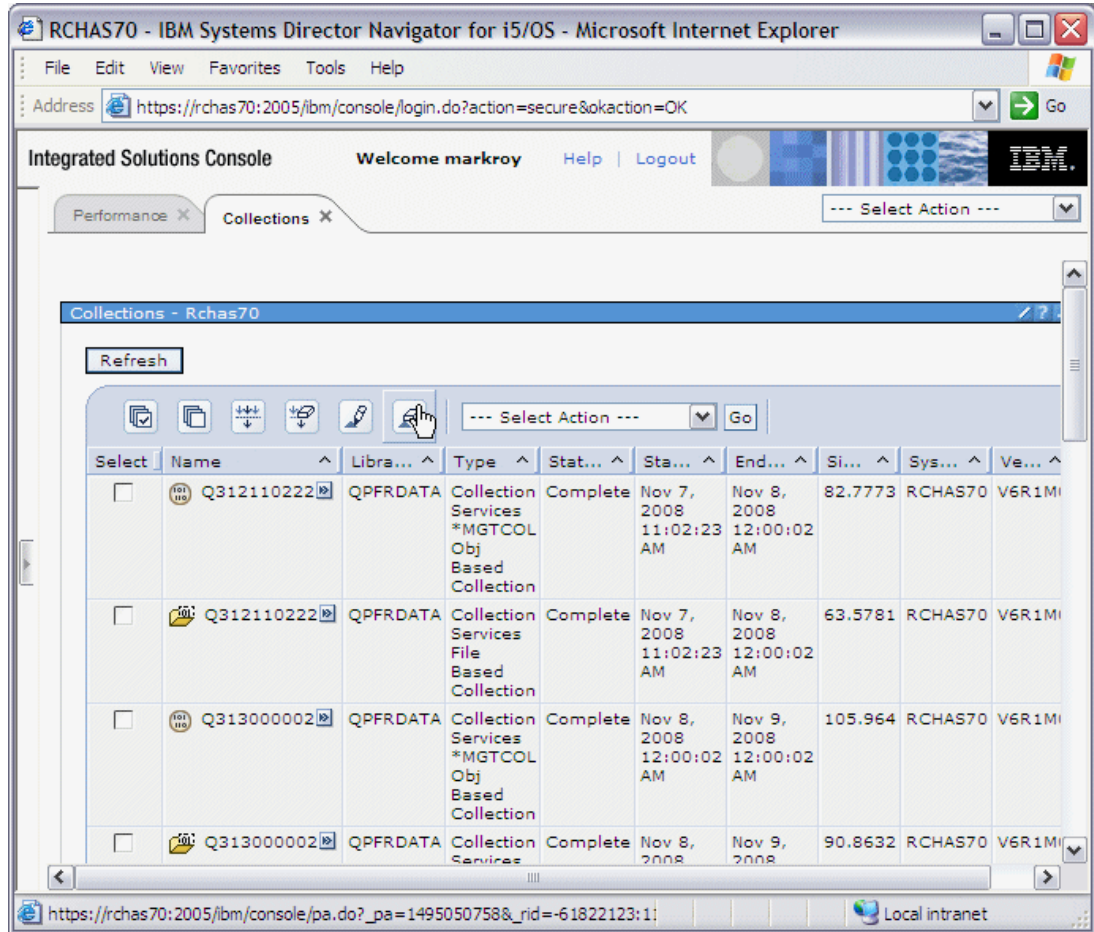


Figure 6-10 Collections list

6.6.2 Investigate Data - Collection Services

Standard IBM i 6.1 provides management and graphical display of Collection Services data through the new Investigate Data interface.

The intent of this support is to provide medium to deep level performance analysis of your partition's performance. One of the primary sets of performance metrics available here but not with the IBM Performance Tools for i5/OS (5761-PT1) are up to 32 *wait conditions*. Some of these wait conditions are actual "waits" such as record or object locks and message or data queue waits. However, others are actually "time spent doing something" such as being dispatched to a CPU (processor) or doing journaling.

The remainder of this section describes how to access the Investigate data web pages and provides examples of some of the graphs available.

On the IBM Systems Director Navigator for i Performance pane, select **Investigate Data** as shown in Figure 6-11.

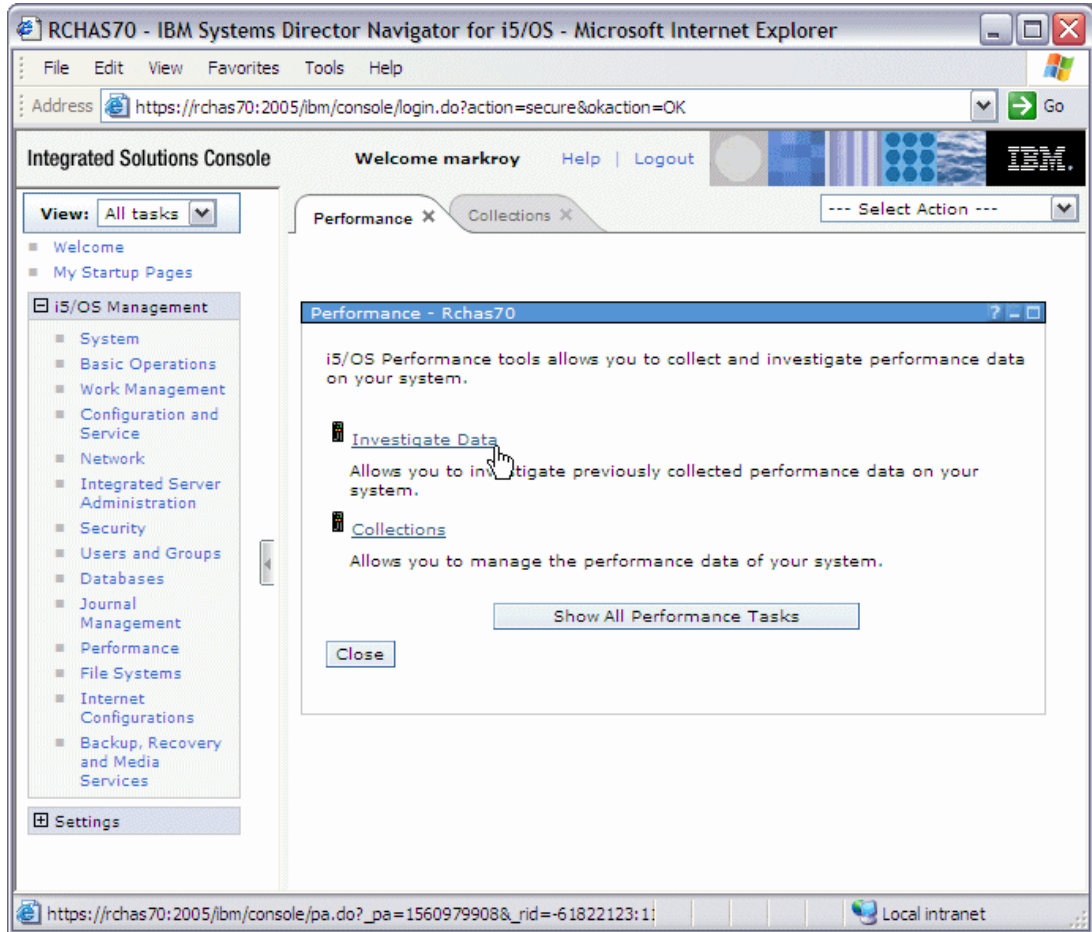


Figure 6-11 Accessing Investigate Data

The Investigate pane is displayed. From here you can select the type of collected data you want to investigate, as shown in Figure 6-12. Select **Collection Services**.

Note: Disk Watcher and Job Watcher are shown only if the Investigate Performance Data support detects that the associated IBM Performance Tools for i5/OS (5761-PT1) options (1 and 3) have been installed.

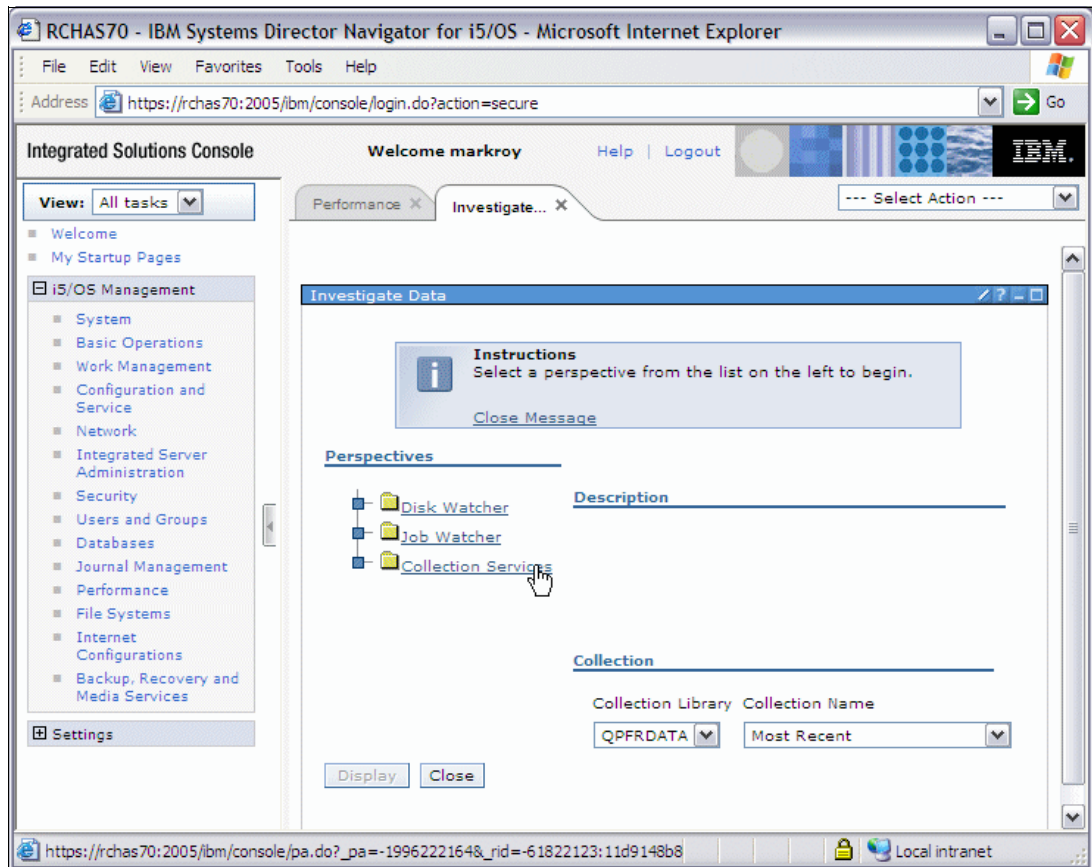


Figure 6-12 Selecting the type of data to investigate

On the Investigate pane:

1. Select the Collection Services data set by first selecting the **Collection Library** using the pull-down menu, or by typing in the name.
2. Then select the required Collection Name using the pull down menu or by typing in the name.
3. Then select the graph to display. In our example, Resource Utilization Overview.
4. Finally, select **Display**.

Figure 6-13 shows an overview for Selecting Resource Utilization.

Note: On IBM Systems Director Navigator for i windows, you can use the arrow button on the line between the left and right panes to either hide or reveal the left pane. In subsequent windows in this example we used this to hide the left pane. The pane remains hidden on following windows until it is explicitly revealed again.

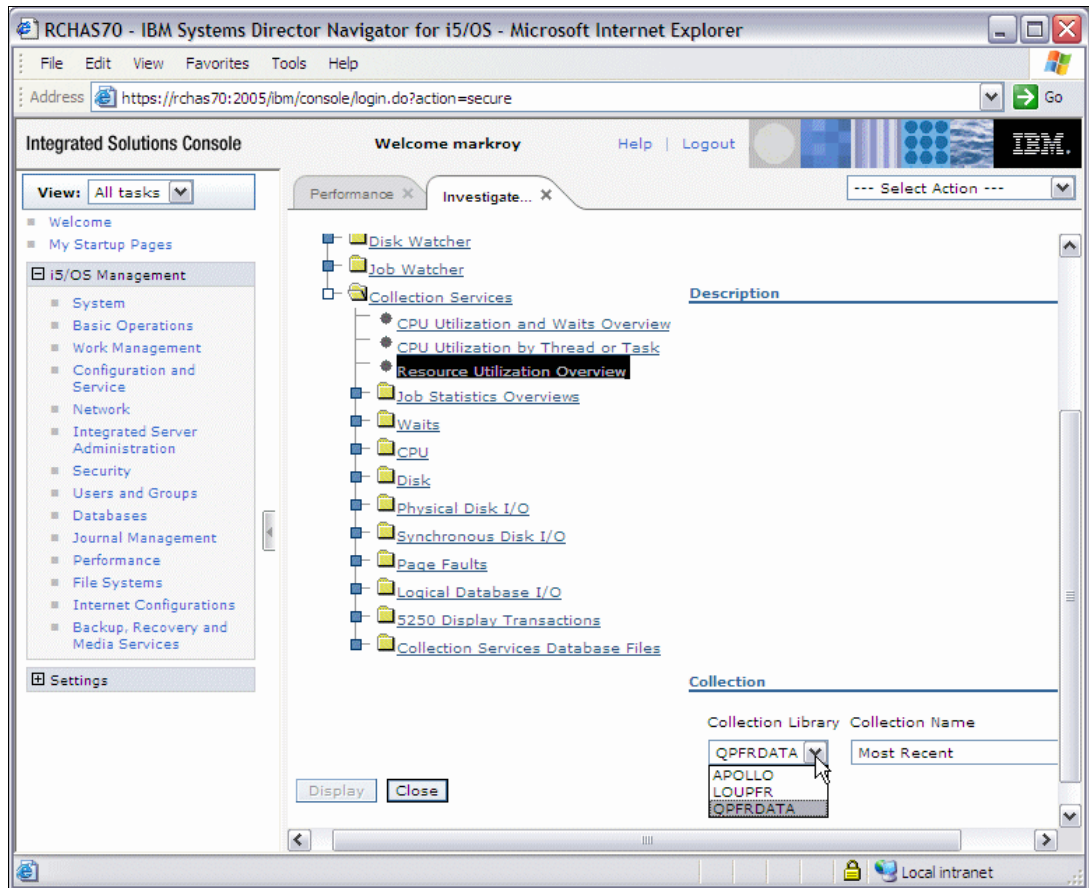


Figure 6-13 Selecting Resource Utilization Overview

The Resource Utilization Overview pane is displayed. This pane contains two graphs, Resource Utilization Percentages and Resource Utilization Rates. Figure 6-14 shows only the former. The graph also contains five buttons in the upper left corner, which can be used to manipulate the view. They are, from left to right:

- ▶ Arrow head - Select
- ▶ Hand - Pan
- ▶ Magnifying glass - Zoom
- ▶ Callout box - Show tool tips
- ▶ Crossed arrows - Full zoom out

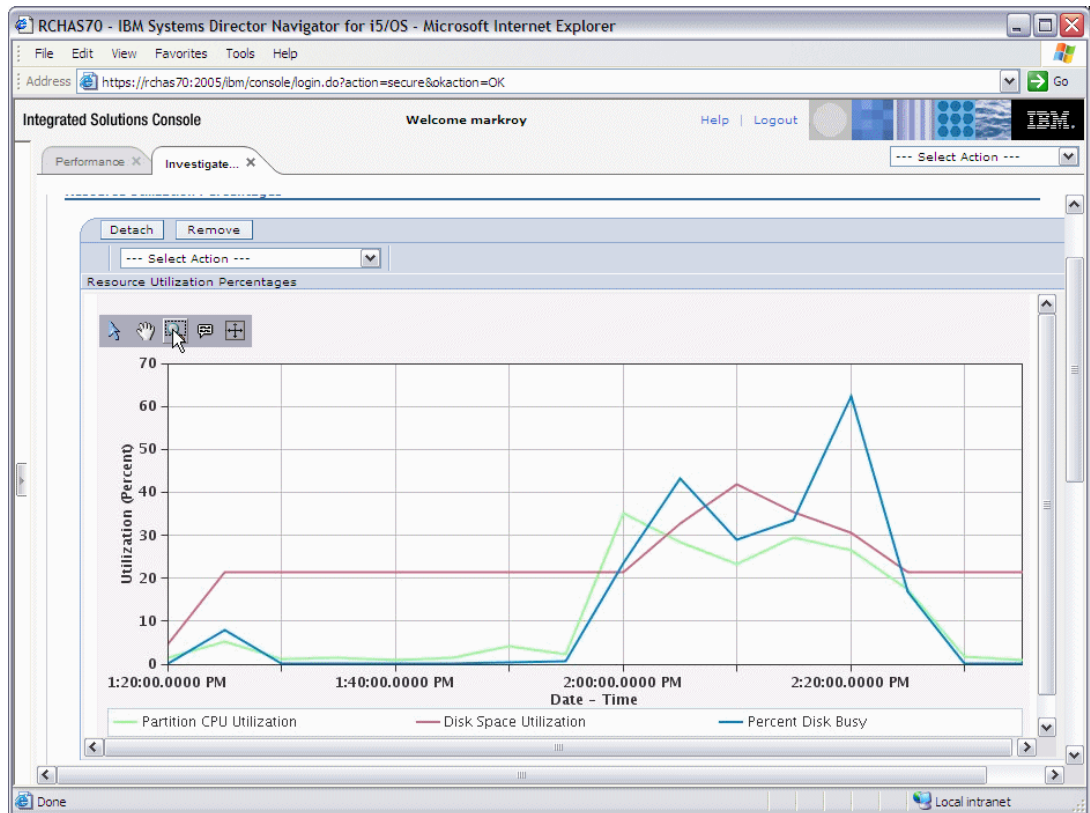


Figure 6-14 Resource utilization overview

Select **Done** to return to the Investigate Data pane and then select another perspective you are interested in, then select **Display**. In our example we selected CPU Utilization by Thread or Task to view the following graph. Figure 6-15 shows an example of CPU Utilization in a task/thread graph.

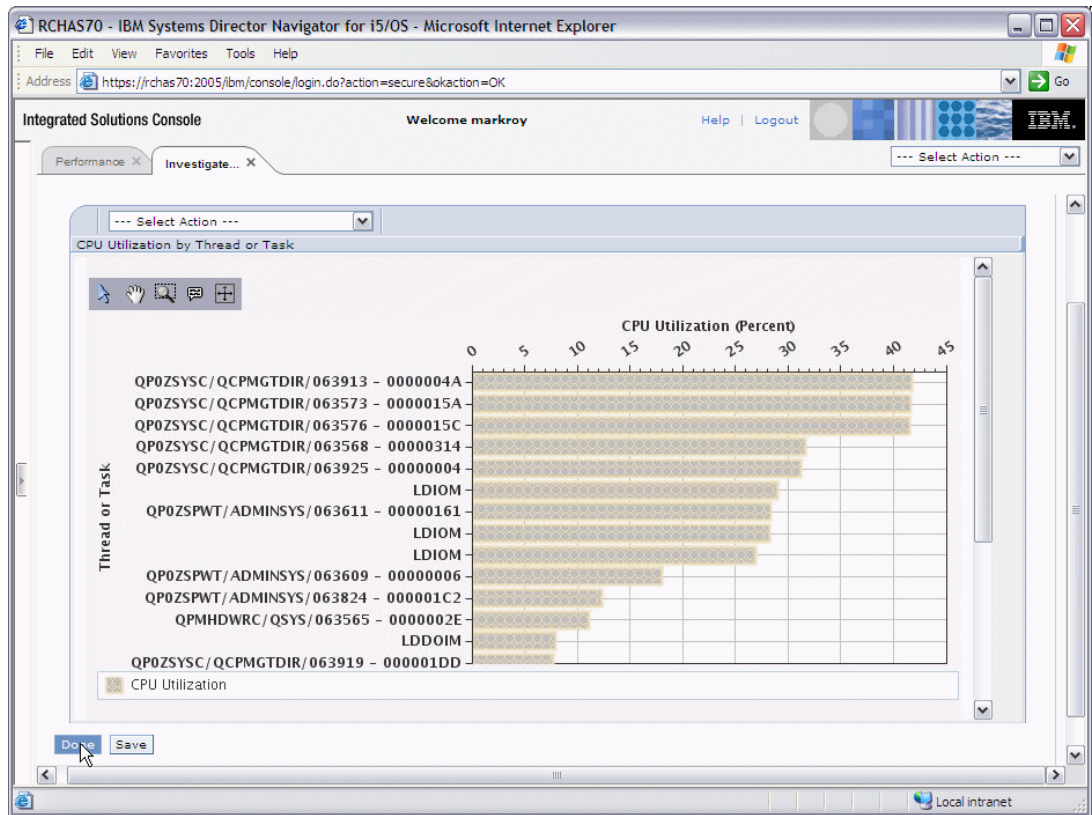


Figure 6-15 CPU Utilization by task/thread

We then returned to the Investigate Data pane and displayed the CPU Utilization and Waits Overview perspective. This can be a useful view because it not only shows the overall CPU utilization but it also shows the time jobs were waiting and how long each category of wait took. Selecting a wait for a particular interval will display the actual wait time. Figure 6-16 shows an example of a CPU Utilization and Waits Overview graph.

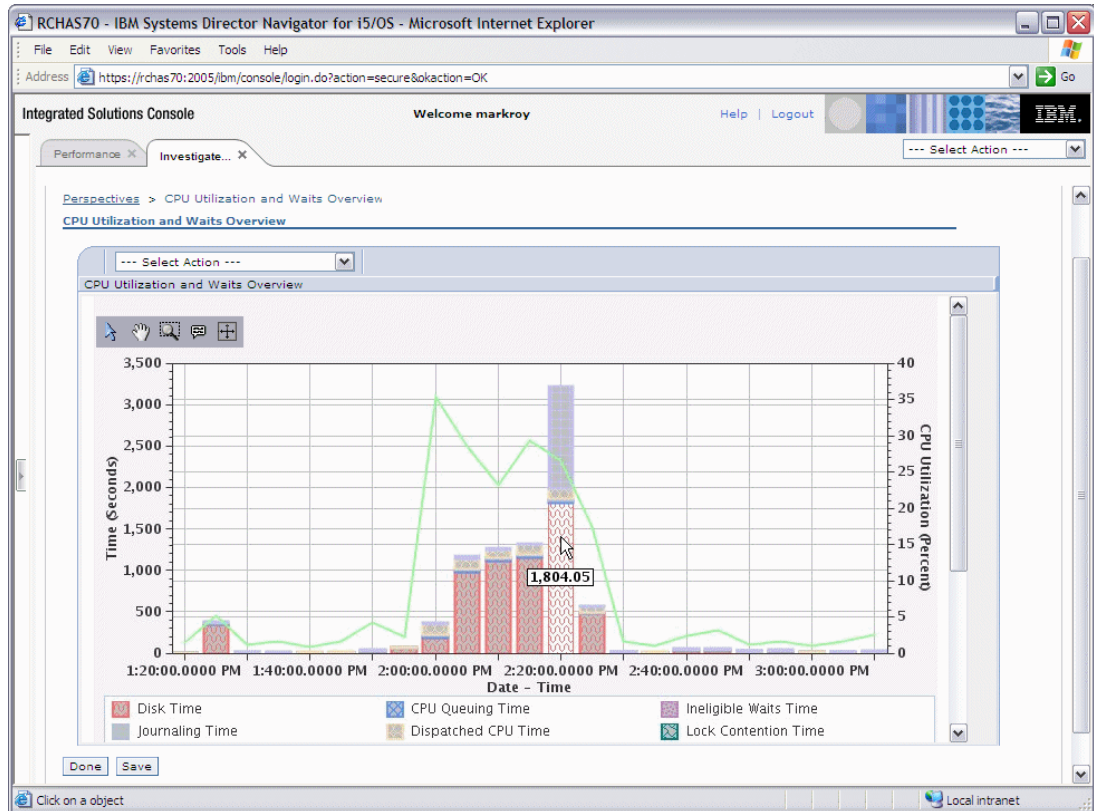


Figure 6-16 CPU Utilization and Waits Overview

There are many other perspectives that can be displayed. They are grouped into:

- ▶ Utilization Overview (as we saw in our examples)
- ▶ Job Statistics Overviews
- ▶ Waits
- ▶ CPU
- ▶ Disk
- ▶ Physical Disk I/O
- ▶ Synchronous Disk I/O
- ▶ Page Faults
- ▶ Logical Database I/O
- ▶ 5250 Display Transactions
- ▶ Collection Services Database Files

Table 6-2 on page 155 provides a full list of perspectives.

Table 6-2 Perspectives

Functional area	Perspective
Job Statistics Overviews	
	Job Statistics by Job or Task
	Job Statistics by Generic Job or Task
	Job Statistics by Job User Profile
	Job Statistics by Job Current User Profile
	Job Statistics by Subsystem
	Job Statistics by Server Type
	Job Statistics by Job Pool
Waits	
	Waits Overview
	Seizes and Locks Waits Overview
	Contention Waits Overview
	Disk Waits Overview
	Journal Waits Overview
	Classic JVM Waits Overview
	All Waits by Thread or Task Sorted by Dispatched CPU
	Waits by Job or Task Sorted by Dispatched CPU
	Waits by Generic Job or Task Sorted by Dispatched CPU
	Waits by Job User Profile Sorted by Dispatched CPU
	Waits by Job Current User Profile Sorted by Dispatched CPU
	Waits by Pool Sorted by Dispatched CPU
	Waits by Subsystem Sorted by Dispatched CPU
	Waits by Server Type Sorted by Dispatched CPU
	Waits by Job Priority Sorted by Dispatched CPU
CPU	
	CPU Utilization Overview
	Interactive Capacity CPU Utilization
	CPU Utilization by Job or Task
	CPU Utilization by Thread or Task
	CPU Utilization by Generic Job or Task
	CPU Utilization by Job User Profile
	CPU Utilization by Job Current User Profile
	CPU Utilization by Subsystem

Functional area	Perspective
	CPU Utilization by Server Type
	CPU Utilization by Pool
Disk	
	Disk Overview for System Disk Pool
	Disk Overview by IOP Name
	Disk Overview by Adapter Name
	Disk Overview by Disk Pool
	Disk Overview by Disk Unit
	Disk Overview by Disk Path
	Disk Details by IOP Name
	Disk Details by Adapter Name
	Disk Details by Disk Pool
	Disk Details by Disk Unit
	Disk Details by Disk Path
	Disk Overview for IOPs
	Disk Overview for Adapters
	Disk Overview for Disk Pools
	Disk Overview for Disk Units
	Disk Overview for Disk Paths
	Disk Details for IOPs
	Disk Details for Adapters
	Disk Details for Disk Pools
	Disk Details for Disk Units
	Disk Details for Disk Paths
Physical Disk I/O	
	Basic Views
	Detailed Views
	Advanced Views
Synchronous Disk I/O	
	Synchronous Disk I/O Overview
	Synchronous Disk I/O by Job or Task
	Synchronous Disk I/O by Thread or Task
	Synchronous Disk I/O by Generic Job or Task
	Synchronous Disk I/O by Job User Profile

Functional area	Perspective
	Synchronous Disk I/O by Job Current User Profile
	Synchronous Disk I/O by Pool
	Synchronous Disk I/O by Subsystem
	Synchronous Disk I/O by Server Type
	Synchronous Disk I/O by Job Priority
Page Faults	
	Page Faults Overview
	Page Faults by Job or Task
	Page Faults by Thread or Task
	Page Faults by Generic Job or Task
	Page Faults by Job User Profile
	Page Faults by Job Current User Profile
	Page Faults by Subsystem
	Page Faults by Server Type
	Page Faults by Pool
Logical Database I/O	
	Logical Database I/O Overview
	Logical Database I/O by Job
	Logical Database I/O by Generic Job
	Logical Database I/O by Job User Profile
	Logical Database I/O by Job Current User Profile
	Logical Database I/O by Subsystem
	Logical Database I/O by Server Type
5250 Display Transactions	
	5250 Display Transactions Overview
	5250 Display Transactions by Job
	5250 Display Transactions by Generic Job
	5250 Display Transactions by Job User Profile
	5250 Display Transactions by Job Current User Profile
	5250 Display Transactions by Subsystem
Collection Services Database Files	
	QAPMARMTRT
	QAPMDISK
	QAPMDOMINO

Functional area	Perspective
	QAPMDPS
	QAPMETH
	QAPMHTTPB
	QAPMHTTPD
	QAPMIOPD
	QAPMISUM
	QAPMJOBMI
	QAPMJOBOS
	QAPMJOBWTD
	QAPMJSUM
	QAPMLPARH
	QAPMMOIP
	QAPMPOOLB
	QAPMPOOLT
	QAPMPPP
	QAPMSYSCPU
	QAPMSYSTEM
	QAPMTCP
	QAPMTCPIFC
	QAPMUSRTNS
	QAPMWASAPP
	QAPMWASCFG
	QAPMWASEJB
	QAPMWASPRSC
	QAPMWASSVR

Note: In IBM i 5.4 the default number of Wait Buckets was 16. In 6.1 the number has been increased to 32. If performance data was collected on a system prior to 6.1, and converted to run on 6.1, the Wait Bucket (Analysis) database files will not convert. In this case, any options that require the Wait data will not be available for selection.

Disk Watcher

Disk Watcher provides for the collection of disk performance data to diagnose disk-related performance problems.

IBM i Disk Watcher allows you to obtain data concerning I/O operations to disk units, along with frequently needed run-time data to determine which objects, files, processes, threads, and tasks are being accessed. This tool surfaces data beyond what is provided by such tools as WRKDSKSTS, WRKSYSSTS, and WRKSYSACT. Disk Watcher provides a mechanism to

use short and longer duration traces to collect disk I/O data along with the associated task and object name.

Some potential uses of this tool are:

- ▶ Evaluating the performance of I/O operations on multi-path disk units.
- ▶ Evaluating the performance of I/O queuing.
- ▶ Determining how performance may be improved by respreading data across units.
- ▶ Determining the optimal placement of devices, IOAs, or buses.

Disk Watcher can be configured and managed through the IBM Systems Director Navigator for i Performance interface or CL commands.

The steps to start and analyze a Disk Watcher collection include:

- ▶ Add a Disk Watcher definition (this step is optional; you can use a pre-existing IBM or user-defined definition).
- ▶ Start Disk Watcher.
- ▶ End Disk Watcher.
- ▶ Use IBM Systems Director Investigator for i to view the Disk Watcher data.

Note: A Disk Watcher definition identifies the performance data that is to be collected during a Disk Watcher collection. When starting a Disk Watcher collection, you nominate which definition to use.

In the following examples, we use IBM Systems Director Navigator for i to demonstrate how to start a Disk Watcher collection and then provide some sample graphs.

Once you are logged on IBM Systems Director Navigator for i, on the left pane, expand i5/OS Management and select **Performance** to open the Performance pane.

As shown in Figure 6-17, we select **Show All Performance Tasks** from the Performance pane.

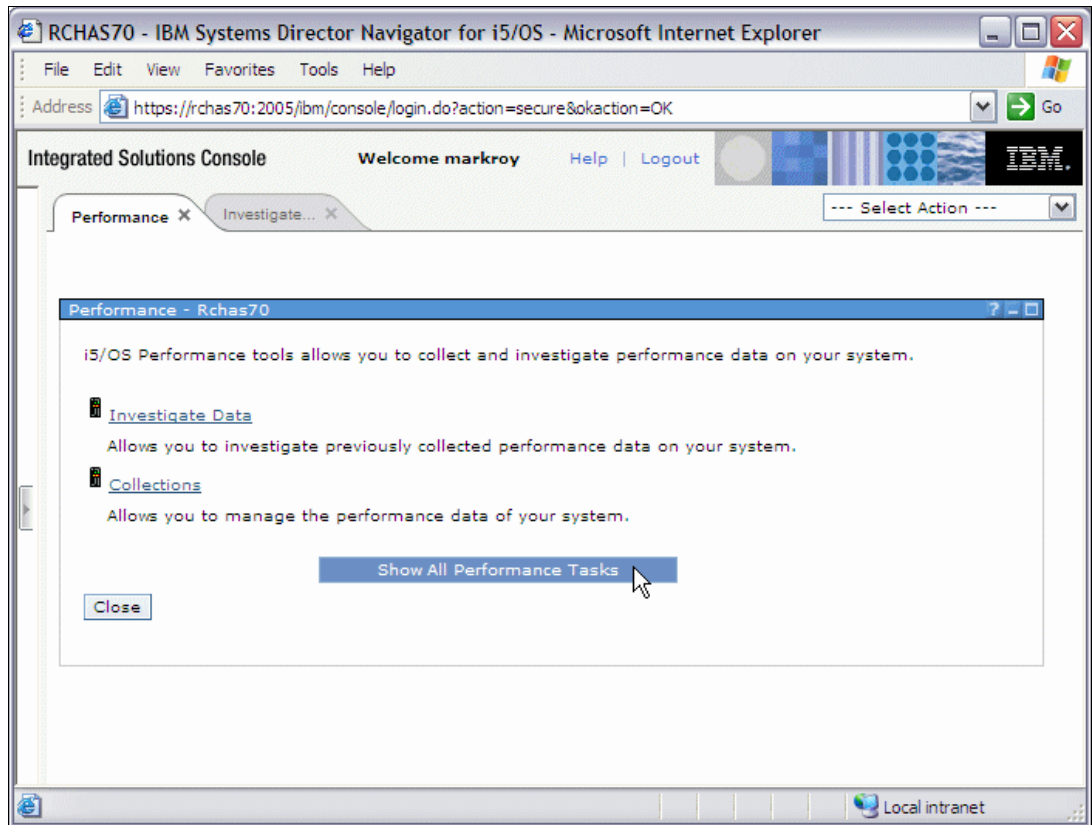


Figure 6-17 Displaying all available performance tasks

There we select the **Start Disk Watcher** option from the tasks menu as shown in Figure 6-18. The panel refreshes and shows that the line has been selected.

The first pane to be displayed is the Welcome pane, where we select **Next** to continue.

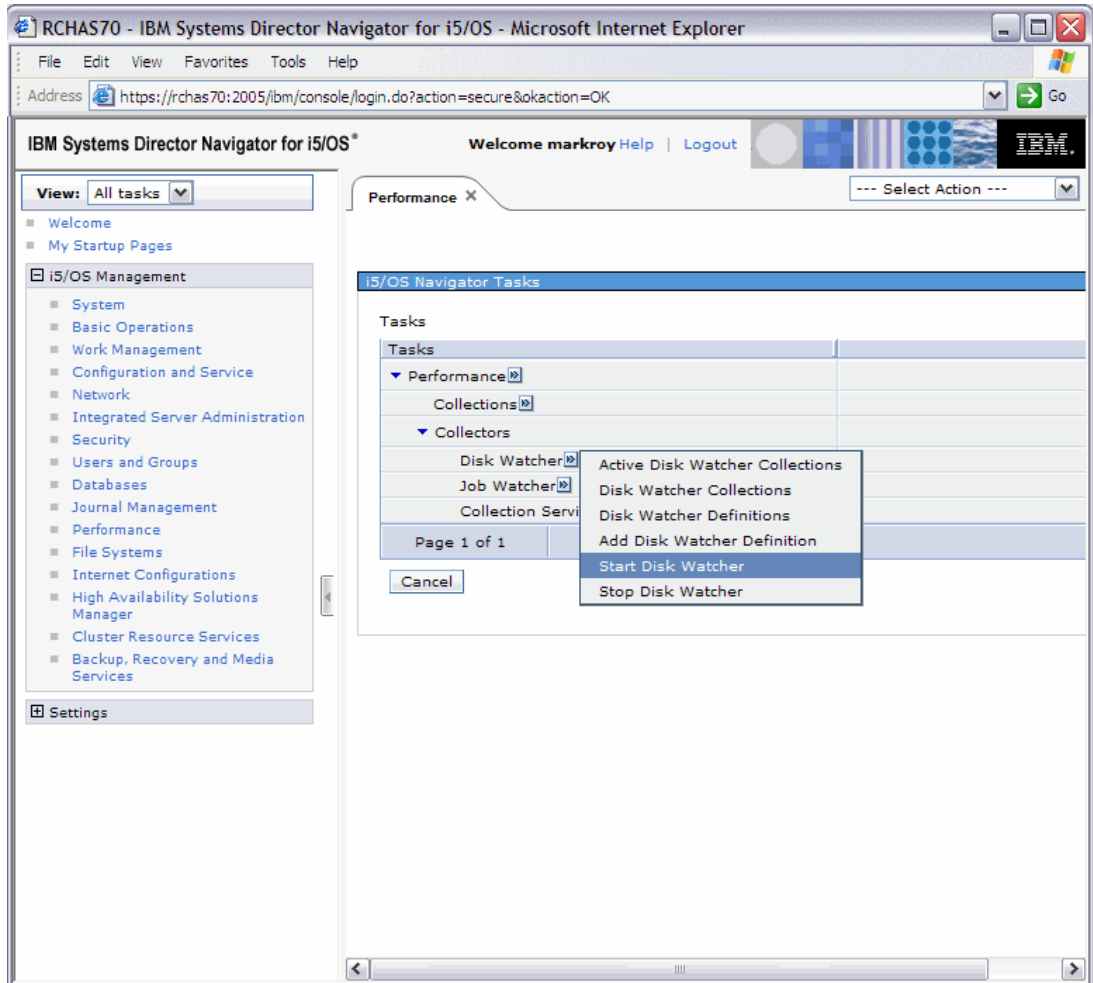


Figure 6-18 Disk watcher collection tasks

We are then presented with the Definition Selection pane where you can select from either IBM-supplied definitions or user-defined definitions. In our example we select the default **Qstats**, as shown in Figure 6-19, and then select **Next** to continue.

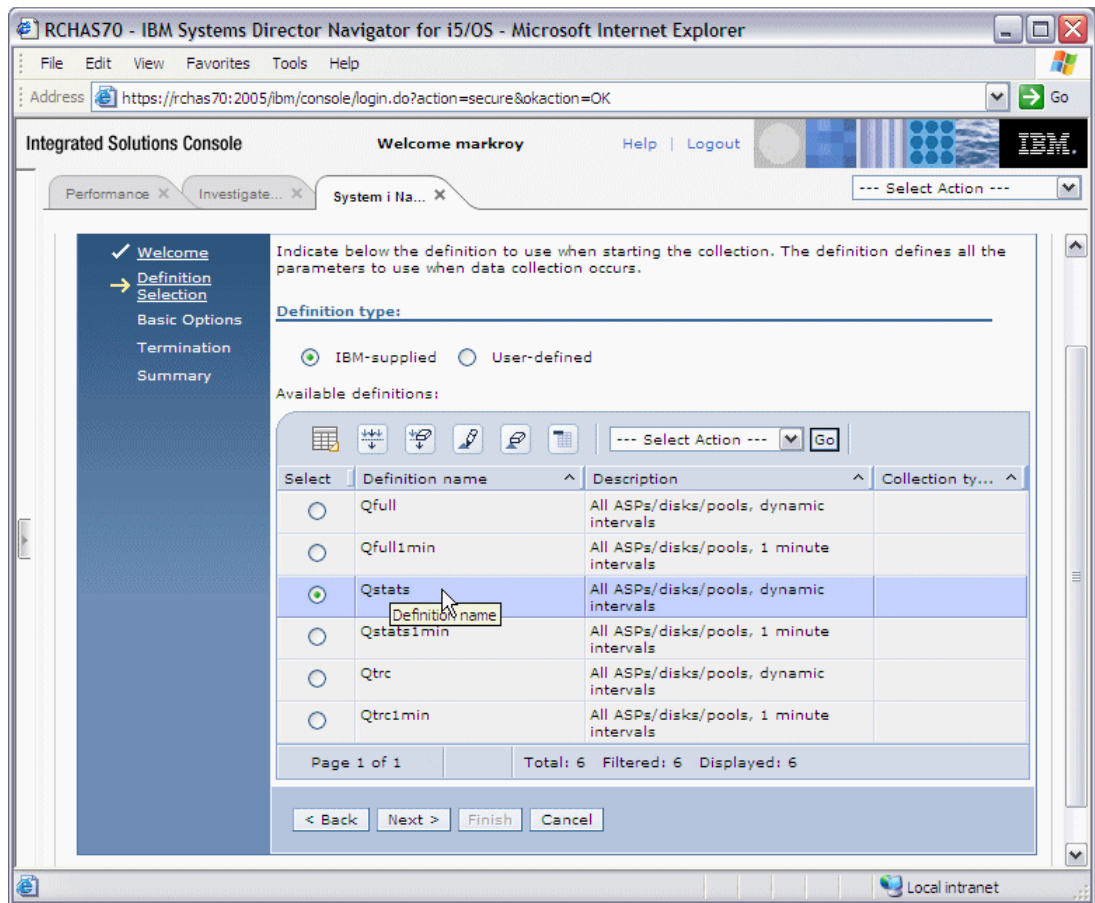


Figure 6-19 Definition Selection

On the Basic Options pane shown in Figure 6-20, we enter a name and a description for the collection, then select **Next** to continue.

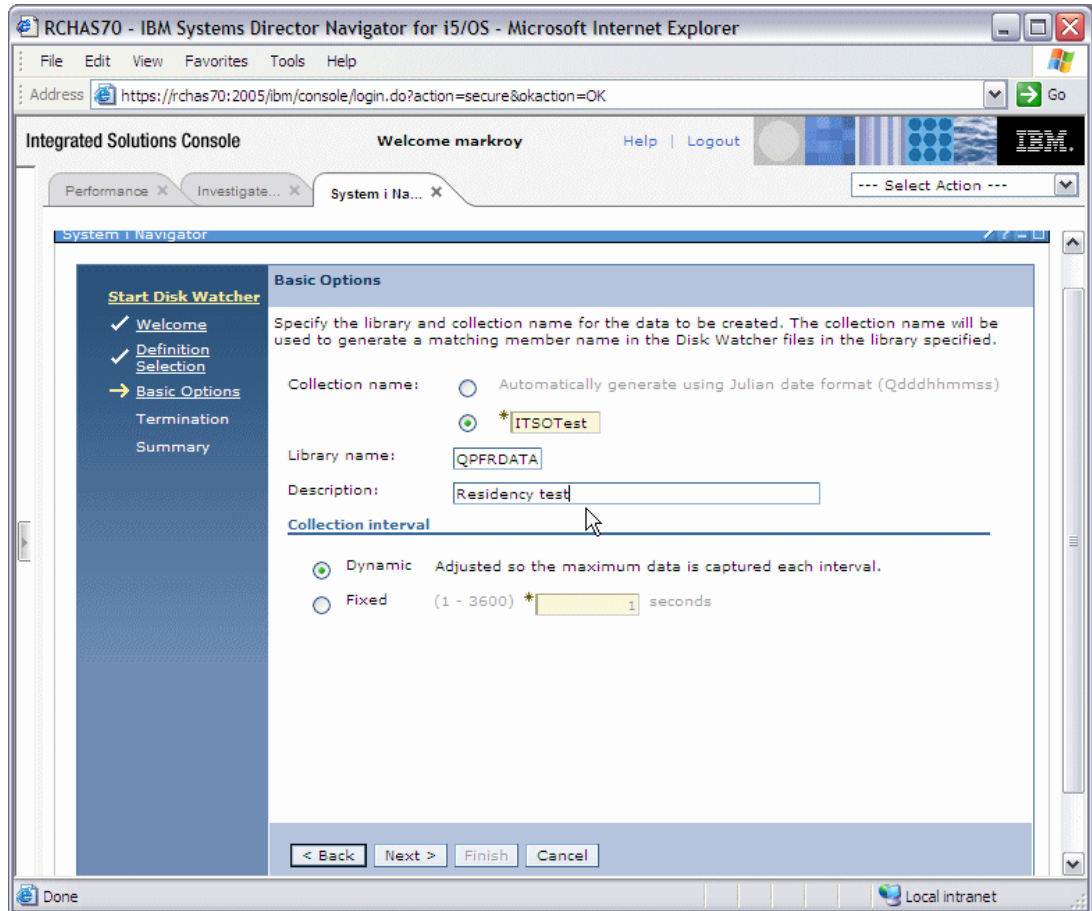


Figure 6-20 Basic Options

On the Termination pane, we specify when the collection should end. We can specify one or more conditions. When the first one of those conditions is satisfied, the collection will end. The conditions are based on:

- ▶ Disk space used by the collection.
- ▶ The number of intervals to be collected.
- ▶ The length of time the collection should run.

For our example we leave the defaults and select **Next** to continue, as shown in Figure 6-21.

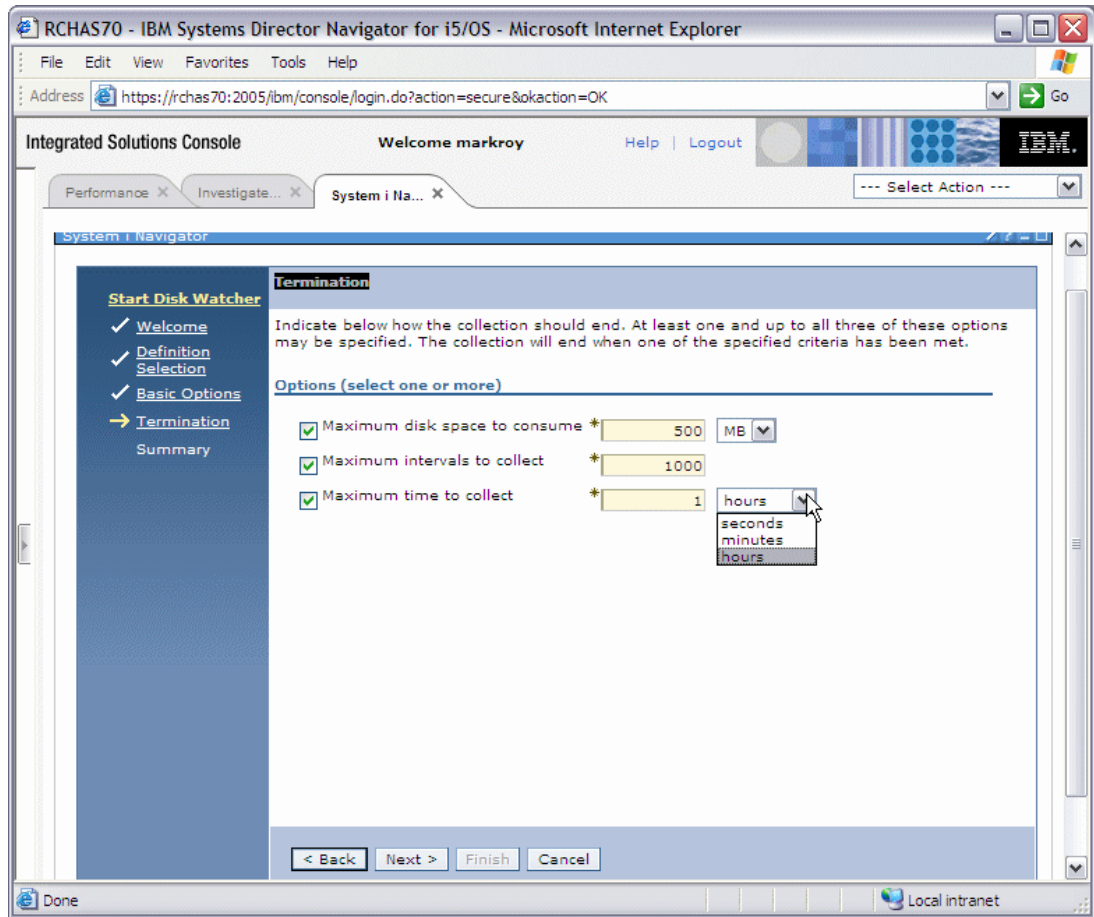


Figure 6-21 Termination

The summary pane displays information about the collection we are about to start, as shown in Figure 6-22. Select **Finish** to submit the request to start the Disk Watcher collection.

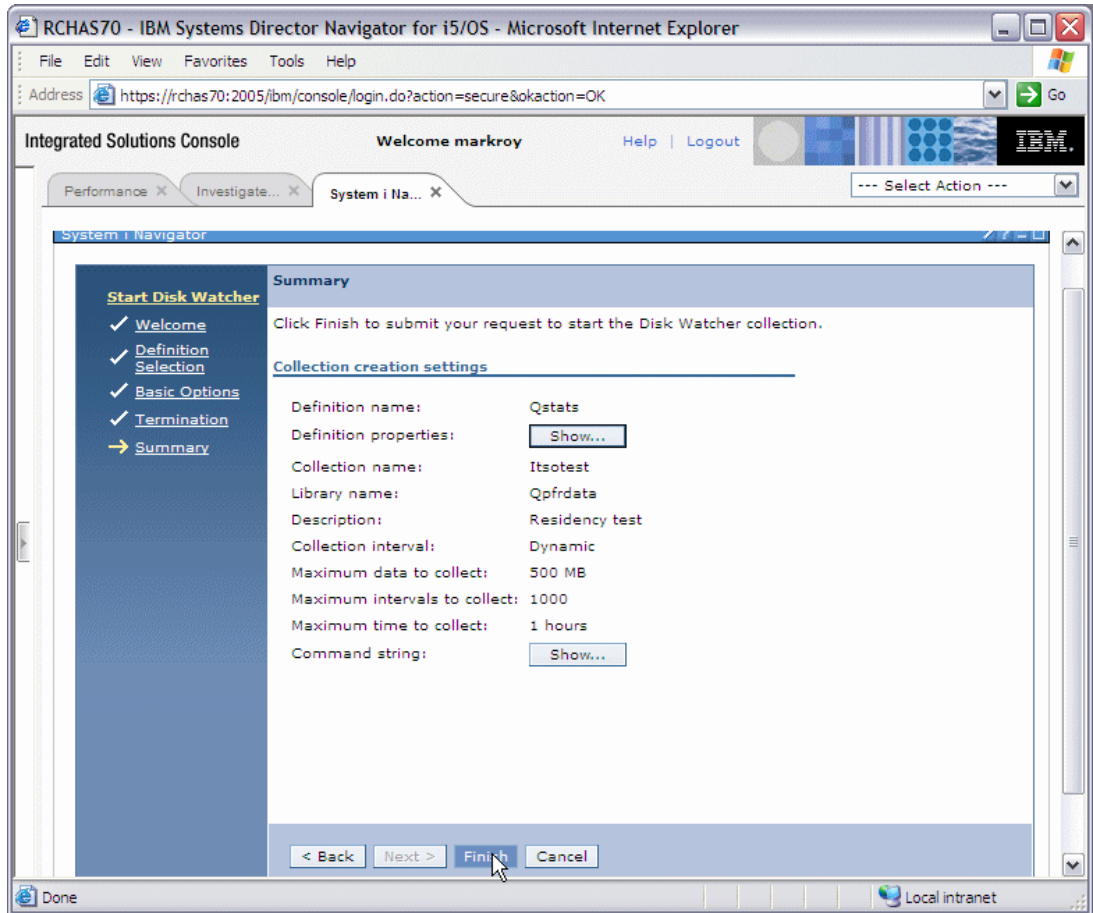


Figure 6-22 Summary

Next we want to confirm that the Disk Watch collection is running. We select **Active Disk Watcher Collections** on the tasks menu as shown in Figure 6-23.

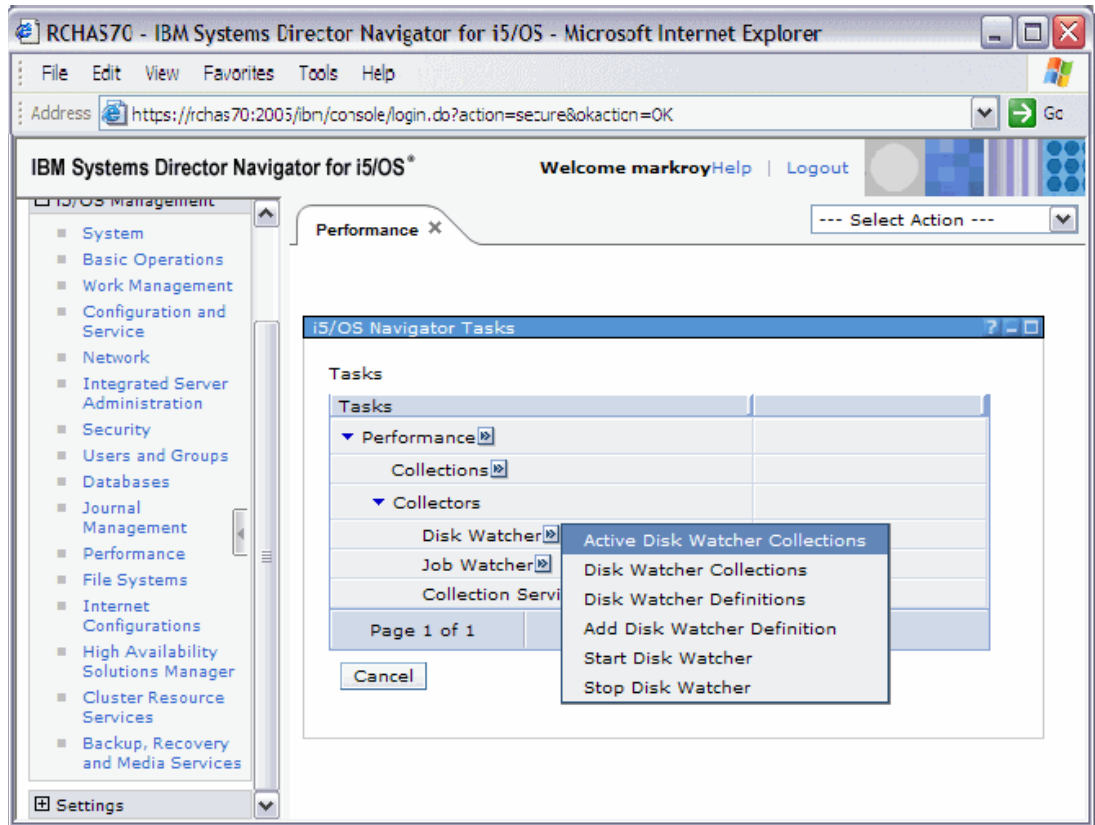


Figure 6-23 Verifying the collection status

The Active Disk Watcher Collections pane is displayed and we can see our recently started collection. Click **Close** to return to the list of perspectives, as shown in Figure 6-24.

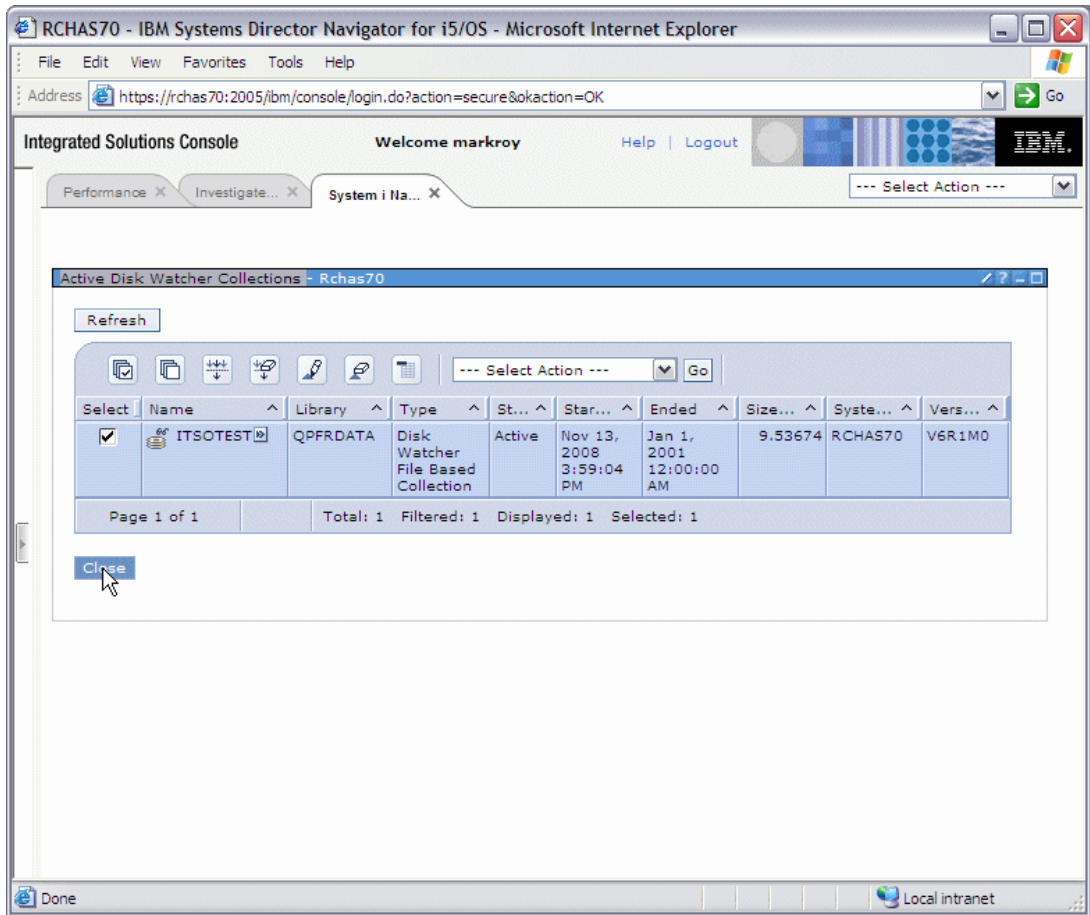


Figure 6-24 Active Disk Watcher Collections

Once the collection has terminated, return to the Investigate Data window, where we will investigate some of the collected data.

On the Investigate page, we expand the Disk Watcher tree to see all available perspectives, as shown in Figure 6-25. We then select **Disk Statistical Overview**.

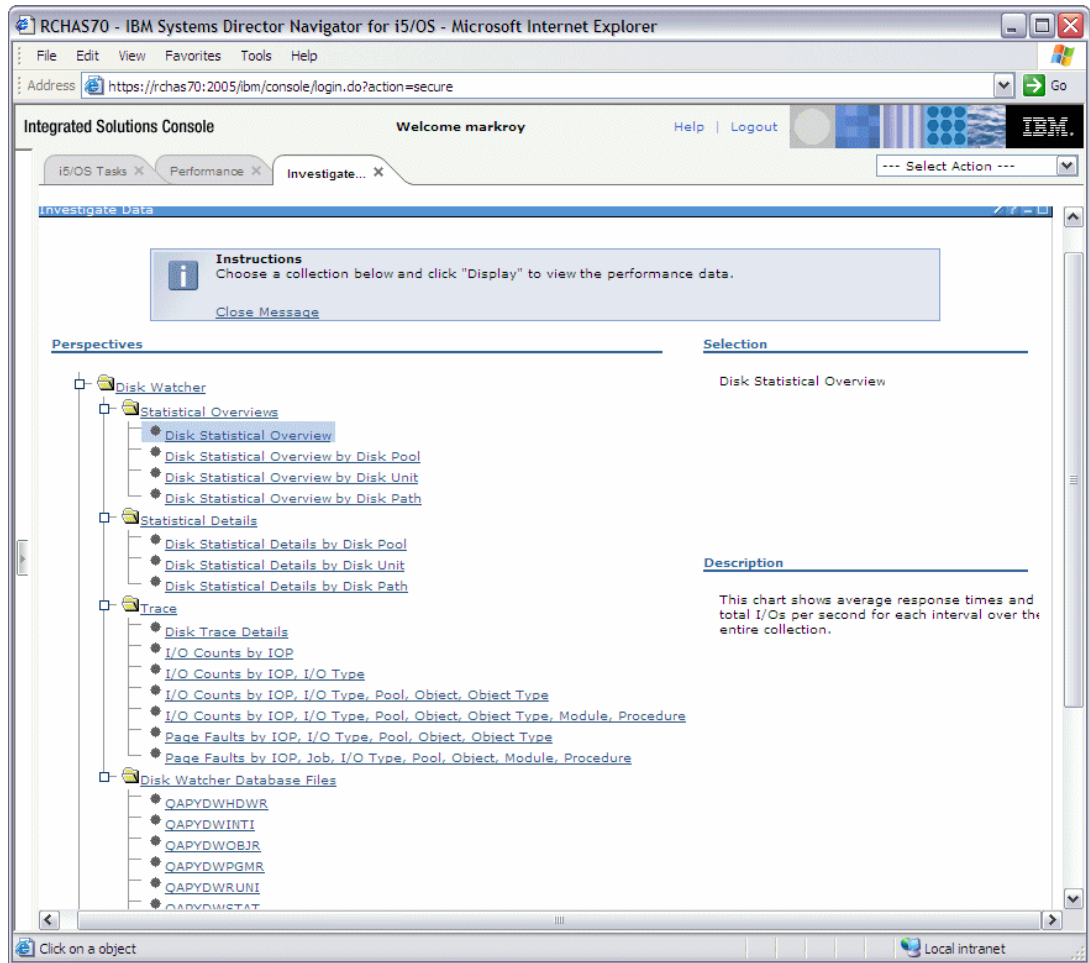


Figure 6-25 Disk Watcher perspectives

On the Disk Statistical Overview graph we can see the number of reads, writes, and other I/O. Using the drop-down list, we can select other Statistical Overview and also Detail graphs. We select **Disk Statistical Details by Disk Unit** to see whether the I/Os were spread evenly amongst the disks, as shown in Figure 6-26 on page 169.

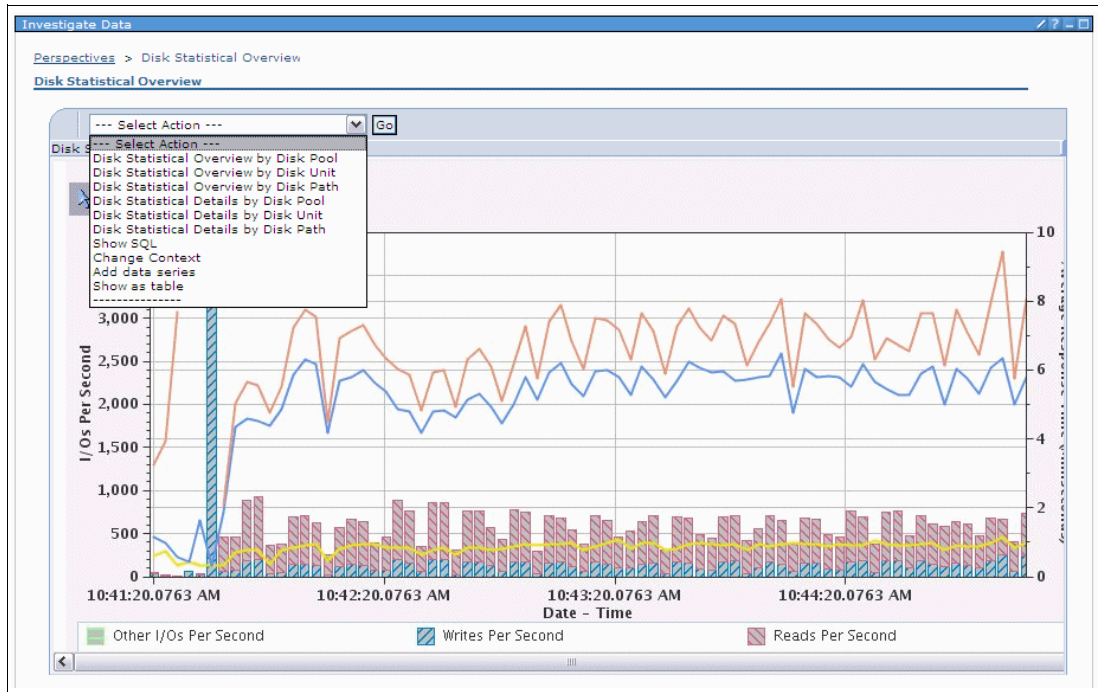


Figure 6-26 Overview chart

In our example, shown in Figure 6-27, we can see there is something unusual about disk unit 4 that might warrant further investigation using other detailed Disk Watcher perspectives.

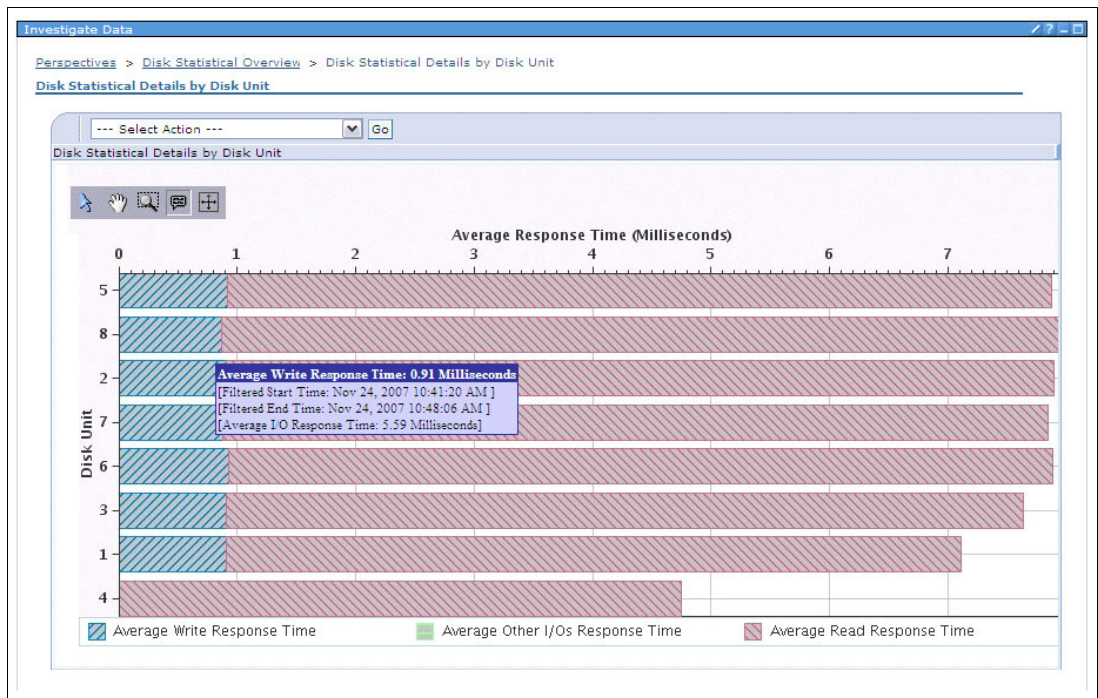


Figure 6-27 View by disk unit

Job Watcher

IBM i Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM

statistics, wait statistics, and more, which are used to diagnose job-related performance problems.

Job Watcher is similar in sampling function to the CL commands WRKACTJOB and WRKSYSACT in that each refresh computes delta information for the ending interval. The data collected from the jobs, threads, or tasks being watched is done so in a nonintrusive manner.

Job Watcher can be configured and managed through the IBM Systems Director Navigator for i Performance interface or CL commands.

The steps to start and analyze a Job Watcher collection are:

- ▶ Add a Job Watcher definition (this step is optional; you can use a pre-existing IBM or user-defined definition).
- ▶ Start Job Watcher.
- ▶ End Job Watcher.
- ▶ Use IBM Systems Director Investigator for i to view the Job Watcher data.

Note: A Job Watcher definition identifies the performance data that is to be collected during a Job Watcher collection. When starting a Job Watcher collection, you nominate which definition to use.

In the following examples, we use IBM Systems Director Navigator for i to demonstrate how to start a Job Watcher collection and then provide some sample graphs.

Note: You need to do some preliminary performance data investigation at the Collection Services level before using the Job Watcher functions. Otherwise you may collect significantly more performance data than you actually need and thus take up unnecessary disk storage.

First, we select **Show All Performance Tasks** from the Performance pane to get the Job Watcher tasks. We then select **Add Job Watcher Definition** to start the wizard, as shown in Figure 6-28.

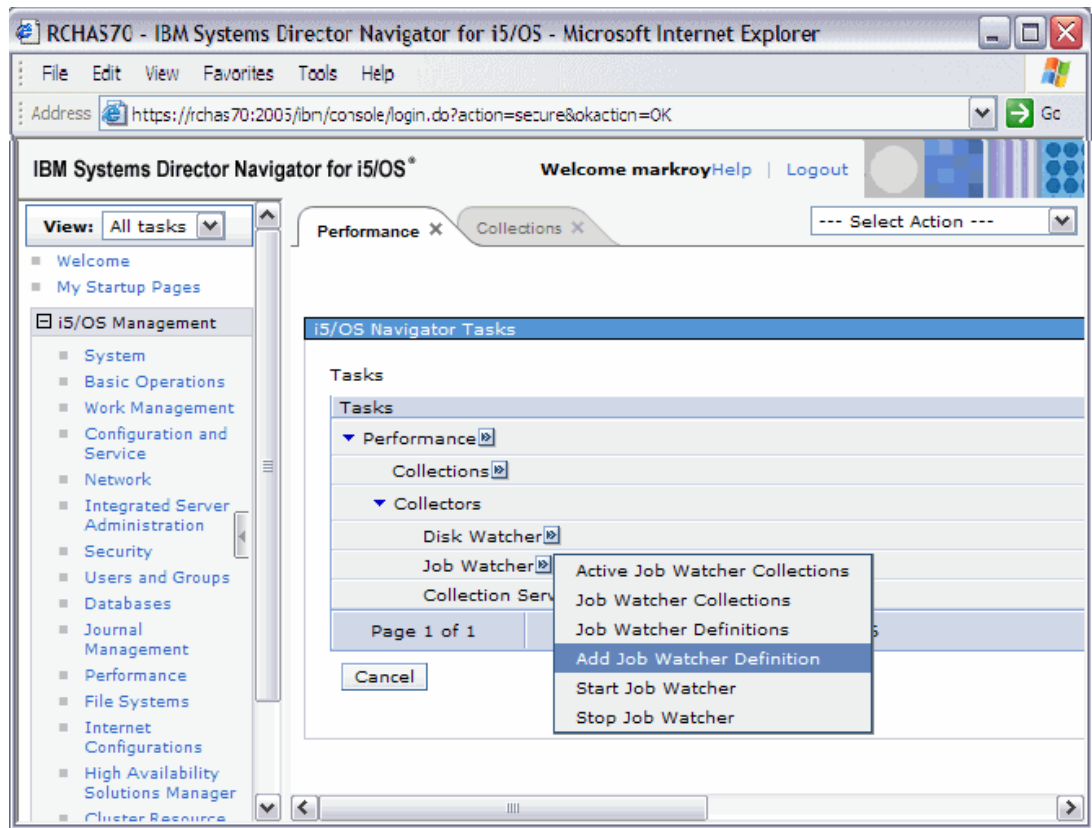


Figure 6-28 Setting up Job Watcher

On the Basic Options pane, as shown in Figure 6-29, we enter a name for the definition and a description. On this page you can change some of the collection options and time intervals if desired. In our example we leave the defaults. Select **Next** to continue.

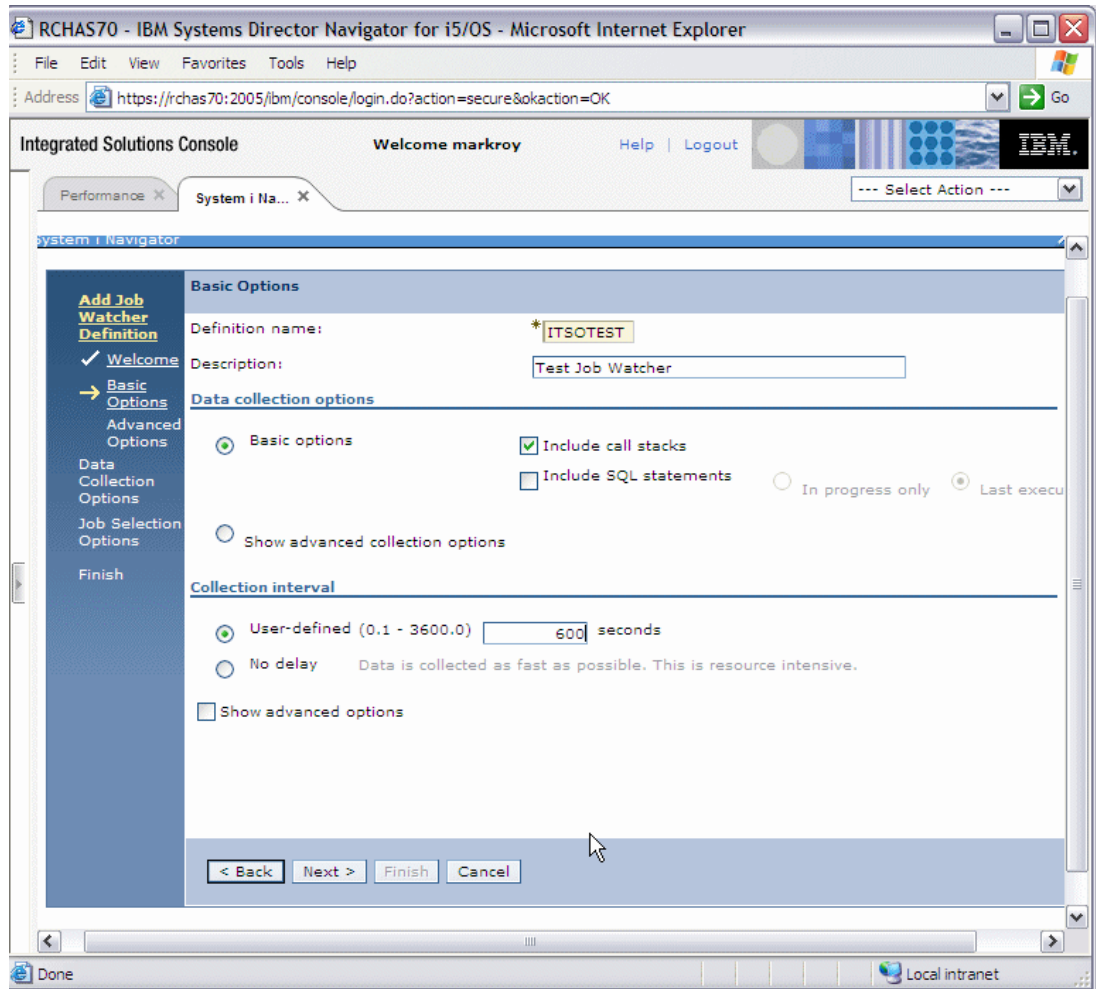


Figure 6-29 Naming the definition

At the Jobs and Task Selection page, select the **Include all jobs and tasks** option, as shown in Figure 6-30, and click **Next**.

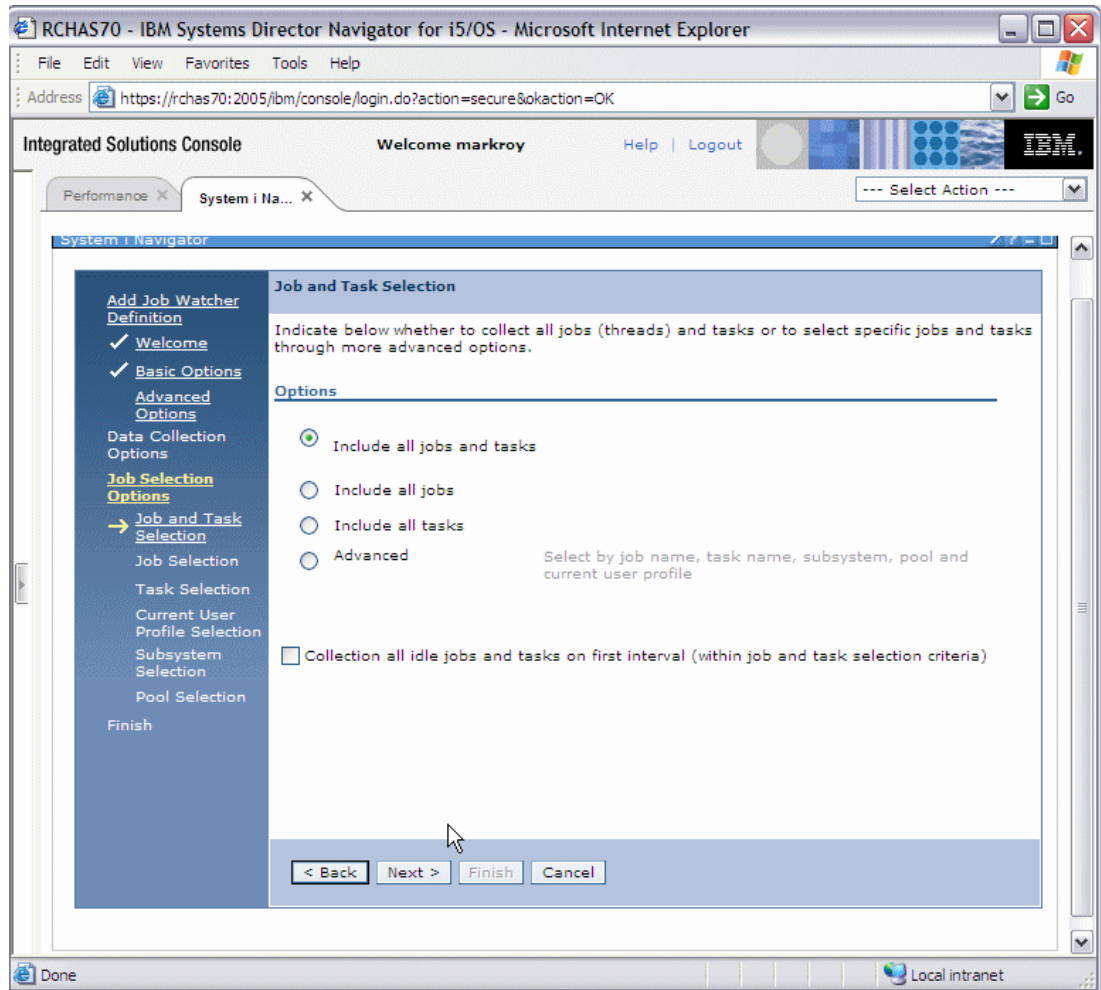


Figure 6-30 Selecting jobs to monitor

The summary pane displays information about the definition creation settings, as shown in Figure 6-31. Select **Finish** to create the definition.

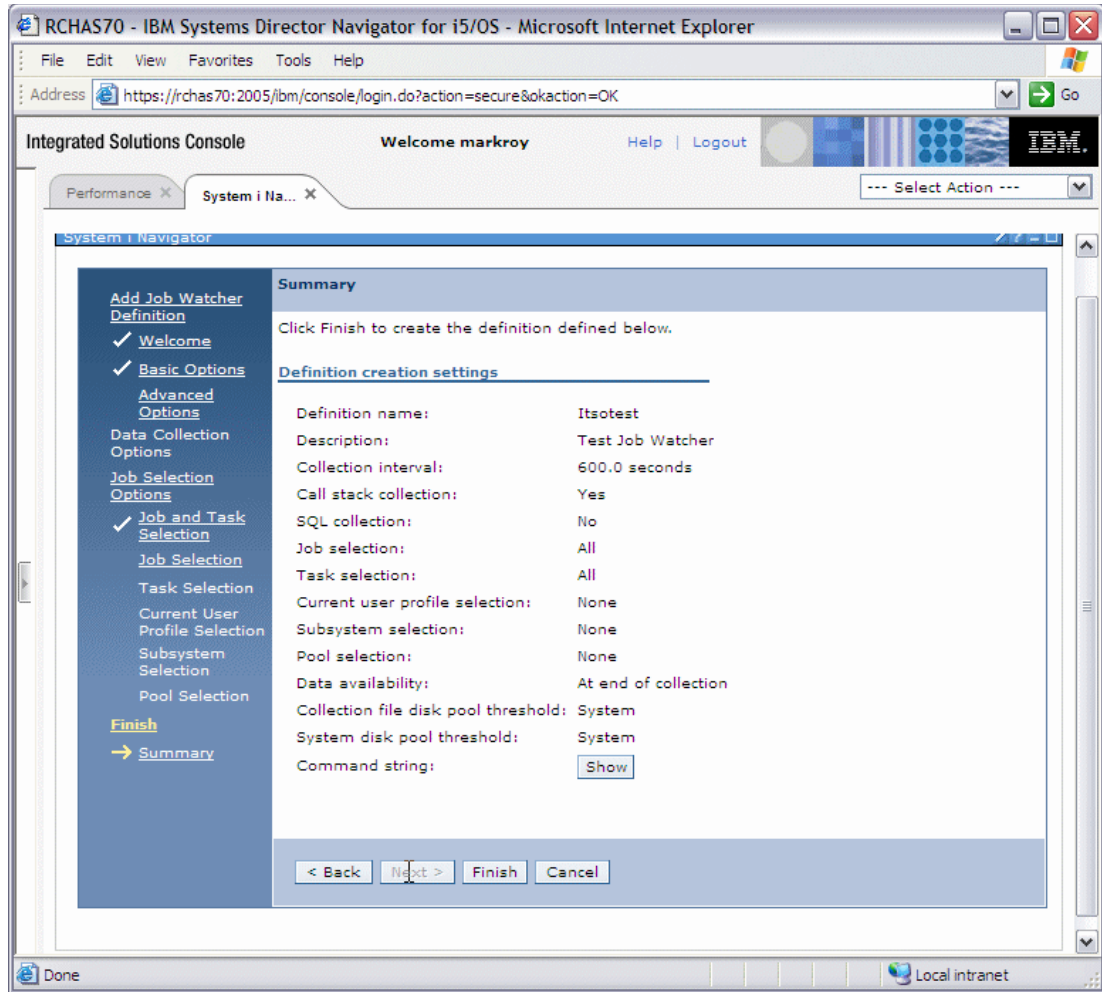


Figure 6-31 Summary and creating the definition

The next step is Start Job Watcher, as shown in Figure 6-32.

The first pane to be displayed is the Welcome pane where we select **Next** to continue.

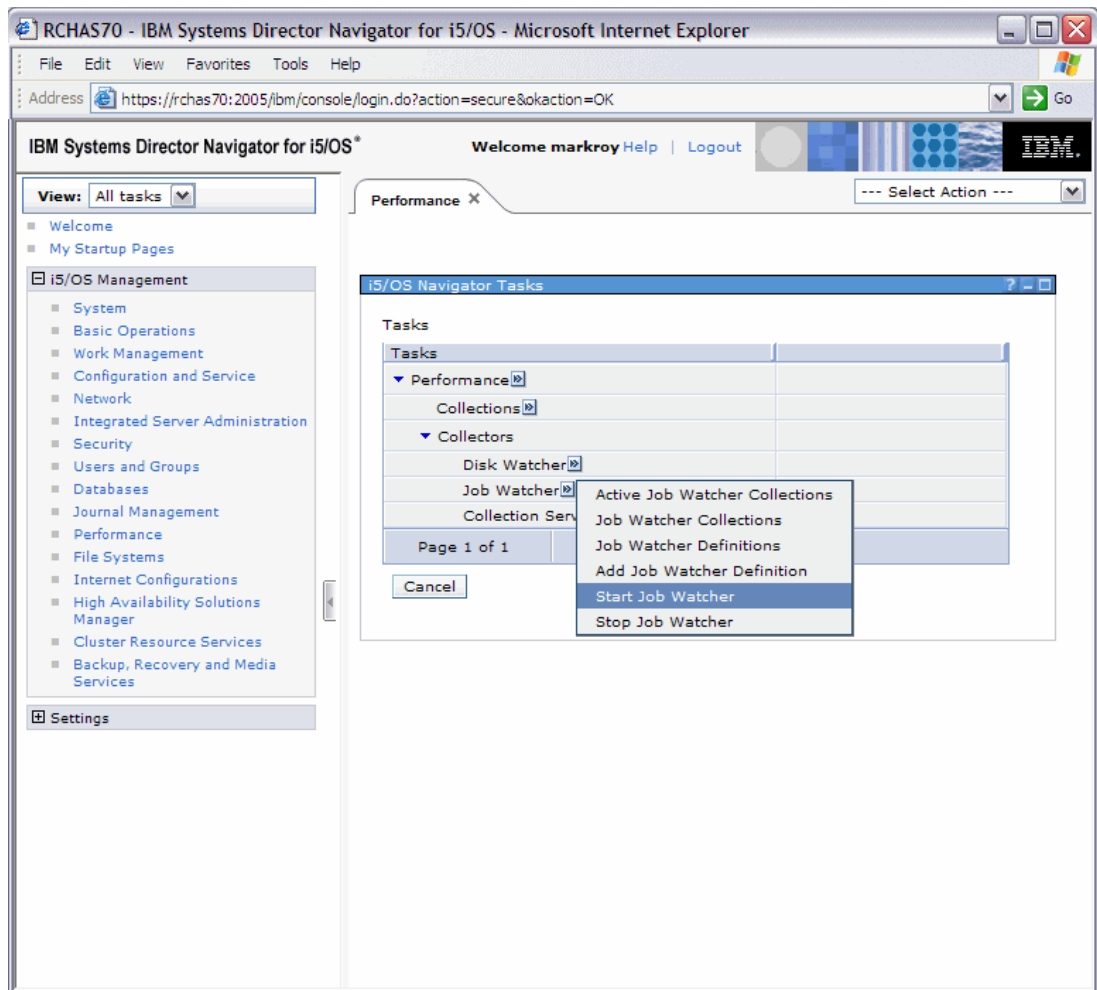


Figure 6-32 Starting Job Watcher

We are then presented with the Definition Selection pane where you can select from either IBM-supplied or user-defined definitions. In our example we select the **Q10sec** option, as shown in Figure 6-33, and select **Next** to continue.

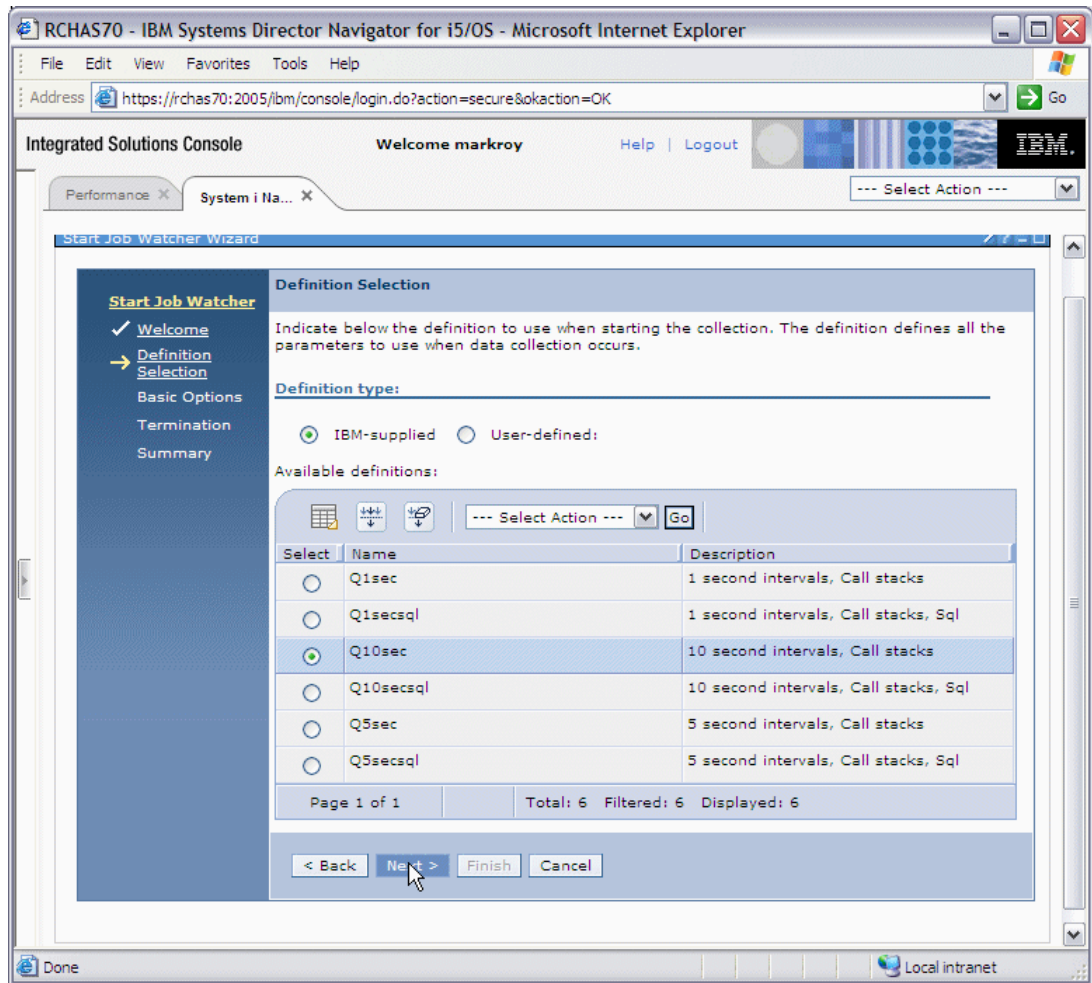


Figure 6-33 Specifying the intervals

On the Basic Options pane, shown in Figure 6-34, we enter a name for the collection name and a description. On this page you can change the collection interval if desired. In our example we change the collection interval to 10 seconds. Select **Next** to continue.

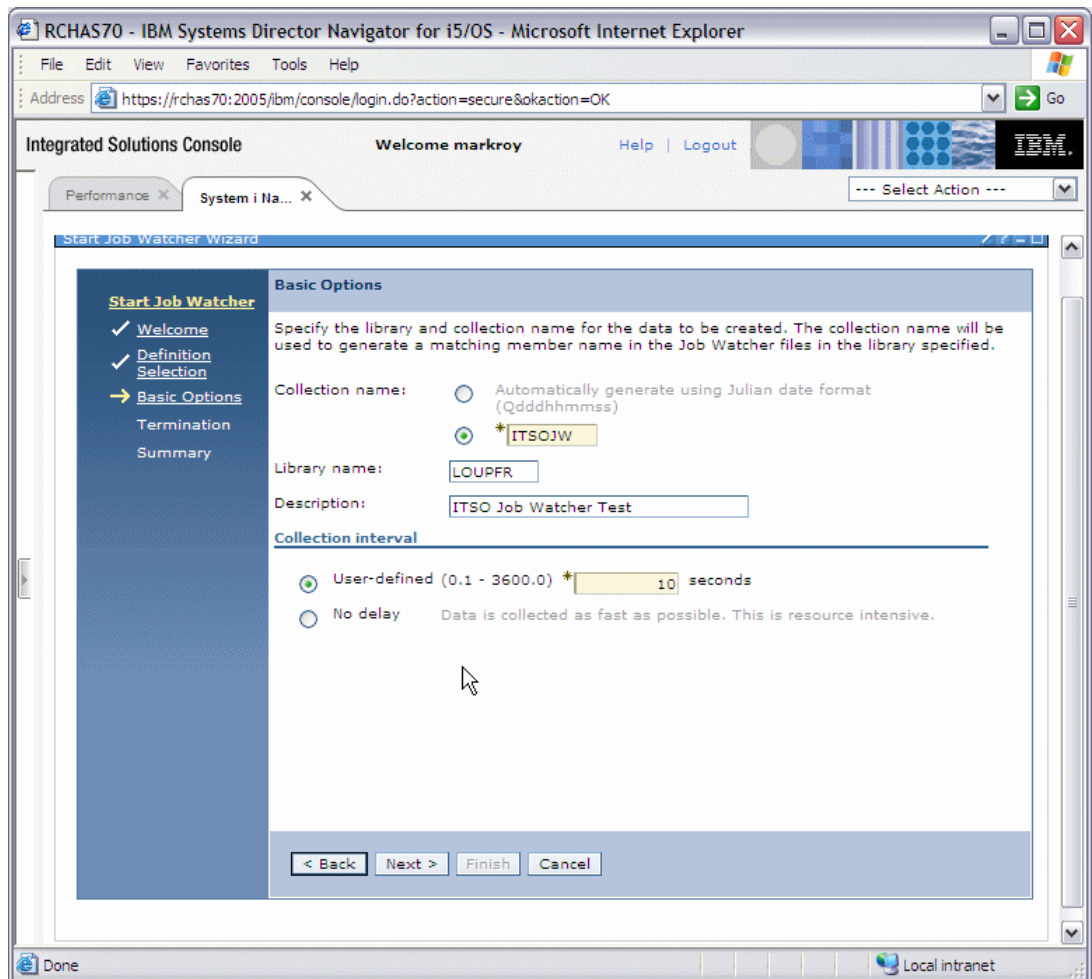


Figure 6-34 Basic options

Note: If the library you specify does not exist, you will be presented with a page where you can nominate to create it.

Select the ending options shown in Figure 6-35.

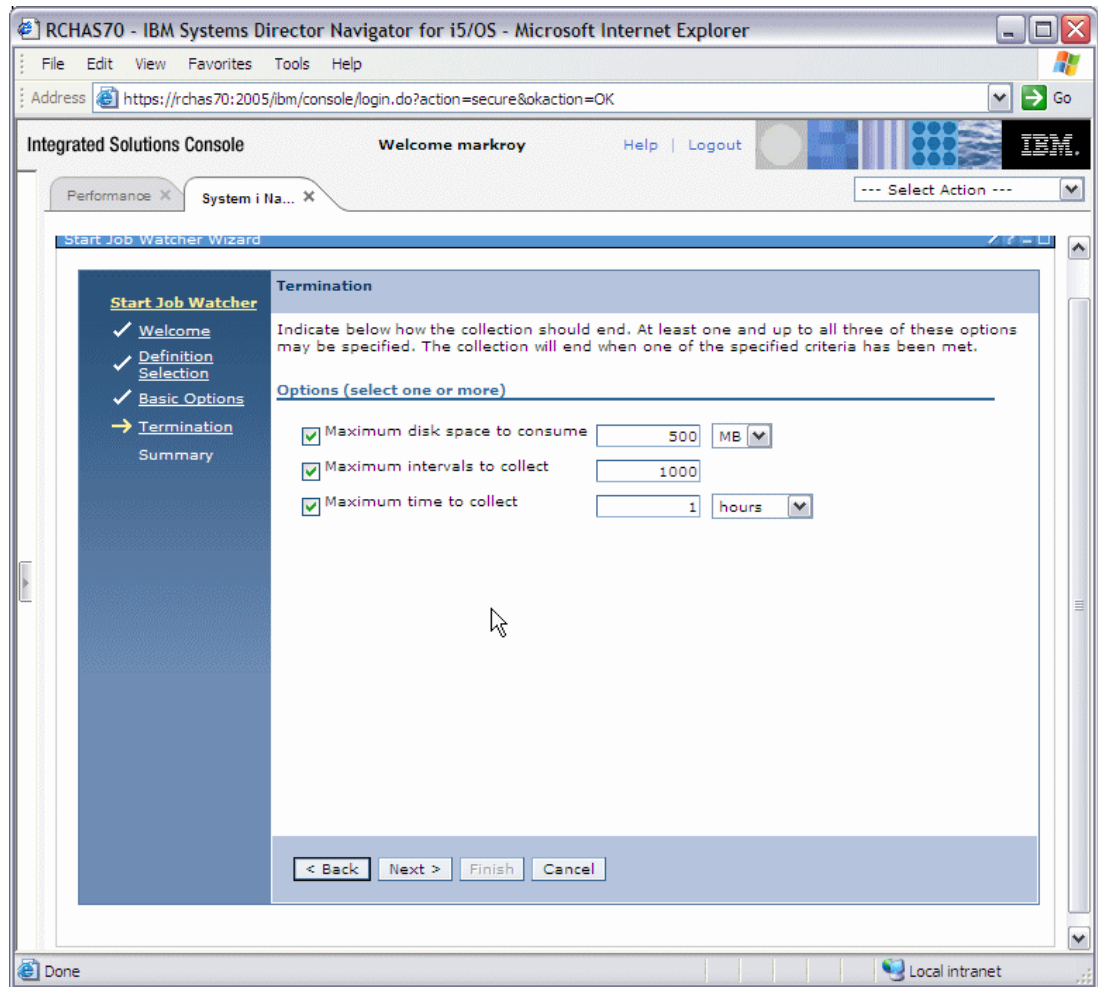


Figure 6-35 Ending options

The summary pane displays information about the collection creation settings, as shown in Figure 6-36 on page 179. Select **Finish** to create the collection.

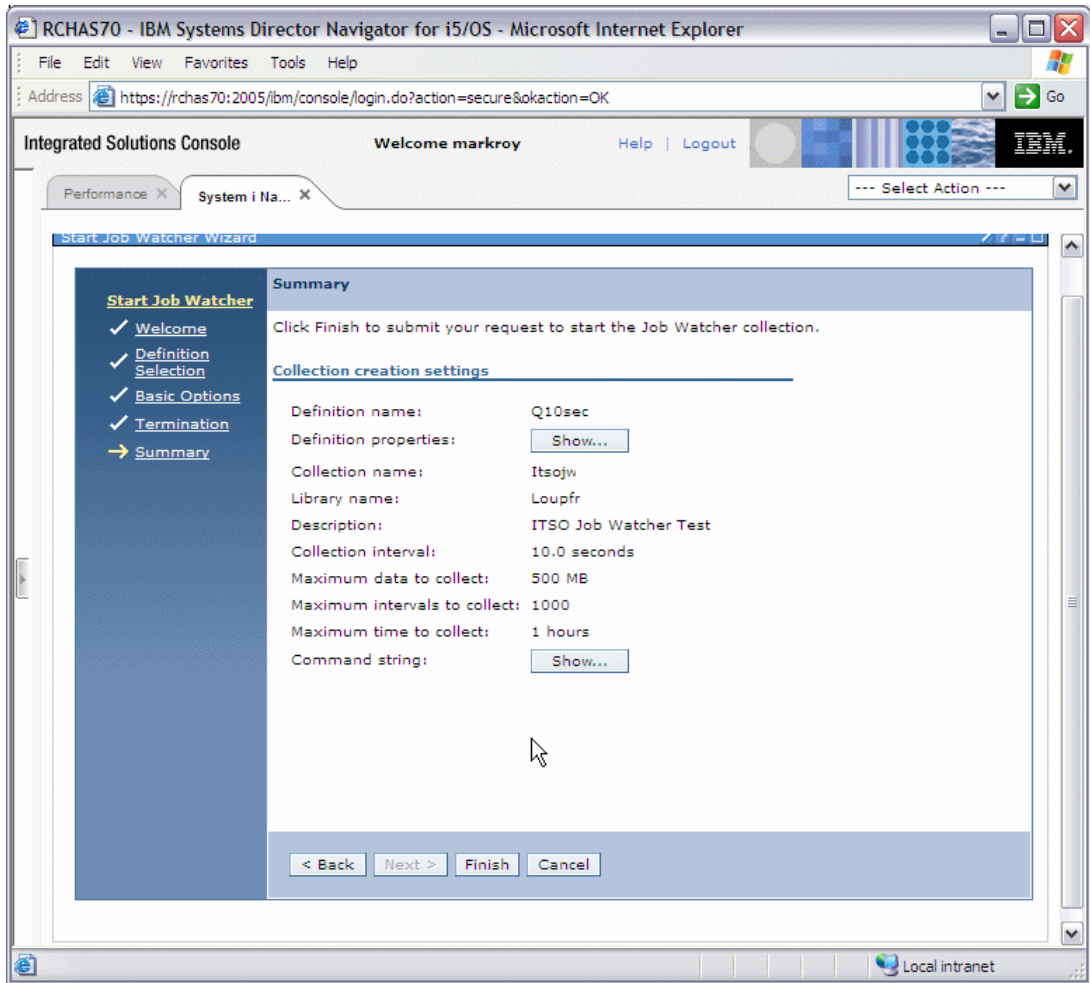


Figure 6-36 Summary and collection initiation

Figure 6-37 shows a CPU utilization and Wait Overview graph.




Figure 6-37 CPU Utilization and Wait overview

Note: iDoctor-Job Watcher, part of the iDoctor performance services offering remains and is enhanced for IBM i 6.1. Much but not all of iDoctor-Job Watcher capabilities have been ported to IBM i under the Performance Data Investigator interfaces through IBM Systems Director Navigator for i.

To get further information, refer to the *Systems Management Performance* reference information document at:

<http://publib.boulder.ibm.com/infocenter/iseri/v6r1m0/topic/rzahx/rzahxtbls.pdf>

6.7 Work with System Activity (WRKSYSACT)

The output of the WRKSYSACT command in IBM i 6.1 shows the number of Virtual Processors configured for the partition. In previous releases this was referred to as the Number of CPUs.

When a very large number of jobs are running in a partition, each using less than 0.1% of the CPU, they do not appear in the output of WRKSYSACT. When the number of this type of jobs is high enough, the system can still be very busy while none of them is showing up in WRKSYSACT. This behavior can be avoided in 6.1 by changing the Job/Task CPU Filter value. The minimum value can be set to 0.01% in 6.1.

Figure 6-38 shows an output of the Work with System Activity (WRKSYSACT) command.

Work with System Activity						RCHAS65			
						11/10/08		14:43:53	
Automatic refresh in seconds						5			
Job/Task CPU filter10			
Elapsed time : 00:00:02						Average CPU util :		4.8	
Virtual Processors : 2						Maximum CPU util :		9.0	
Overall DB CPU util : .0						Minimum CPU util :		.6	
						Current processing capacity: 2.00			
Type options, press Enter.									
1=Monitor job 5=Work with job									
	Job or					Total	Total	DB	
Opt	Task	User	Number	Thread	Pty	CPU Util	Sync I/O	Async I/O	CPU Util
	QPADEV000P	MARKROY	313392	00000304	1	4.7	14	0	.0
	QUMEPRVAGT	QSECOFR	245542	0000000C	50	.2	0	0	.0
	RMTMSAFETA				0	.2	0	0	.0

Figure 6-38 Work with System Activity (WRKSYSACT)

6.8 Performance Explorer

The Add PEX Definition (ADDPEXDFN) and Change PEX Definition (CHGPEXDFN) commands have been enhanced in IBM i 6.1.

6.8.1 Job name (JOB)

The job parameter has an additional element, *ALLSBS. Specifying *ALLSBS instructs PEX to search all subsystems to find jobs matching the specified job name, user, and number.

6.8.2 Maximum storage to use (MAXSTG)

The maximum storage for a trace has been increased from 40,000,000 to 250,000,000 kilobytes.

6.8.3 Program events (PGMEVT)

The following special values have been removed because they are no longer used:

- ▶ *MIPRECALL
- ▶ *MIPOSTCALL
- ▶ *JVAPRECALL
- ▶ *JVAPOSTCALL

6.8.4 Base events (BASEVT)

A new base event, *CPUSWT (CPU switch), has been added.

6.8.5 Lock events (LCKEVT)

Three new lock events have been added:

- ▶ *OBJLOCK
- ▶ *DBLOCK
- ▶ *SPCLOCK

6.8.6 Operating system events (OSEVT)

One of the previously unnamed special values, *MIEV24, has been renamed *ARMTRC.

6.8.7 Synchronization events (SYNCEVT)

A new synchronization event, *MTXCLEANUP (Mutex cleanup), has been added.

6.8.8 Virtual I/O events (VRTIOEVT)

The new event can be used to instruct PEX to collect virtual I/O events.

6.8.9 Program bracketing events (PGMBKTEVT)

Starting with IBM i 6.1, the special value *PRC will behave the same as *MIENTRYEXIT.

6.8.10 Convert Performance Collection (CVTPFRCOL)

The convert Performance Collection (CVTPFRCOL) command converts performance data from a previous release to the format needed for processing by the current release of Performance Tools for IBM i.

The command first determines the release level at which the data was collected. Then the members of all necessary files are converted. The conversion may be done in the same library where the current data resides. To avoid the risk of destroying the old data if the command ends abnormally, convert the data into a different library (TOLIB), and later, delete the data from the old library (FROMLIB).

If the conversion is done in a different library, the old data remains in the current library (FROMLIB) and the new data resides in the new library (TOLIB). If a new library is specified for the newly converted data, all files are copied to the new library, including those files which do not need to be converted.

6.9 Collection Services

IBM i 6.1 brings some changes in how performance collection objects are managed by IBM i.

Collection Services has never run automatically as shipped. However, PM for System i was dependent on Collection Services data and the PM Agent would automatically start it. In OS/400 V5R2, the PM Agent changed to always collect data even if the customer was not sending it to IBM. For all practical purposes Collection Services has been automatically running on all customer systems unless PM Agent was shut down.

IBM i 6.1 changes the focus from Collection Services being secondary or manually started, to Collection Services being a key system function that IBM intends to be running. Although the collector can still be manually stopped and started, the automatic IPL startup feature will not be affected. Additionally, PM for System i will no longer restart Collection Services and it will respect your decision to turn Collection Services off or change the configuration if you have done so.

6.9.1 Create Performance Data (CRTPFRDTA)

Additional indices in the performance database are necessary to facilitate a quick response time of the Systems Director Navigator for i performance Investigate Data tasks. These additional files are not part of the base CRTPFRDTA function. To get this new summary data created, you must specify this under the new 6.1 parameter "Create performance summary" on the Configure Performance Collection (CFRPFCOL) command.

If the option to create summary data was not selected when the performance database files were created, use the new CRTPFRSUM command to create these summary files. The performance summary data is based upon the already generated QAPMnnnn performance database files created either through:

- ▶ Configure the Collection Services "create database files" parameter.
- ▶ The IBM i CRTPFRDTA command.

Figure 6-39 on page 183 shows Create Performance Data (CRTPFRDTA) command options.

```

                                Create Performance Data (CRTPFRDTA)

Type choices, press Enter.

From collection . . . . . Name, *ACTIVE, *SELECT
  Library . . . . . QPFRDATA Name
To member . . . . . *FROMMGTCOL Name, *FROMMGTCOL
To library . . . . . *FROMMGTCOL Name, *FROMMGTCOL
Text 'description' . . . . . *SAME

Categories to process . . . . . *FROMMGTCOL Name, *FROMMGTCOL, *APPN...
      + for more values
Time interval (in minutes) . . . *FROMMGTCOL *FROMMGTCOL, 0.25, 0.5, 1...
Starting date and time:
  Starting date . . . . . *FROMMGTCOL Date, *FROMMGTCOL
  Starting time . . . . . Time
Ending date and time:
  Ending date . . . . . *FROMMGTCOL Date, *FROMMGTCOL, *ACTIVE
  Ending time . . . . . Time

```

Figure 6-39 Create Performance Data (CRTPFRDTA)

6.9.2 Configure Perf Collection (CFGPFRCOL)

In IBM i 6.1, two new parameters have been added to the Configure Perf Collection (CFGPFRCOL) command.

Figure 6-40 shows Configure Perf Collection command options.

```

                                Configure Perf Collection (CFGPFRCOL)

Type choices, press Enter.

Default interval . . . . . 15.00 *SAME, .25, .50, 1.0, 5.0...
Collection library . . . . . QPFRDATA Name, *SAME
Default collection profile . . . *STANDARDP *SAME, *MINIMUM, *STANDARD...
Cycle time . . . . . 000000 Time, *SAME
Cycle interval . . . . . 24 *SAME, 1-24 hours
Collection retention period:
  Number of units . . . . . 00120 *SAME, 1-720, *PERM
  Unit of time . . . . . *HOURS *HOURS, *DAYS
Create database files . . . . . *YES *SAME, *YES, *NO
Create performance summary . . . *NO *SAME, *YES, *NO

                                Additional Parameters

Standard data retention (days) 000000010 Number, *SAME, *PERM

```

Figure 6-40 Configure Perf Collection (CFGPFRCOL)

Standard data retention (days) (STDDTARET)

This parameter specifies the retention period for standard file data. Standard file data is that data normally collected for a collection profile on an ongoing basis that has been exported to performance data base files (QAPM*). The standard file data retention period is used to determine how long performance database file collections should be retained on the system. When standard data in the configured collection library is older than the current retention period, it will be automatically deleted by the Collection Services server job (QYSPFRCOL) the next time the collector is started or cycled. The retention period is specified in days.

Notes:

- ▶ Collection Services supports expiration and deletion of both *MGTCOL and DB file collections (that is, the members of the QAPMXXXX performance database files) that exist in the configured collection library QPFRDATA as system-owned data. Collections and performance database files in other libraries are not managed.
- ▶ An expired DB collection will not be deleted if PM Agent is running and still needs to process it.

Create performance summary (CRTPFRSUM)

The parameter specifies whether or not to create additional performance database summary data as supported by the CRTPFRSUM command. This data facilitates quicker processing of the performance database data by tools such as the “Investigate Data”. If this option is enabled, and CRTDBF(*YES) was specified, the CRTPFRSUM function will be performed for the database collection created by the CRTPFRDTA job.

6.10 Performance Management Agent

The portion of the IBM Performance Management for System i5® offering that runs on IBM i has been named Performance Management Agent (PM Agent).

PM Agent no longer starts Collection Services. The database collection cleanup function has been removed and added to Collection Services and is configurable.

The PM Agent CL commands have new names, as listed in Table 6-3.

Table 6-3 PM Agent CL new command names

Old CL command	New CL command	Description
CFGPM400	CFGPMAGT	Configure PM Agent
PMLINMON	CFGPMLIN	Configure PM Agent Line

6.11 System i Navigator On Demand Performance Center enhancements

System i Navigator On Demand Performance Center has improved some functions such as customizable reports, copy results to a spreadsheet, enhanced plan cache flexibility, and analysis. Here is the list of enhancements:

- ▶ Spreadsheet Integration
- ▶ Column Customization

- ▶ Index Advice Condenser (PTF)
- ▶ Fast monitor summary compare
- ▶ Plan Cache Resize & Event Monitor
- ▶ Enhanced Plan Cache Analysis
- ▶ Visual Explain While Running
- ▶ Special Client Register

For further information about how to improve database performance using System i Navigator tools, refer to *System Management Performance* reference information at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzahx/rzahxtb1s.pdf>

6.12 Workload Estimator

The IBM Systems Workload Estimator (a.k.a. the Estimator or WLE) is a Web-based sizing tool for IBM Power Systems™, System i, System p, and System x, available at:

<http://www.ibm.com/systems/support/tools/estimator/>

You can use this tool to size a new system, to size an upgrade to an existing system, or to size a consolidation of several systems. The Estimator allows measurement input to best reflect your current workload and provides a variety of built-in workloads to reflect your emerging application requirements. Virtualization can be reflected in the sizing to yield a more robust solution, by using various types of partitioning and virtual I/O. The Estimator provides current and growth recommendations for processor, memory, and disk that satisfy the overall customer performance requirements.

6.12.1 Version availability

The intent is to provide a new version of the IBM Systems Workload Estimator three to four times a year. Each version includes an update message after approximately three months.

The IBM Systems Workload Estimator is available in two formats, on-line and as a download. Both are described at:

<http://www.ibm.com/systems/support/tools/estimator/>

The on-line version is the preferred way to access the Estimator.

It is also highly recommended that there should be involvement of IBM Sales or IBM Business Partners before making any purchasing decisions based on the results obtained from the Estimator.

The approximate size requirements are about 16.5 MB of hard disk space for the Estimator and 60 MB for the server setup. A rough expectation of the time required to install the entire tool (server and Estimator) is 20 minutes. A rough estimate of the time required for installing the update to Estimator only, assuming the Server was set up previously, is about 5 minutes.

Note: Running the server on the client may not run as fast as over the Internet due to the additional processing requirements on the client.

While releases of WLE are not tied to IBM i releases, we have included the 2008 enhancements to WLE.

Note: The following WLE functions were provided as part of v2008.4 and 2008.4.1 (October 2008).

6.12.2 New and enhanced IBM Power system models

WLE supports enhancements to the powerful line of unified servers based on industry-leading POWER6 processor technology. These models are now fully converged supporting AIX, IBM i, and Linux operating systems. Make sure that you utilize WLE to help make the best model selection to satisfy the requirements for your overall workload, especially if you exploit virtualization.

- ▶ Power 520 (8203-E4A), 4.2 GHz, with 1-4 cores
- ▶ Power 550 (8204-E8A), 3.5 and 4.2 GHz, with 2-8 cores
- ▶ Power 560 (8234-EMA), 3.6 GHz, with 4-16 cores
- ▶ Power 570 (9117-MMA), 3.5, 4.2, 4.4, 4.7, and 5.0 GHz, with 2-32 cores
- ▶ Power 595 (9119-FHA), 4.2 and 5.0 GHz, with 8-64 cores

6.12.3 New System x models

Support for the most recent IBM System x models was added.

6.12.4 Energy estimation

WLE is now able to provide an energy consumption estimate for most recommended Power System configurations. Notice the Energy Estimate function on the Selected System panel. WLE estimates typical power requirements (Watts and BTU/hour) for a specific system configuration under normal operating conditions. Having this function integrated with WLE allows you to easily obtain this estimate. The IBM Systems Energy Estimator, a standalone Web-based version exists at:

<http://www.ibm.com/systems/support/tools/estimator/energy/>

6.12.5 Disk storage system sizing support for DS5000 and XIV

WLE was recently enhanced with sizing support for Disk Storage Systems (a.k.a SAN). WLE will make a high-level recommendation using disk storage systems to satisfy your disk I/O requirements for your overall solution if you specify external storage in the WLE User Options; then you can further tailor your SAN solution with the External Storage function on the Selected System panel. WLE now has sizing support for the DS5000 family as well as the IBM XIV® Storage System.

6.12.6 License and Web standard

The WLE license has been updated and now can be viewed in multiple languages. Also, the WLE user interface was updated to use more recent IBM.com standards. You may notice the look and feel changed slightly, but it does not affect any functions.

6.12.7 Import customer performance data

When WLE imports measured customer data (through PM or PTX/Topasout), in some cases not enough information was passed along to find an exact match to a unique model or feature. Under those circumstances, WLE will now post to the user what it can determine and ask the user to make the final unique model selection. For performance analysis of upgrades or consolidations, it is critical that the existing system be identified.

6.12.8 Workload Developer Toolkit enhancements

The Workload Developer allows others to build workload plug-ins for the Workload Estimator to characterize applications (like the IBM sizing guides referred to in the next section). Refer to the *What's New* information found in the help menu of the Developer for the latest enhancements. Remember, the Developer is included in the download version of Workload Estimator.

6.12.9 Workloads and sizing guides

WLE supports nearly 200 workload plug-ins now for a variety of IBM and ISV products. With each new WLE version, many workloads are updated. Some workloads are provided internally to WLE, while many can be found at:

<http://ibm.com/servers/sizing>

6.12.10 Support for IBM i and Linux on Power 595 (9119-FHA)

WLE now supports the powerful lineup of unified servers based on industry-leading POWER6 processor technology with support for IBM AIX, IBM i, and Linux operating systems.

- ▶ Power 520 (8203-E4A), with WLE sizing support for AIX and Linux
- ▶ Power 520 Express, i Edition (9407-M15 and 9408-M25) with WLE sizing support for IBM i, AIX, and Linux
- ▶ Power 550 (8204-E8A), with WLE sizing support for AIX and Linux
- ▶ Power 550 Express, i Edition (9409-M50), with WLE sizing support for IBM i, AIX, and Linux
- ▶ Power 570 (9117-MMA), with WLE sizing support for AIX, IBM i, and Linux
- ▶ Power 595 (9119-FHA), with WLE sizing support for AIX, IBM i, and Linux

6.12.11 New System x models

Support for the most recent IBM System x models was added.

6.12.12 Sizing capability for disk storage systems

WLE has been enhanced with sizing support for Disk Storage Systems (a.k.a SAN). WLE will make a high-level recommendation using disk storage systems to satisfy your disk I/O requirements for your overall solution if you specify this in the WLE User Options. Make sure you set the Disk Attachment Type parameter to *external*. WLE will also provide you the ability to tailor your DS n solution by selecting a specific DS n system model or by selecting a particular DS n family, or by selecting a business category (entry, midrange, or enterprise) of DS n systems. On the WLE Selected System panel, you will notice the summarized disk

storage system recommendation as well as the individual recommendations for each partition. You can investigate some what-if questions and further tailor your SAN solution by using the External Storage function on the Selected System panel. If you wish to do a more detailed SAN sizing analysis, you may also use the Disk Magic sizing tool with the existing WLE or Disk Magic interface also found on the External Storage function.

6.12.13 WebSphere Process Server and WebSphere Enterprise Service Bus

A WPS/WESB workload has been added to WLE for IBM i covering WebSphere Process Server (WPS) and WebSphere Enterprise Service Bus (WESB) version 6.1. This workload will allow you to size your Service-Oriented applications.

6.12.14 Performance Management interface to Workload Estimator

WLE was enhanced to be more flexible with customer data import. If WLE does not have enough information to properly identify the model or LPAR configuration, you are given a drop-down to help match the data to the actual system model. More information is now displayed in the detailed data and the advanced option sections to help validation.

6.12.15 System Planning Tool linkage

WLE has previously supported linkage to provide SPT with the WLE recommendation by using the *Plan Immediate/Growth System* function on the Selected System panel. However, this function would only be available if you started the experience from a client running SPT. Many of you have requested a more robust interface here. WLE was enhanced to generate a file to import to SPT at your own convenience. Generate this by using the same Plan Immediate/Growth System function.

6.12.16 Ability to tailor the PDF contents

WLE continues to offer a PDF version of the recommendation. With this release you can customize the contents of the PDF. You will notice an additional navigation step when the PDF function is used on the Selected System panel.

6.12.17 Ability to limit sizing recommendations to POWER6 models

Each time the WLE sizing engine does a sizing, it examines all possible models. The system models that satisfy the requirements will appear in the drop-down; the models that do not meet the requirements appear in the rejected system list. If you know that you want to limit solutions to POWER6 models, you can improve the performance of the WLE tool by limiting the models to evaluate. In WLE User Options, set Family to POWER6.

6.12.18 Workload Developer toolkit enhancements

The Workload Developer allows others to build workload plug-ins for the Workload Estimator to characterize applications (like the IBM sizing guides referred to in the next section). Support has been added to better allow the mixing of various workload types in a more robust way. Refer to the *What's New* information found in the help menu of the Developer for the latest enhancements. Remember, the Developer is included in the download version of Workload Estimator.

6.12.19 New IBM Power System servers

IBM System i and IBM System p are unifying the value of their servers into this single, powerful lineup of servers based on industry-leading POWER6 processor technology with support for IBM AIX, IBM i, and Linux operating systems.

- ▶ Power 520 (8203-E4A), with WLE sizing support for AIX and Linux.
- ▶ Power 520 Express, i Edition (9407-M15 and 9408-M25) with WLE sizing support for IBM i, AIX, and Linux.
- ▶ Power 550 (8204-E8A), with WLE sizing support for AIX and Linux.
- ▶ Power 550 Express, i Edition (9409-M50), with WLE sizing support for IBM i, AIX, and Linux.
- ▶ Power 570 (9117-MMA), with WLE sizing support for AIX, IBM i, and Linux.
- ▶ Power 595 (9119-FHA), with WLE sizing support for AIX (IBM i this summer, and Linux at a future date).

6.12.20 Converged sizing enablement for System i and System p

In prior versions, you may have started a WLE sizing experience with a function called *New System i Estimation* or *New System p Estimation*. These two brand personalities have been converged within WLE; now you may start a sizing experience with the function called *New IBM Power Systems Estimation*. Simply select the operating system for a system or partition and add workloads as you have always done in WLE. Just as before, WLE will recommend the best fit solution. If you want to change the model that WLE recommended, the model drop-down will include all systems (Power Systems, System i, System p, blades) that meet the operating system and performance requirements.

6.12.21 New POWER6 blade for IBM BladeCenter JS12 Express

WLE sizing support was added for this new 2-core POWER6 blade. WLE user options exist such that an estimation can have a solution form factor preference of blade, non-blade, or all models (the default). If all models are considered for the solution and if the blade meets all the workload requirements, you can select this blade with the model drop-down on the selected system panel. VIOS is required if there are any partitions running IBM i.

6.12.22 New IBM System x models

Support for the most recent IBM System x models was added.

6.12.23 New Windows Server 2008 support

WLE now supports sizing workloads that use the Windows Server 2008 Operating System. This new operating system, which just started shipping, is of major strategic importance to System x and the necessary support has just been added to WLE.

6.12.24 New disk attachment types, new SAS storage

Two new disk attachment type families were added to WLE. These families include the recent SCSI and SAS CCINs as well as the new SAS CCINs. POWER6 DAS includes CCINs 572C, 572A, 57B9; and POWER6 Cached DAS includes 571E, 571F, 57B7, 57B8, and 572B. Disk

attachment types are groupings of storage types that have similar performance characteristics. The desired storage attachment type for the WLE recommendation can be set in the WLE User Options. The associated help text can assist you in the selection with a list of all the CCINs, features codes, and a summary of the relative performance characteristics. Note that the default disk attachment type has changed to be POWER6 Cached DAS.

6.12.25 Updates to POWER6 virtualization for IBM Power Systems

Characterizations in WLE for virtualization support are updated to best reflect the performance of IBM Power Systems.

6.12.26 Increased session time-out

The Web-based WLE previously had session time-out settings of 20 minutes. Some users requested that this be increased to avoid losing the progress of their estimations while they were interrupted, perhaps with a long phone call. With this version of WLE, sessions will be extended significantly. Note that WLE continues to allow estimations to be saved (Save All Workloads on the Selected System panel) and restored (Restore Estimation on the Workload Selection panel).

6.12.27 Performance management for Power Systems

Support for WLE to import measured customer performance data has been enhanced again for PM for System i and PM for System p. This support is now called PM for Power Systems. Give this a try sometime and see how WLE can utilize historic measured data to consider trending for assistance in sizing your next upgrade.

6.12.28 Workload Developer enhancements

The Workload Developer allows others to build workload plug-ins for the Workload Estimator to characterize applications (like the IBM Sizing Guides referred to previously). Support has been added to allow modeling of various software virtualization programs including VMWare. Be sure to read the *What's New* information found in the help menu of the Developer for the latest enhancements. Remember, the Developer is included in the download version of Workload Estimator.

6.12.29 IBM Sizing Guides for Power Systems, System i, System p, System x

There are frequent updates to the list of workloads supported at *IBM Sizing Guides*:

<http://ibm.com/servers/sizing>

One of the most notable updates was with VMWare, which supports consolidating multiple servers.

6.12.30 Workload updates

Many WLE workload plug-ins were updated with this release to better characterize the workload with POWER6 and Power Systems. The IBM i WebSphere workload was updated to add support for sizing WebSphere Application Server 6.1 using IBM Technology for Java (64-bit). Others that were updated include Domino, Portal, HATS, and Web Facing 7.1.

Note: The following items originally shipped as part of v2008.1 (January 2008).

6.12.31 New IBM System p POWER6 models

Announce-day support was added to the Workload Estimator (WLE) for the p520-8203-E4A (up to four cores) and the p550-8204-E8A (up to eight cores). Continue to use WLE in the usual way; the recommendation will include these models when appropriate.

6.12.32 New IBM System x models

Support for the most recent IBM System x models was added.

6.12.33 IBM i Support on the POWER6 blade

With IBM i 6.1 and VIOS client support, IBM i runs on the IBM BladeCenter JS22 blade. WLE user options exist such that an estimation can have a solution form factor preference of blade, non-blade, or all models (the default). If all models are considered for the solution and if the blade meets all the workload requirements, you can select this blade with the model drop-down on the selected system panel.

6.12.34 New disk models

Several new disk attachment types were added (CCIN: 57B7, 57B8), including the new SAS RAID Enablement feature. You may specify the desired storage attachment type for the estimation in the User Options. The associated help text can assist you in the selection with a list of all the CCINs, features codes, and a summary of the relative performance characteristics.

6.12.35 IBM i 6.1 and Virtual I/O Server client support

This new IBM i major release has much new function including significantly enhanced performance for environments that use Java, WebSphere, and SQL. With this new release and with POWER6 processors, an IBM i partition now has VIOS client support. This means that disk I/O and/or Ethernet I/O can be redirected to the VIOS partition. WLE allows the user to optionally select VIOS support for workloads in AIX, Linux, or IBM i partitions; this support can be activated by answering the workload questions in the workload definition tab.

6.12.36 Logical Partitioning and Virtual I/O Server algorithm updates

The characterization of virtualization support in WLE has been updated based on new POWER6 performance data and to be more customer representative. When sizing System i or System p with virtualization configurations, be sure to consider using WLE to help build a configuration recommendation; even if you have used another source for sizing a given workload or product, you can utilize WLE (perhaps using the generic workload) to size the system solution using virtualization.

6.12.37 Restore and quick validation

Previously when you would restore a prior estimation in WLE, you would have to navigate through each panel for each workload to review and verify the answers to the workload questions. For large configurations with many workloads and/or many partitions, this can lead to using the continue function many times. With this new version, WLE attempts to validate all the workloads at the time of the restore. You will notice the presence of a green check or a yellow caution symbol indicating whether or not the quick validation occurred. If you wish, just navigate to the workloads with the yellow marking to complete the validation; when complete, the Selected System tab will appear and you can proceed directly to the recommendation.

6.12.38 Performance Management quality improvement

Performance Management for System i and System p was updated. The capacity planning interface to Size Next Upgrade was improved to better handle old and new formats for disk data as well as other system information. Consider using this WLE interface as data input for planning your next upgrade or consolidation. Note that you can import system or partition data multiple times to build a consolidation estimation.

6.12.39 Workload Developer

The Workload Developer allows others to build workload plug-ins for the Workload Estimator to characterize applications (like the IBM Sizing Guides referred to previously). Be sure to read the *What's New* information found in the help menu of the Developer for the latest enhancements. Remember, the Developer is included in the download version of Workload Estimator.

6.12.40 IBM i Workloads

Several workloads were updated including Domino to support Sametime® 7.5, WebSphere for improved performance, HATS, and WebFacing.

IBM Sizing Guides for System i, p, and x.

There are frequent updates to the list of workloads supported at *IBM Sizing Guides* at:

<http://ibm.com/servers/sizing>

One of the most notable updates was with the Exchange 2007 workload to allow multiple sites.

6.12.41 3Com Telephony workload

The 3Com IP Telephony workload has been updated to add support for Domino and Sametime integration.

6.12.42 Disk adapters and I/Os per second ratings

Workload Estimator contains a table detailing some performance characteristics of different disk attachment families. The table can assist you in determining which disk attachment family to consider. You can select the family row that best matches your configuration requirements. The throughput values are listed in terms of *system disk operations* per second per disk drive at 40% busy in a RAID5 storage protection environment. The term system disk

operations includes all disk operations issued by the operating system including those for software mirroring. This would typically match up with what operating system performance tools would report; it does not include additional disk operations generated by the disk subsystem for RAID or for hardware mirroring. The typical values for read/write ratios and cache efficiencies correspond to the throughput values listed for each attachment type. So you would expect those with higher cache efficiency values to handle higher throughput rates for system disk operations. A cache efficiency value of 0 implies there is no cache.

The Estimator uses these values for internal disk calculations; significant swings in these values can greatly impact the potential throughput of the disk IOA. Table 6-4 will assist you in making a high-level decision.

Note: Actual performance may vary significantly.

Table 6-4 Disk adapters and I/Os per second ratings

Disk Attachment Type Family Name (write or read cache)	IOA CCIN (some may not be listed)	Orderable Feature Code (FC) (some may not be listed)	Sys Disk Ops/sec per arm @ 40% busy with RAID5	%Read / %Write	IOA Read Cache Efficiency	IOA Write Cache Efficiency
			RPMs 7.2K / 10K / 15K			
Original IOA/IOP	2778	2778	50 / 60 / 92	66 / 34	0	70
▶ IOA Feature #5705 ▶ IOA Feature #1974	1974, 5702, 570A, 570B, 573E	0624, 1974, 5702, 5705, 5710, 5712, 5715	38 / 46 / 68	75 / 25	0	0
▶ IOA Feature #5703 (40 MB write cache) ▶ IOA Feature #5709 (16 MB write cache) ▶ IOA Feature #5727 (40 MB write cache)	1908, 1975, 1976, 2782, 5703, 5709, 5727, 573D	0619, 0628, 1907, 1908, 1975, 1976, 2782, 4510, 5703, 5709, 5711, 5726, 5727, 5728	60 / 75 / 115	66 / 34	0	73
IOA Feature #2757 (235/757 MB write cache)	2757	0618, 2757, 5581(w/aux. write cache)	69 / 84 / 127	66 / 34	0	79
IOA Feature #2780 (235/757 MB write cache)/(256 MB/ 1GB read cache)	2780	0627, 2780, 5554 (w/mirroring), 5555 (w/mirroring), 5556 (w/mirroring), 5580 (w/aux. write cache)	73 / 89 / 134	66 / 34	7	79
IOA Feature #5736	1912, 571A	1912, 5736	42 / 51 / 76	75 / 25	0	0
IOA Feature #5737 90 MB write cache)	1913, 571B,	1913, 5737	66 / 83 / 127	66 / 34	0	75
POWER6 DAS	572C, 572A, 57B9	0664, 5779, 5900, 5906, 5909, 5912	NA / 51 / 76	75 / 25	0	0
POWER6 Cached DAS (390 MB / 1.5GB write cache)/(415 MB/ 1.6GB read cache)	571E, 571F, 57B7, 57B8, 572B Some with aux. write cache	0649, 0650, 0651, 0654, 5582, 5583, 5679, 5738, 5739, 5778, 5781, 5782, 5799, 5800, 5886, 5902,	90 / 110 / 165	66 / 34	7	79
Ultra SCSI 3	NA	NA	50 / 60 / 80	66 / 34	0	0

Disk Attachment Type Family Name (write or read cache)	IOA CCIN (some may not be listed)	Orderable Feature Code (FC) (some may not be listed)	Sys Disk Ops/sec per arm @ 40% busy with RAID5	%Read / %Write	IOA Read Cache Efficiency	IOA Write Cache Efficiency
			RPMs 7.2K / 10K / 15K			
Ultra SCSI 2	NA	NA	35 / 48 / 60	66 / 34	0	0
Serial ATA	NA	NA	40 / 47 / 72	67 / 33	0	0
Serial SCSI	NA	NA	75 / 90 / 120	67 / 33	7	50



Logical partitioning

This chapter discusses the following topics:

- ▶ “PowerVM” on page 196
- ▶ “IBM i 6.1 host and client partitions” on page 207

7.1 PowerVM

To assist understanding the new LPAR capabilities in IBM i 6.1, we start with an overview of the underlying technology that supports logical partitioning, PowerVM™.

PowerVM is a brand that includes a range of hardware and software features that enable a system to be more useful in adapting various workloads.

PowerVM is the new brand name for Power Systems virtualization. It includes elements such as those shown in Table 7-1.

Table 7-1 PowerVM functions

Function	Description
Hypervisor	Supports multiple operating environments on a single system.
Micro-Partitioning™	Fractional processor capacity. Enables up to 10 partitions per processor core.
Dynamic Logical Partitioning	<ul style="list-style-type: none"> ▶ Processor, memory, and I/O resources can be dynamically moved between partitions. ▶ Partitions can use dedicated or shared (capped or uncapped) processor resources. ▶ Processor resources can automatically move between partitions based on workload demands.
Multiple Shared Processor Pools	Processor resources for a group of partitions can be capped helping to reduce software license requirements.
Virtual I/O Server (VIOS)	Virtualizes I/O resources (disk, tape, Ethernet) for client partitions, simplifying management and reducing system costs.
Integrated Virtualization Manager (IVM)	Simplifies partition creation and management for entry Power servers.
PowerVM Lx86	Supports running many x86 Linux applications in a Linux on Power partition.
Live Partition Mobility	Running AIX and Linux partitions can be moved to another system, helping to avoid planned downtime.
N_Port ID Virtualization (NPIV)	Simplifies the management and improves performance of Fibre Channel SAN environments—a fiber channel industry standard method for virtualizing a physical fiber channel port (for example, to multiple partitions).
System Planning Tool	Simplifies the planning for and installation of Power servers with PowerVM. Places hardware adapters according to desired partition configuration, resulting file input to the IBM configurator.

Note: Multiple Shared Processor Pools, Live Partition Mobility and NPIV are only supported on selected POWER 6 models.

Not all of these features are relevant in many IBM i environments. The following section provides more details about the more important features from an IBM i perspective.

7.1.1 Hypervisor

The POWER Hypervisor™ is the foundation for virtualization on POWER systems, enabling the hardware to be divided into multiple partitions, and ensuring isolation between them.

Always active on POWER servers, the POWER Hypervisor is responsible for dispatching the logical partition workload across the shared physical processors. The POWER Hypervisor also enforces partition security, and provides inter-partition communication that enables the Virtual I/O Server's virtual SCSI and virtual Ethernet function. Figure 7-1 shows a sample Hypervisor scenario.




Figure 7-1 Hypervisor

7.1.2 Shared processor pools

Shared-processor partition capability was introduced with IBM POWER5-based servers.

These shared-processor partitions, known as micro partitions, are allocated processor resources from a single pool of physical processors.

The amount of processor capacity that is allocated to a micro partition (its entitled capacity) may range from ten percent (10%) of a physical processor up to the entire capacity of the physical shared processor pool. The physical shared processor pool can have from one physical processor up to the total installed processor capacity of the system.

7.1.3 Multiple Shared-Processor Pools

On POWER6-based systems there is support for Multiple Shared-Processor Pools. This new capability allows a systems administrator to create a set of micro partitions with the purpose of controlling the processor capacity that they can consume from the physical shared-processor pool.

Multiple Shared-Processor Pools use processor capacity from the physical Shared Processor Pool—there can only be one physical shared processor pool in the system. Each Shared-Processor Pool has an associated Entitled Pool Capacity which is consumed by the set of micro partitions in that Shared-Processor Pool.

The micro partitions in a Shared-Processor Pool are guaranteed to receive their entitled capacity. In addition, unused processor cycles in their Shared-Processor Pool are harvested and are redistributed to eligible micro partitions within the same Shared-Processor Pool.

The source of additional processor cycles can be the Reserved Pool Capacity, which is processor capacity specifically reserved for a Shared-Processor Pool but not assigned to any of the micro partitions in the Shared-Processor Pool.

When the set of micro partitions in a Shared-Processor Pool are heavily loaded, they can consume additional processor capacity (assuming they are uncapped) from outside their Shared-Processor Pool up to a defined maximum. Processor capacity distributed in this way has been ceded by underutilized sets of micro partitions in their Shared-Processor Pools.

There is a default Shared-Processor Pool that is used for all micro partitions if the system administrator has not created other Shared-Processor Pools. The default Shared-Processor Pool has some special properties that ensure operational consistency with POWER5-based micro-partitioning and the physical shared processor pool. Figure 7-2 shows an example of Multiple Shared-Processor Pools.

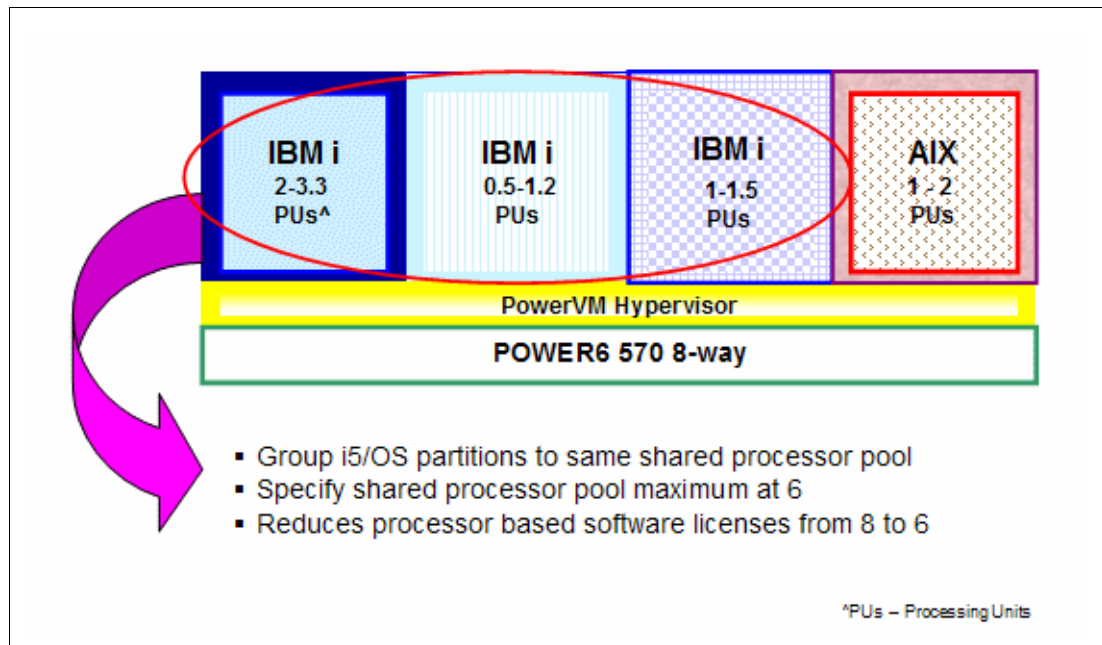


Figure 7-2 Multiple Shared Processor Pools example

Using Multiple Shared Processor Pools is intended to control processor utilization of “related uncapped partitions” by grouping them to use a common shared pool with a maximum processor unit capacity for all associated partitions. This results in control of maximum OS license requirements when all partitions are attempting to consume maximum processor capacity concurrently.

Additional details on Multiple Shared Processor Pool setup and usage include:

- ▶ The default shared processor pool (the only shared processor pool previous to multiple shared processor pool support availability).
- ▶ 63 additional shared processor pools are available for use.
- ▶ You must configure the shared processor pool attributes (defaults provided) before an additional shared processor pool can be used.

Additional shared processor pool attributes include:

- Name
Must be a unique name on the system.

- Identifier
Unique number that identifies the pool.
- Reserved processors
Whole number. A logical partition can use reserved processing units only when the logical partition is an uncapped logical partition that requires more than its assigned amount of processor resources. Reserved processing units cannot be used by logical partitions in other shared processor pools. Also, reserved processing units cannot be assigned to a specific logical partition.
- Maximum processors
Whole number. Equal to the maximum number of processing units that the logical partitions in each shared processor pool can use.

Note: IBM System p and AIX announced support of multiple shared processor pools during 4Q 2007. This support is provided by the POWER6 Hypervisor and associated firmware. Thus the capabilities are available to AIX 5.3 or later, the Linux release levels supported on POWER6, and IBM i 5.4 with LIC 5.4.5 or later.

Figure 7-3 shows how to access the multiple shared pool support from the HMC. You must select the tasks pull-down menu while displaying the system itself, as shown, to see the Shared Processor Pool Management option.

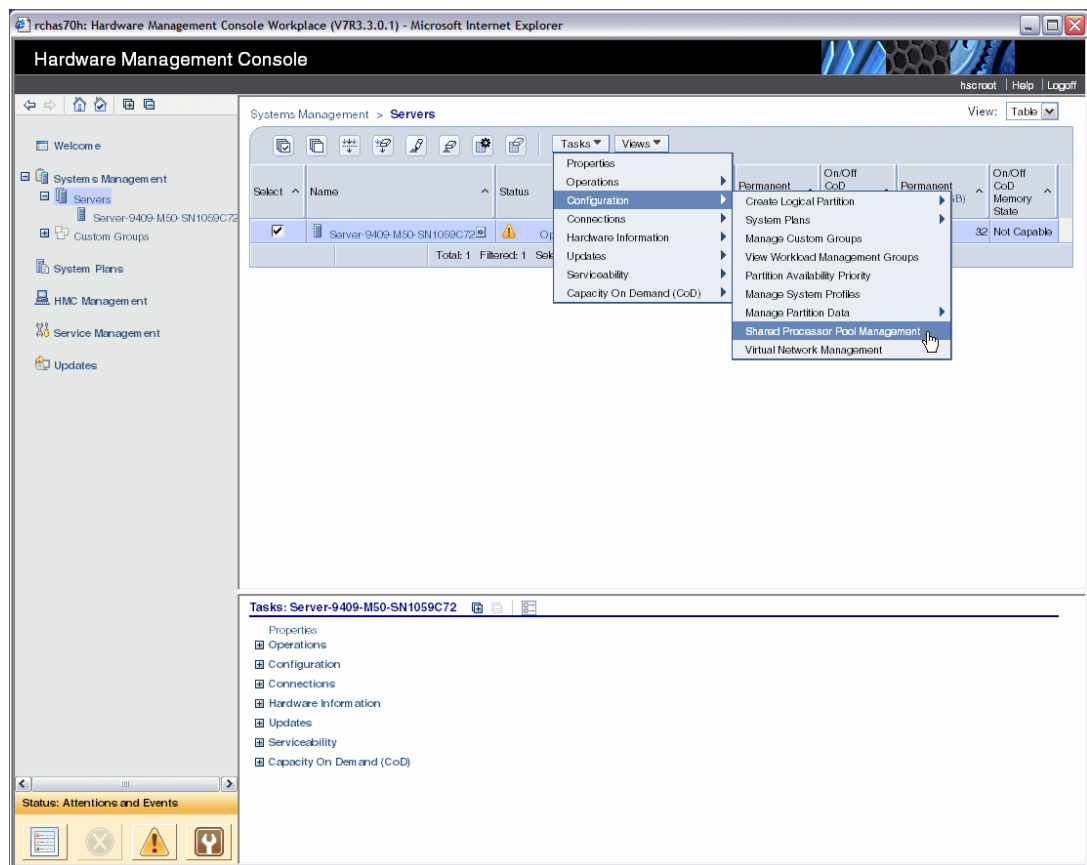


Figure 7-3 Accessing the multiple shared pool support from the HMC

Selecting **Shared Processor Pool Management** you get the list of shared processor pools as shown in the upper left of Figure 7-4. The 63 additional shared processor pools have pre-assigned names which you can change.

You must set up a shared processor pool before you assign existing partitions to use the pool. Figure 7-4 shows the list of shared pools. Select shared pool SharedPool63 (not shown in upper left window). We change its pool name to ITSOTEST, specify its processor capacities and click **OK** to get the lower left list of shared processor pools redisplayed, showing pool ITSOTEST.

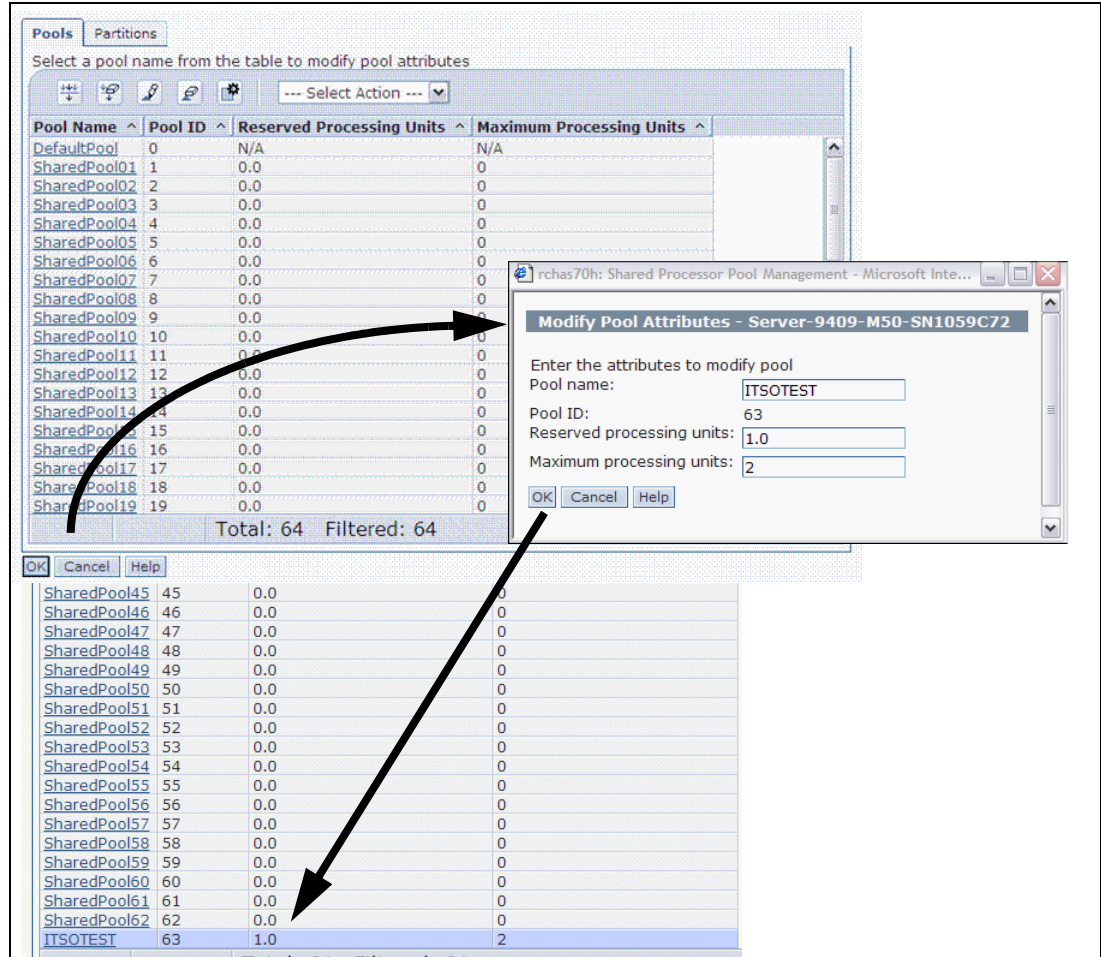


Figure 7-4 Example: Selecting, renaming, configuring shared processor pool 63

As shown in Figure 7-5 on page 201, we selected the Partitions tab showing existing partitions ucc42(5) and ucc46(1) as currently sharing the default shared pool. In the lower right window we show assigning shared pool ITSOTEST to partition ucc42. In our simple example, we could similarly assign shared processor pool ITSOTEST to partition ucc46.

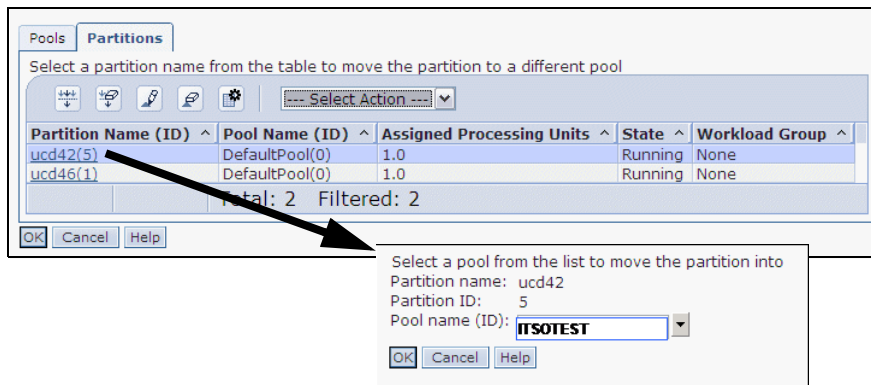


Figure 7-5 Example: Assigning a partition to shared processor pool 63 (ITSOTEST)

For more information on shared processor pools, go to the IBM Systems Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp>

and search with “shared processor pools.”

7.1.4 Dedicated processor sharing

POWER6 systems support the option to share the dedicated processors of a partition when it is inactive. When the partition normally running with dedicated processors is shut down, you specify that its processor capacity is available to shared processor partitions. If you allow the dedicated processors to be shared, they are still committed to this logical partition while they are available to shared processor partitions.

Dedicated processors can be shared only if they are idle. They are still committed to the active dedicated processor partition, even if you allow the dedicated processors to be shared. The dedicated processor partition can reclaim its dedicated processors at any time it needs them.

Changing this setting can be done on the partition profile settings to the logical partition by activating the logical partition using this partition profile. These changes can also be performed directly by changing the logical partition properties. Direct changes to the logical partition properties take effect immediately.

7.1.5 Dynamic logical partitioning

The POWER Hypervisor allows you to define logical partitions, which can be a subset of the overall system resources. It is possible to dynamically move system resources between partitions without rebooting. These resources can be:

- ▶ Physical processors (for dedicated processor LPARs)
- ▶ Virtual processors or entitled capacity (for micro-partitions)
- ▶ Memory
- ▶ Physical I/O adapters
- ▶ Virtual I/O adapters

This is known as dynamic reconfiguration or dynamic LPAR.

7.1.6 Virtual I/O Server

As part of PowerVM there is an appliance server with which you can associate physical resources and which allows you to share these resources amongst a group of logical partitions. The Virtual I/O Server (Figure 7-6 on page 203) can use both virtualized storage and network adapters, making use of the virtual SCSI and virtual Ethernet facilities.

For storage virtualization, these back-end devices can be used:

- ▶ Direct-attached entire disks from the Virtual I/O Server
- ▶ SAN disks attached to the Virtual I/O Server
- ▶ Logical volumes defined on either of the previous disks
- ▶ File-backed storage, with the files residing on either of the first two disks
- ▶ Optical storage devices

For virtual Ethernet we can define Shared Ethernet Adapters (SEA) on the Virtual I/O Server, bridging network traffic from the virtual Ethernet networks out to physical Ethernet networks.

The Virtual I/O Server technology facilitates the consolidation of LAN and disk I/O resources and minimizes the number of physical adapters that are required, while meeting the non-functional requirements of the server.

The Virtual I/O Server can run in either a dedicated processor partition or a micro partition.

Note: IBM i 6.1, as a hosting partition also provides various virtual I/O server hosting features (virtual SCSI, Ethernet bridging, and so forth). Its range of virtualization capabilities is not as extensive as that provided by VIOS, but they do enable IBM i products and capabilities to be used that are not available under VIOS.

For example, IBM i does not support virtualization of Fibre Channel adapters (NPIV) or direct connectivity to IBM Storage Solution products such as the DS4000® class of storage servers, or connectivity through a SAN Virtual Controller (SVC). However, tools for backup, such as the Backup, Recovery and Media Services (BRMS) product, 5761-BR1, can be used to back up IBM i objects, including network server storage objects used by client AIX or Linux partitions and integrated Windows or Linux operating systems on integrated System x servers.

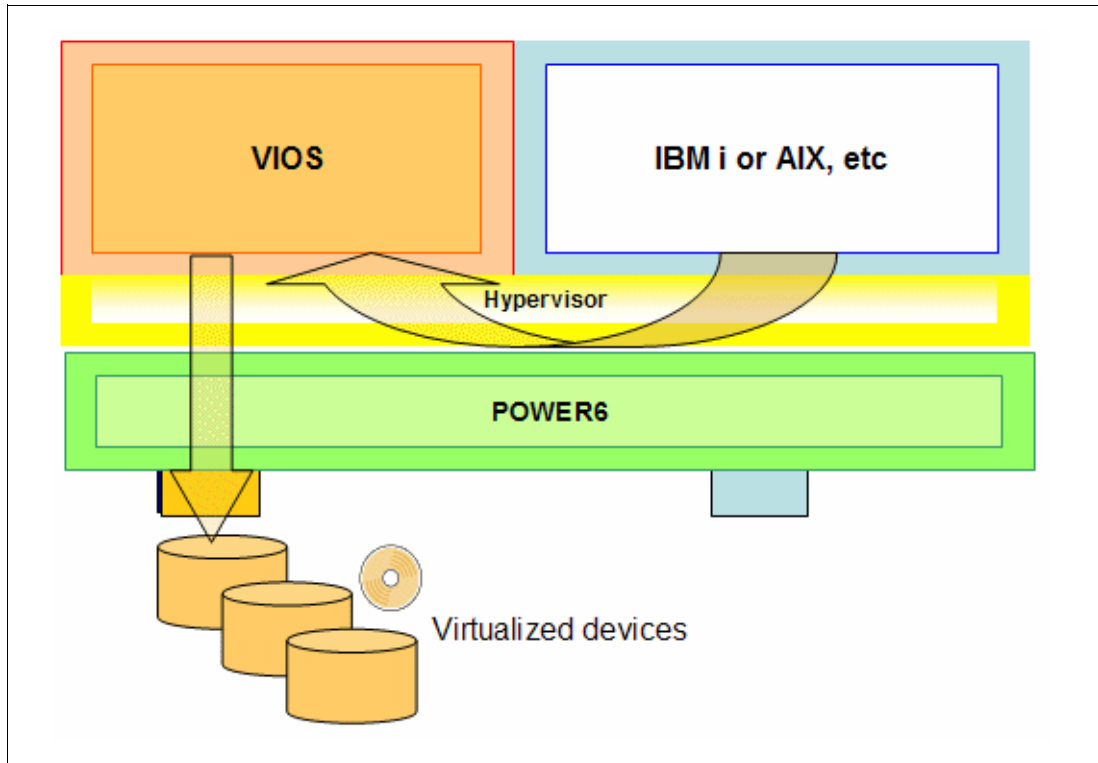


Figure 7-6 Virtual I/O Server (VIOS)

7.1.7 Integrated Virtualization Manager

The Integrated Virtualization Manager (IVM) is a Web browser-based management interface that provides support for certain models of POWER5 and POWER6 systems. The IVM is hosted by a Virtual I/O Server partition.

The HMC has been designed to be the comprehensive solution for hardware management that can be used either for a small configuration or for a multiserver environment. Although complexity has been kept low by design and many recent software revisions support this, the HMC solution might not fit in small and simple environments where only a few servers are deployed or not all HMC functions are required.

There are many environments where there is a need for small partitioned systems, either for test reasons or for specific requirements, for which the HMC solution is not ideal. An example is where there are small partitioned systems that cannot share a common HMC because they are in multiple locations and network connectivity using the HMC is not the appropriate solution.

IVM is a simplified hardware management solution that inherits most of the HMC features. It manages a single server, avoiding the need for an independent personal computer. It is designed to provide a solution that enables an administrator to reduce system setup and to make hardware management easier, at a lower cost.

IVM provides a management model for a single system. Although it does not offer all of the HMC capabilities, it enables the exploitation of IBM Virtualization Engine technology. IVM is targeted for use with the small and medium systems that are best suited for this product. Figure 7-7 on page 204 shows a scenario of IVM and HMC - Partition creation.

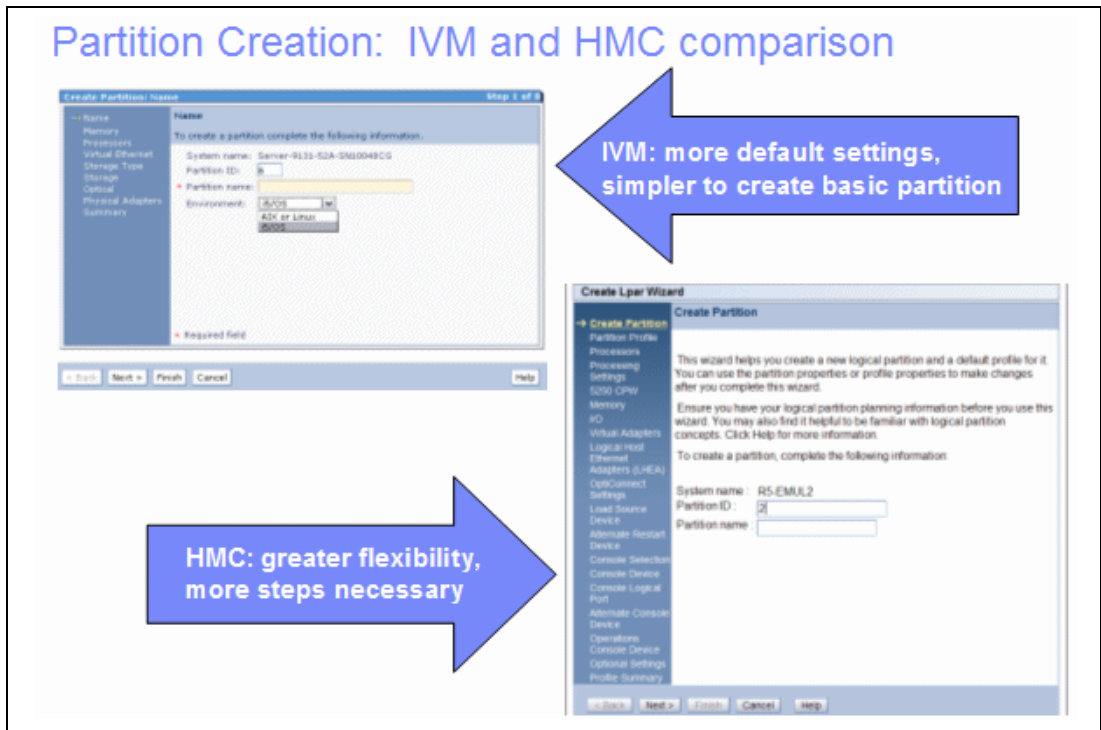


Figure 7-7 IVM and HMC - Partition creation

For a more detailed comparison of IVM and HMC, refer to 20.1, “IVM and HMC comparison at a glance” on page 474.

IVM is an enhancement of the Virtual I/O Server (VIOS), the product that enables I/O virtualization in POWER5 and POWER6 processor-based systems. It enables management of VIOS functions and uses a Web-based graphical interface that enables an administrator to remotely manage the server with a browser. The HTTPS protocol and server login with password authentication provide the security required by many enterprises.

Because one of the goals of IVM is simplification of management, some implicit rules apply to configuration and setup:

- ▶ When a system is designated to be managed by IVM, it must not already be partitioned.
- ▶ The first operating system to be installed must be the VIOS.

The VIOS is automatically configured to own all of the I/O resources and it can be configured to provide service to other LPARs through its virtualization capabilities. Therefore, all other logical partitions (LPARs) do not own any physical adapters and they must access disk, network, and optical devices only through the VIOS as virtual devices. Otherwise, the LPARs operate as they have previously with respect to processor and memory resources.

The system Hypervisor has been modified to enable the VIOS to manage the partitioned system without an HMC. The software that is normally running on the HMC has been rewritten to fit inside the VIOS and to provide a simpler user interface. Because the IVM is running using system resources, the design has been developed to have minimal impact on disk, memory, and processor resources.

The IVM does not interact with a system’s service processor. A specific device named the Virtual Management Channel (VMC) has been developed on the VIOS to enable a direct

Hypervisor configuration without requiring additional network connections. This device is activated, by default, when the VIOS is installed as the first partition.

The VMC enables IVM to provide basic logical partitioning functions:

- ▶ Logical partitioning configuration
- ▶ Boot, start, and stop actions for individual partitions
- ▶ Display of partition status
- ▶ Management of virtual Ethernet
- ▶ Management of virtual storage
- ▶ Basic system management

Because IVM executes on an LPAR, it has limited service-based functions, and ASMI must be used. For example, a server power-on must be performed by physically pushing the server power-on button or remotely accessing Advanced System Management Interface (ASMI) because IVM does not execute while the server power is off. ASMI and IVM together provide a simple but effective solution for a single partitioned server.

LPAR management using IVM is through a common Web interface developed for basic administration tasks. Being integrated within the VIOS code, IVM also handles all virtualization tasks that normally require VIOS commands to be run.

IVM and HMC are two unique management systems: The IVM is designed as an integrated solution designed to lower your cost of ownership, and the HMC is designed for flexibility and a comprehensive set of functions. This provides you the freedom to select the one ideal solution for your production workload requirements.

Note: The internal design of IVM requires that no HMC should be connected to a working IVM system. If a client wants to migrate an environment from IVM to HMC, the configuration setup has to be rebuilt manually. This includes systems that had previous software levels of VIOS running on them, because they would also have been managed by an HMC.

Figure 7-8 depicts a sample IVM page. On the right View/Modify Partitions pane, you can see total system processor and memory maximums, any pending hardware event status, your currently defined partition configuration summaries and select buttons to create, start (activate) and stop (shut down) existing partitions.

On the left navigation pane you see the major groupings of management functions:

- ▶ Partition management
- ▶ I/O adapter management (Host Ethernet Adapters (“Integrated Virtual Ethernet” adapters) Virtual Ethernet, and Physical Adapters)
- ▶ Virtual storage management (the external storage configuration)
- ▶ Management of IVM itself
- ▶ System Plan (for automated LPAR configuration ordering and deployment)
- ▶ Service management (problem determination, firmware updates, logs, backup of the VIOS, partition configuration information, and others.)

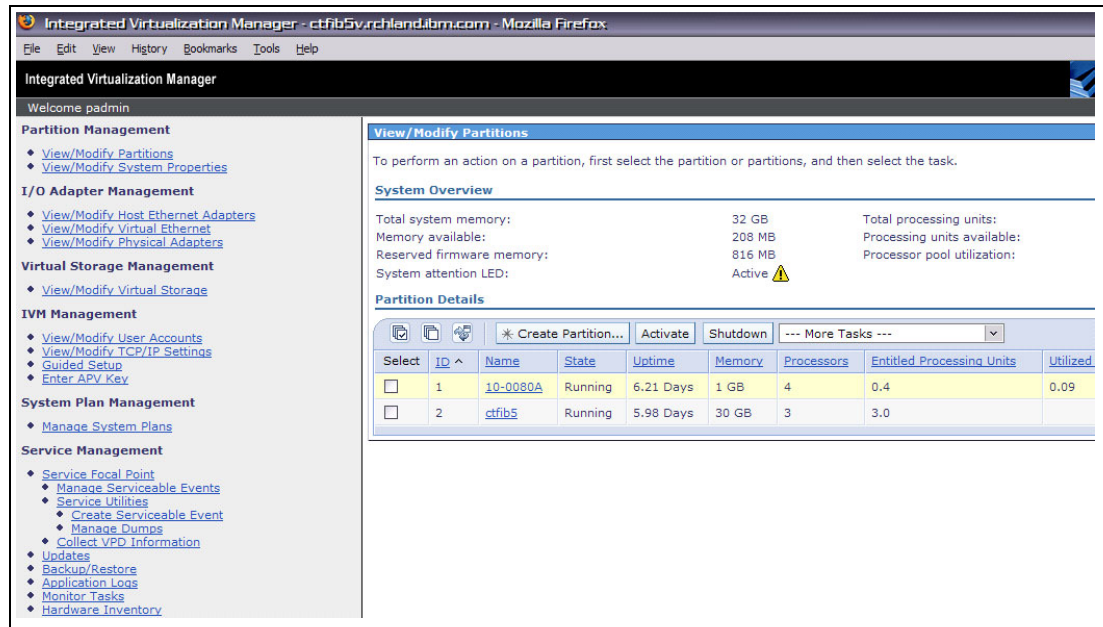


Figure 7-8 Integrated Virtualization Manager

7.2 IBM i 6.1 host and client partitions

In this section we provide an example of creating an IBM i hosted partition with an HMC.

We start with a description of the different types of partitions under virtualization:

- ▶ Host/hosting
- ▶ Hosted

The hosting (sometimes called “server”) partition owns physical resources and facilitates the sharing of the physical resources among the hosted (often called “client”) partitions.

In IBM i 6.1, client partitions can be hosted by a VIOS partition or another IBM i 6.1 partition, as shown in Figure 7-9.

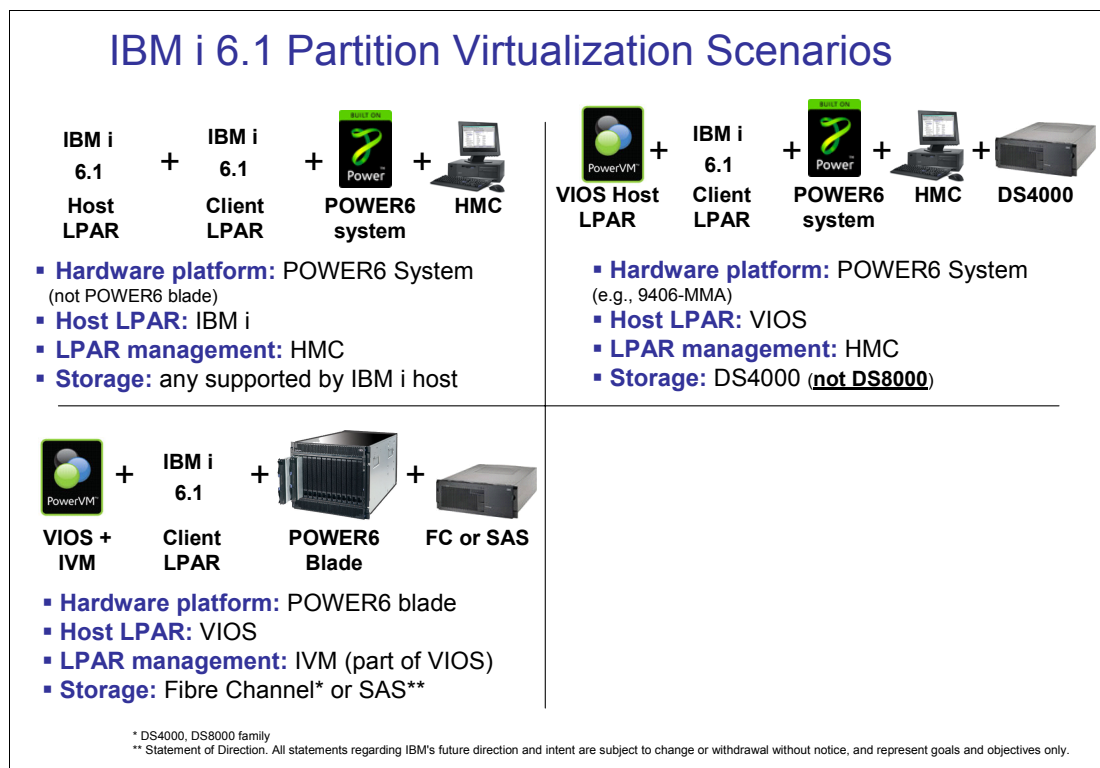


Figure 7-9 IBM i 6.1 Partition Virtualization scenarios

Note: To use IVM, the partition that owns the physical resources must be running VIOS.

To start the configuration of the client, we launch the Create Lpar Wizard on the HMC as shown in Figure 7-10. In the left navigation pane, select the option to create an i5/OS Logical Partition. That will launch the wizard as displayed in the right side of the figure.

On the Create Partition window, enter the Partition ID and name fields and select **Next**.

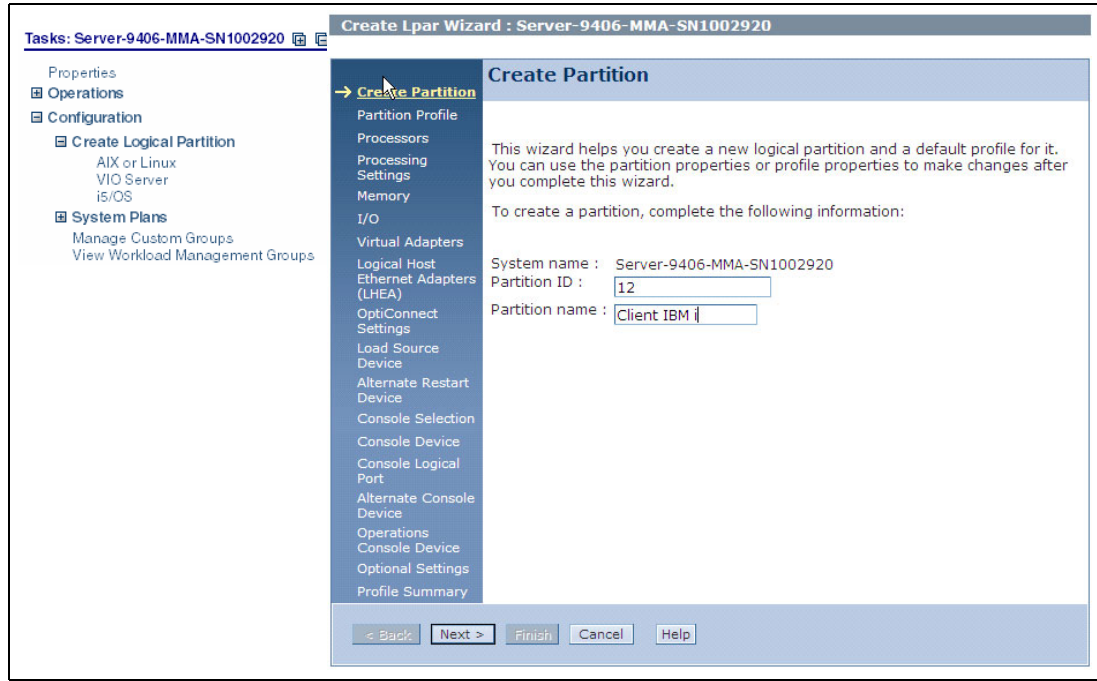


Figure 7-10 Create a partition

As shown in Figure 7-11 on page 209, provide a Profile Name on the Create Partition profile page, but do not check “Use all the resources” in the system checkbox. In our example, the hosting partition is the partition that owns the resources, not our new client partition.

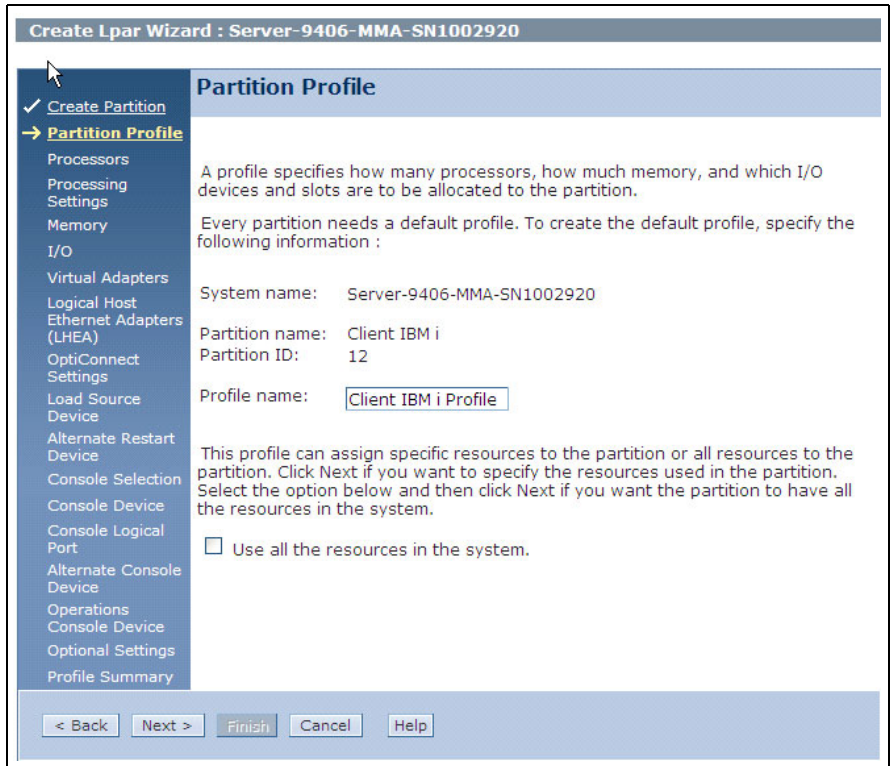


Figure 7-11 Partition profile

Next we define what the processor usage type will be, as shown in Figure 7-12. In our example, we will use shared processors.

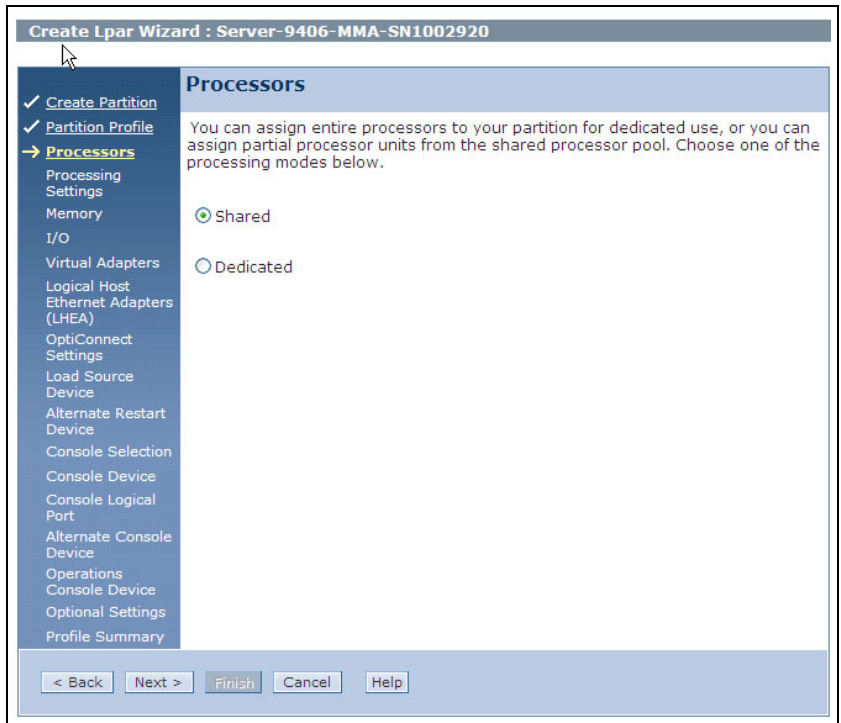


Figure 7-12 Processors

Now that we have specified that we are using shared processors, we need to tell the wizard how we would like the sharing set up. In the Processor Settings page we can enter the processing units, virtual processors, and whether or not the partition is uncapped, as shown in Figure 7-13.

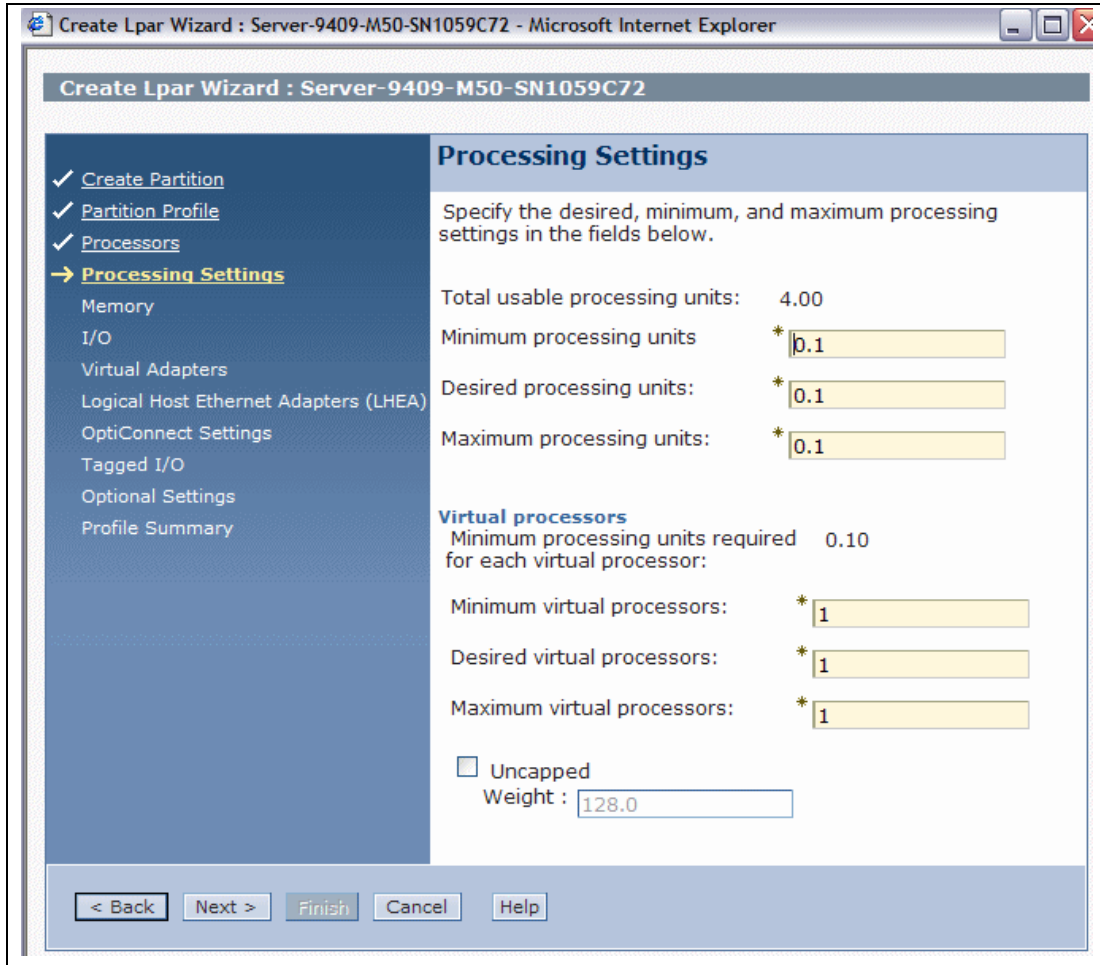


Figure 7-13 Processing settings

Once the processing settings are complete, we move on to the memory settings shown in Figure 7-14 on page 211.

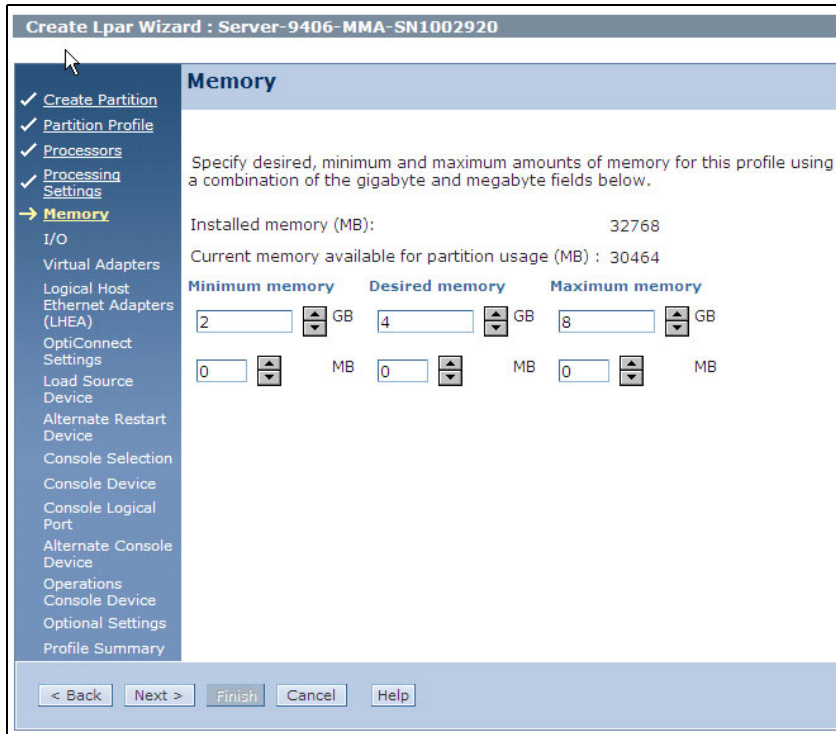


Figure 7-14 Memory

In the next Web page, I/O specifications, we can nominate which, if any, physical resources we want to assign to the new client partition. In our example, we will not be assigning any physical I/O resources—they will all be virtual. This means that all virtual I/O resources that we want to assign to our new client partition have already been assigned as physical resources in the host partition.

We do not select anything on this page other than **Next**. Figure 7-15 shows a list of physical resources which will not be selected in our example.

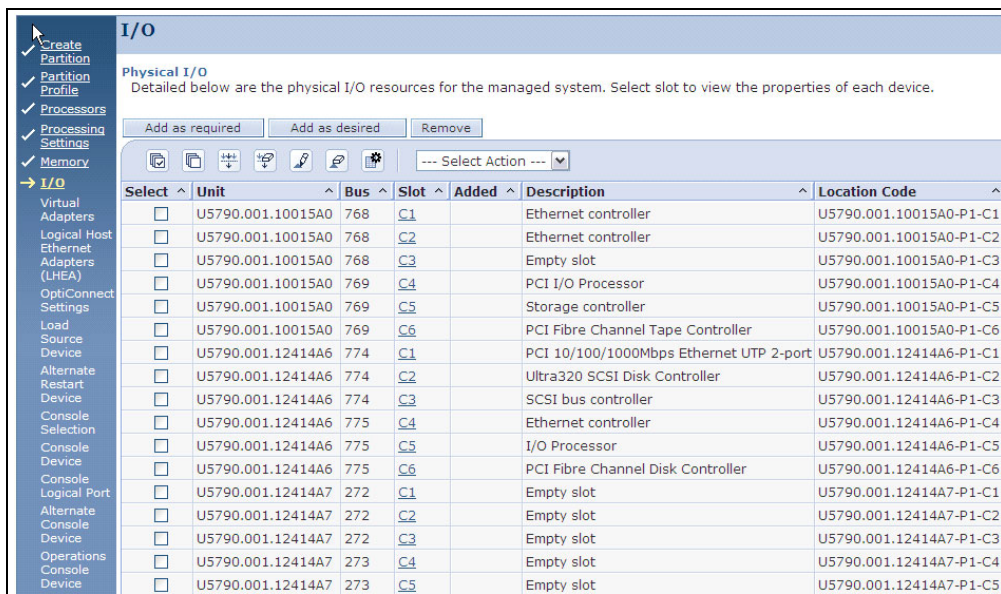


Figure 7-15 I/O

Next we create virtual adapters. These can be Ethernet, SCSI or Serial. In our case, we want a **SCSI Adapter**, as shown in Figure 7-16.

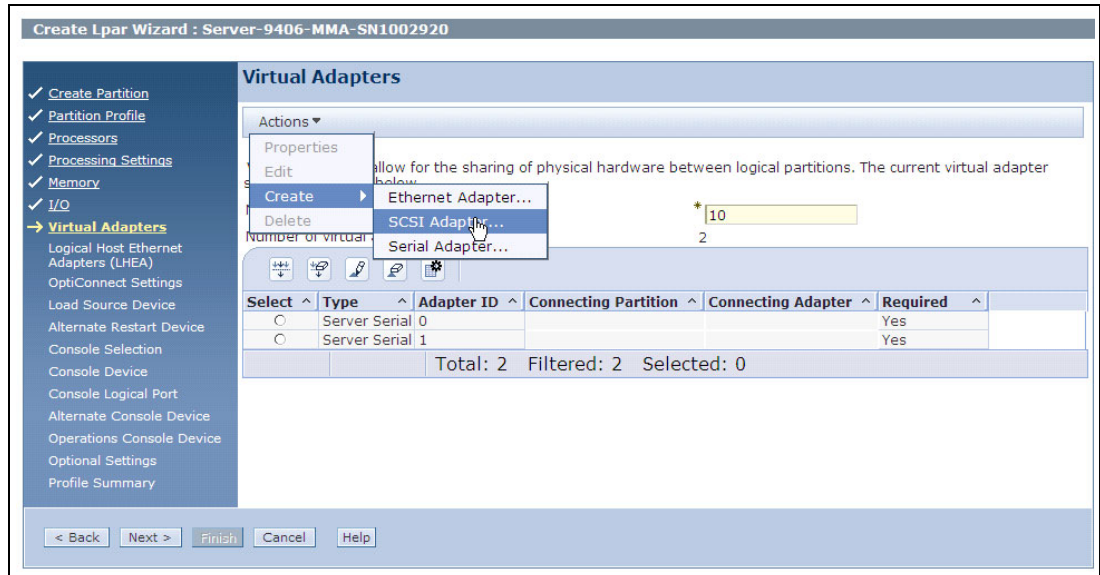


Figure 7-16 Virtual adapters

When adding a virtual adapter, we need to tell the wizard whether we want a server-side or client-side adapter. In our case, we are creating a client partition, so we select **Client** as shown in Figure 7-17.

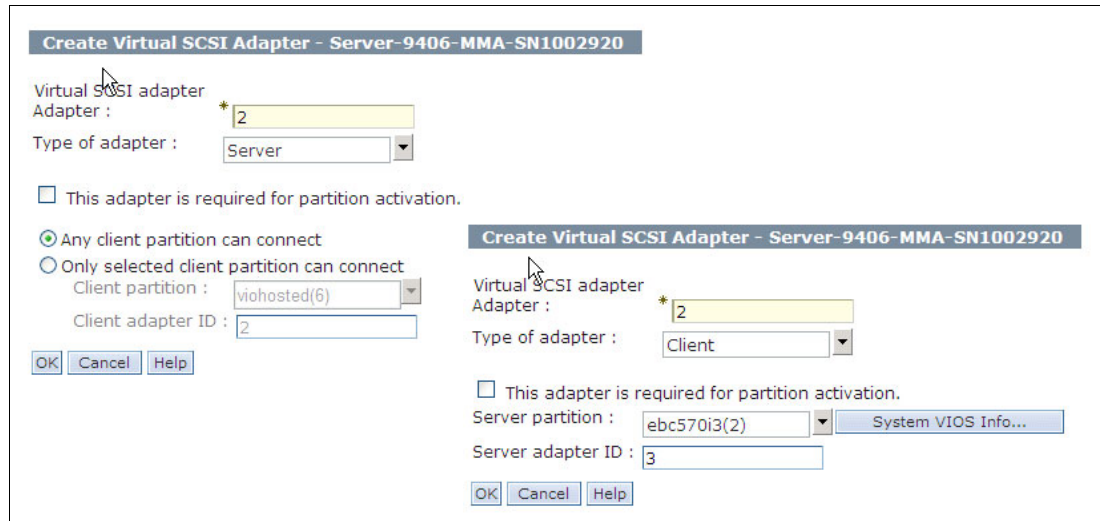


Figure 7-17 Virtual SCSI adapter

Next, we can continue creating any other virtual adapters we need, as shown in Figure 7-18 on page 213.

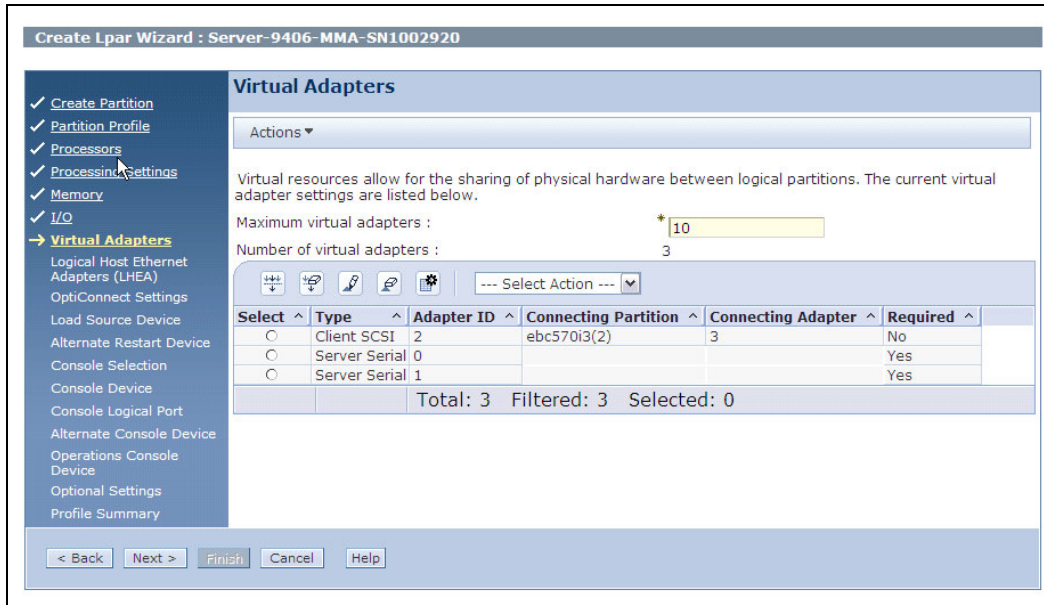


Figure 7-18 Other virtual adapters

Just as we do in physical partitions, we need to tag the load source device. In a client partition, we can use a physical adapter or a virtual adapter. In our scenario, we are only using virtual adapters, so we select the same in the drop-down list, as shown in Figure 7-19.

For further information about assigning a virtual adapter (including Open Storage) as IBM i load source device, refer to the *IBM i Virtualization and Open Storage Read-me First* document, available at:

http://www.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

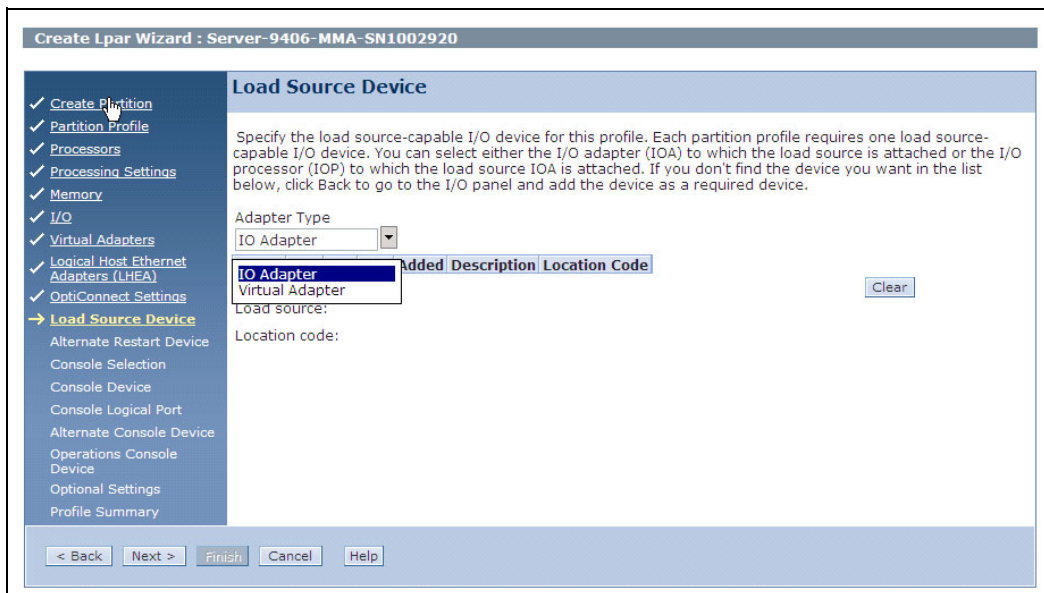


Figure 7-19 Load source device type

After selecting the virtual adapters, we are presented with a list of adapters to choose from. We only have one in our configuration, as shown in Figure 7-20.

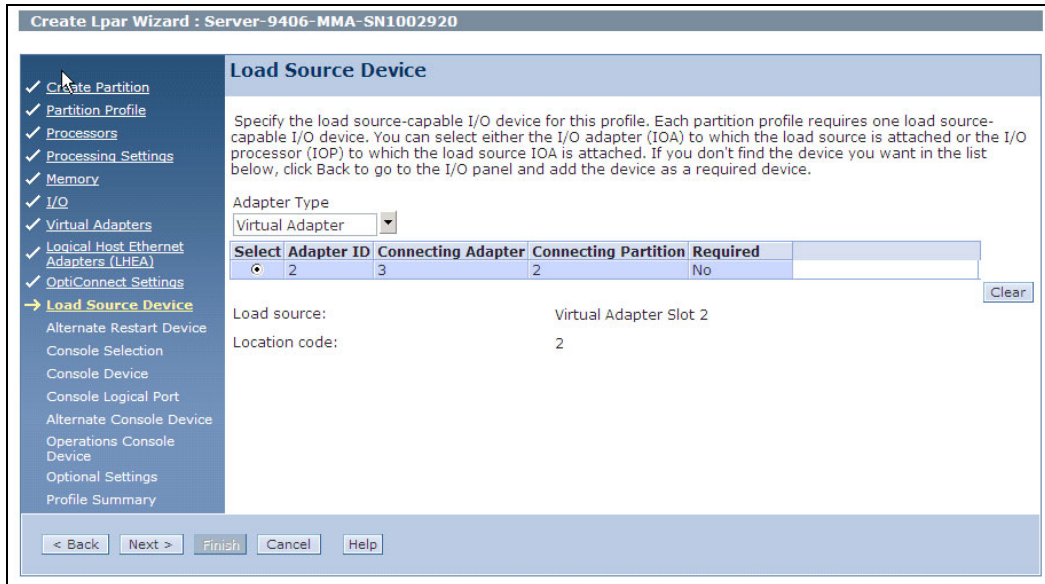


Figure 7-20 Load source device selection

After defining the load source device, we need to nominate an alternate IPL device. As with the load source device, we can select either a physical or a virtual adapter, as shown in Figure 7-21.

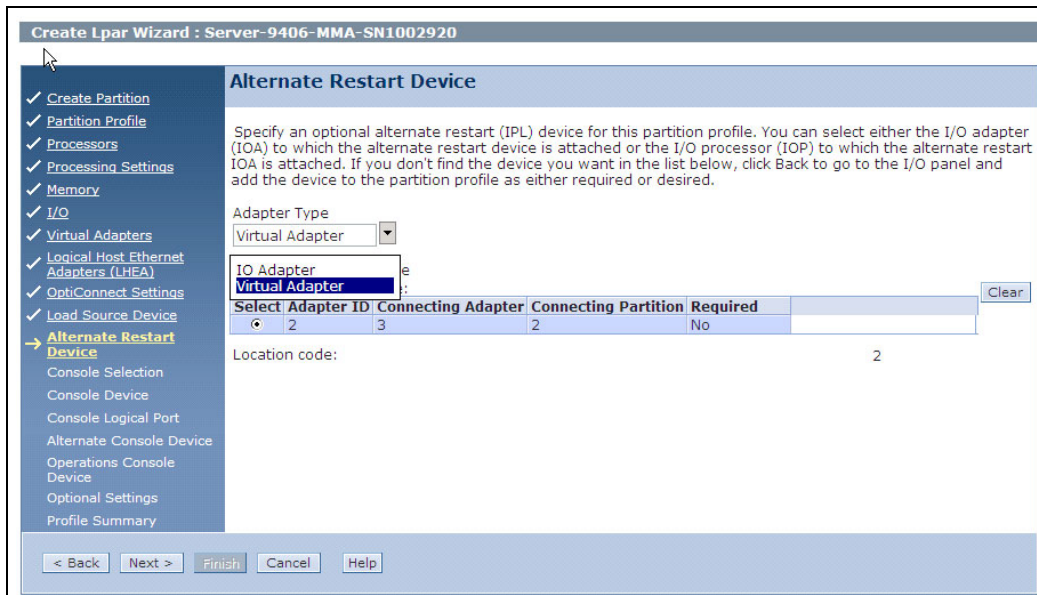


Figure 7-21 Alternate restart device

The last step in defining the configuration is selecting what type of console we wish to use as shown in Figure 7-22.

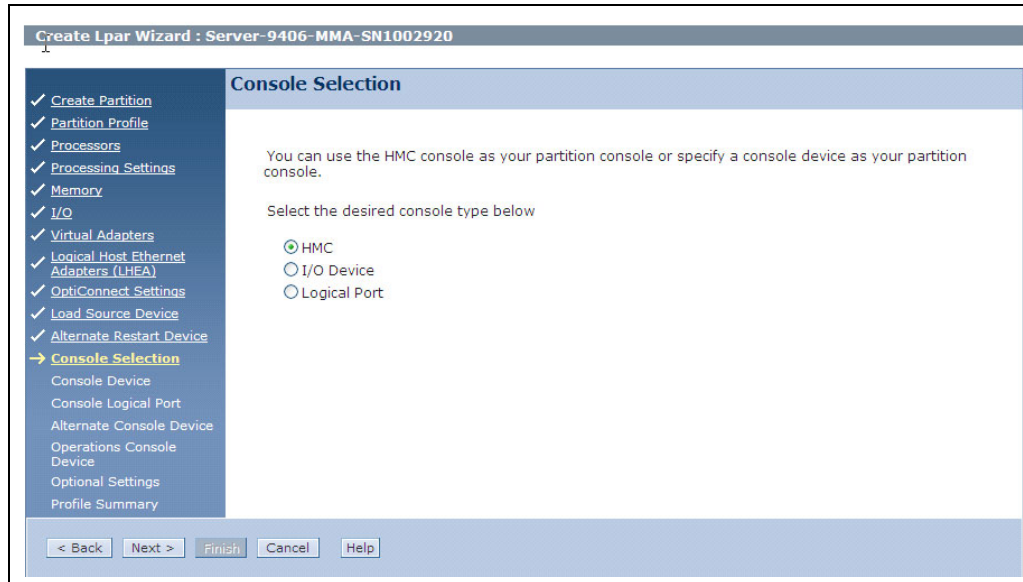


Figure 7-22 Console selection

After we have finished defining the configuration, we are presented with a summary of the profile. From this page, we can go back and modify it or, if it is what we wanted, select **Finish** to create the partition profile. See an example of a Profile summary in Figure 7-23.



Figure 7-23 Profile summary

The partition creation is now complete.

See the example of HMC Contents of All Partitions shown in Figure 7-24.

Sele...	Name	L...	Server	Status	Process... Units	Mem... (GB)	Active Pref...	Environm...	Referenc... Code
<input type="checkbox"/>	AIX 6.1	3	Server-9406-MMA-SN100...	Running	0.2	1.125	AIX61	AIX or Linux	
<input type="checkbox"/>	Client IBM i	12	Server-9406-MMA-SN100...	Not Activ...	0	0		i5OS	00000000
<input type="checkbox"/>	ebc570H	4	Server-9406-MMA-SN100...	Running	1	8	ebc570H	i5OS	00000000
<input type="checkbox"/>	ebc570E	5	Server-9406-MMA-SN100...	Running	0.5	4	ebc570E	i5OS	00000000
<input type="checkbox"/>	ebc570B	2	Server-9406-MMA-SN100...	Running	1	6	ebc570B	i5OS	00000000
<input type="checkbox"/>	ebcconf	9	Server-9406-MMA-SN100...	Not Activ...	0	0	ebcconf	AIX or Linux	00000000
<input type="checkbox"/>	ebcpri	7	Server-9406-MMA-SN100...	Running	0.1	1	ebcpri	AIX or Linux	Linux ppc6...
<input type="checkbox"/>	ebcscs	10	Server-9406-MMA-SN100...	Not Activ...	0	0	ebcscs	AIX or Linux	00000000
<input type="checkbox"/>	ebcsec	8	Server-9406-MMA-SN100...	Running	0.1	1	ebcsec	AIX or Linux	Linux ppc6...

Total: 13 Filtered: 13 Selected: 0

Figure 7-24 All Partitions

For further information about configuration and creation of host partitions and assigning resources, refer to *PowerVM Virtualization on IBM System p: Introduction and Configuration Fourth Edition*, SG24-7940, also available at:

<http://www.redbooks.ibm.com/abstracts/sg247940.html>

For further information about logical partitioning on Power Systems, refer to the *Power Systems Logical partitioning* document available at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>



Networking

This chapter discusses the following topics:

- ▶ “TCP/IP” on page 218
- ▶ “Virtual IP enhancements” on page 222
- ▶ “Open Shortest Path First routing protocol” on page 223
- ▶ “Domain Name System” on page 224
- ▶ “Netserver” on page 226
- ▶ “Virtual Private Network” on page 227
- ▶ “FTP client and server extended FTP commands” on page 227
- ▶ “Mail support enhancements” on page 228
- ▶ “IBM Tivoli Directory Server for i5/OS” on page 230
- ▶ “Advanced program-to-program communications” on page 233

8.1 TCP/IP

The following sections summarize the key TCP/IP enhancements in IBM i 6.1.

8.1.1 IPv6 support enhancements

In IBM i 5.4 TCP/IP support was enhanced to support IPv6, essentially at the user application sockets level interface. That is, you could write your own applications interfacing to TCP/IP IPv6 support at the API/sockets level.

Starting with IBM i 6.1, almost all IBM applications support runs over either IPv4 or IPv6. This includes:

- ▶ HTTP Server for i
- ▶ Telnet
- ▶ File Transfer (FTP)
- ▶ System i Access for Windows
- ▶ System i Navigator
- ▶ IBM Systems Director Navigator for i

TCP/IP setup also expands IPv6 support to include:

- ▶ Virtual IPv6 address
- ▶ Host table
- ▶ Domain Name System server
- ▶ Intrusion detection support

An overview of IBM i 6.1 communications support of IPv6 in IBM applications is provided in this and other chapters of this publication. For example, in this chapter under the topic 8.1.2, “TCP/IP configuration enhancements” on page 218, in 2.1, “Intrusion detection and prevention” on page 6, and in Chapter 17, “System i Access products” on page 429.

For a list of applications that were enabled, see:

<http://www-03.ibm.com/systems/i/software/tcpip/applications6.html>

8.1.2 TCP/IP configuration enhancements

The following functions for the TCP/IP configuration have been enhanced in this release:

- ▶ The EZ-Setup wizard for configuring TCP/IP was removed. You must use the 5250 command character-based interface to configure TCP/IP for the first time.
- ▶ To convert host names and their associated IPv6 addresses, you can configure and use a host table instead of a DNS server.
- ▶ You can configure IPv6 Stateless Address Auto-Configuration and start IPv6 interfaces when the operating system is in restricted state.
- ▶ You can create IPv4 or IPv6 addresses.
- ▶ If you have started TCP/IP without starting IPv6, you can start IPv6 at a later time without ending TCP/IP using the Start TCP/IP (STRTCP) command as shown in Figure 8-1 on page 219.

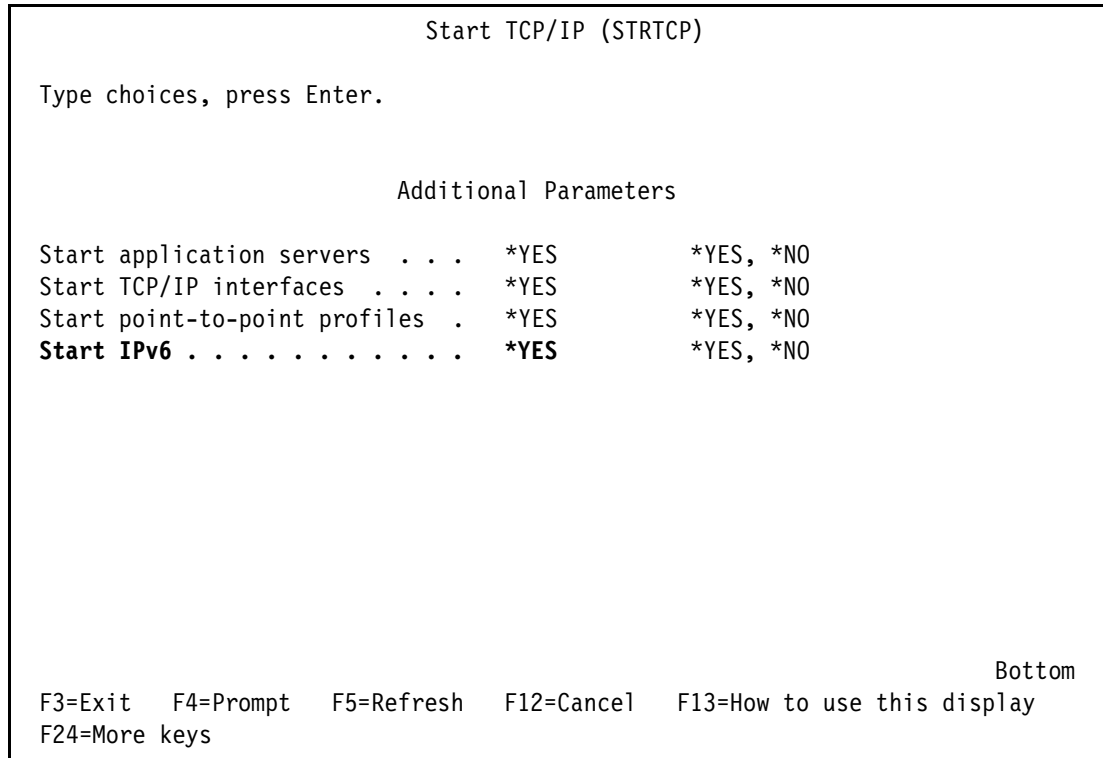


Figure 8-1 Start IPv6

8.1.3 IBM i command enhancements

You can use the IBM i command, in addition to System i Navigator, to configure and customize TCP/IP:

- ▶ Configuring IPv6 Stateless Address Auto-Configuration
- ▶ Manually adding, changing, and removing IPv4 and IPv6 interfaces
- ▶ Starting and stopping IPv4 and IPv6 interfaces
- ▶ Manually adding, changing, and removing IPv4 and IPv6 routes
- ▶ Starting and ending an IPv4 or IPv6 connection

8.1.4 System i Navigator enhancements

System i Navigator provides interfaces to the new 6.1 capabilities and now provides more consistent functions between IPv4 and IPv6, including:

- ▶ IPv6 stateless auto configuration interfaces are now listed in the IPv6 interfaces window. You can start and stop them using the pop-up menu.
- ▶ The preferred line descriptions selection is provided in the Options tab of the IPv6 Interface Properties window for Virtual IPv6 addresses.
- ▶ A new column called Link Status (on-link and off-link status) is displayed in the IPv6 interfaces window.
- ▶ You can change select properties of IPv4 and IPv6 interfaces when they are active. Others will not take effect until the interface is restarted.

- ▶ The Host Table window displays both IPv4 and IPv6 addresses. If you want to add, edit, or remove host names that are related to the same host table entry, you can complete these multiple tasks at the same time.
- ▶ The Neighbor Cache menu was moved from the navigation tree to the pop-up menu of an individual IPv6 interface or IPv6 line.
- ▶ The IPv4 and IPv6 attribute settings now use the same window that contains both IPv4 and IPv6 common property settings.

8.1.5 Systems Director Navigator for IBM i Network interfaces

Many of the System i Navigator functions are also replicated to the new IBM Systems Director Navigator for i Network set of tasks.

The general Systems Director Navigator for IBM i interface and navigation are described in Chapter 8, “Networking” on page 217.

To select a network task you first select **Network** on the i5/OS Management menu, as shown in Figure 8-2.

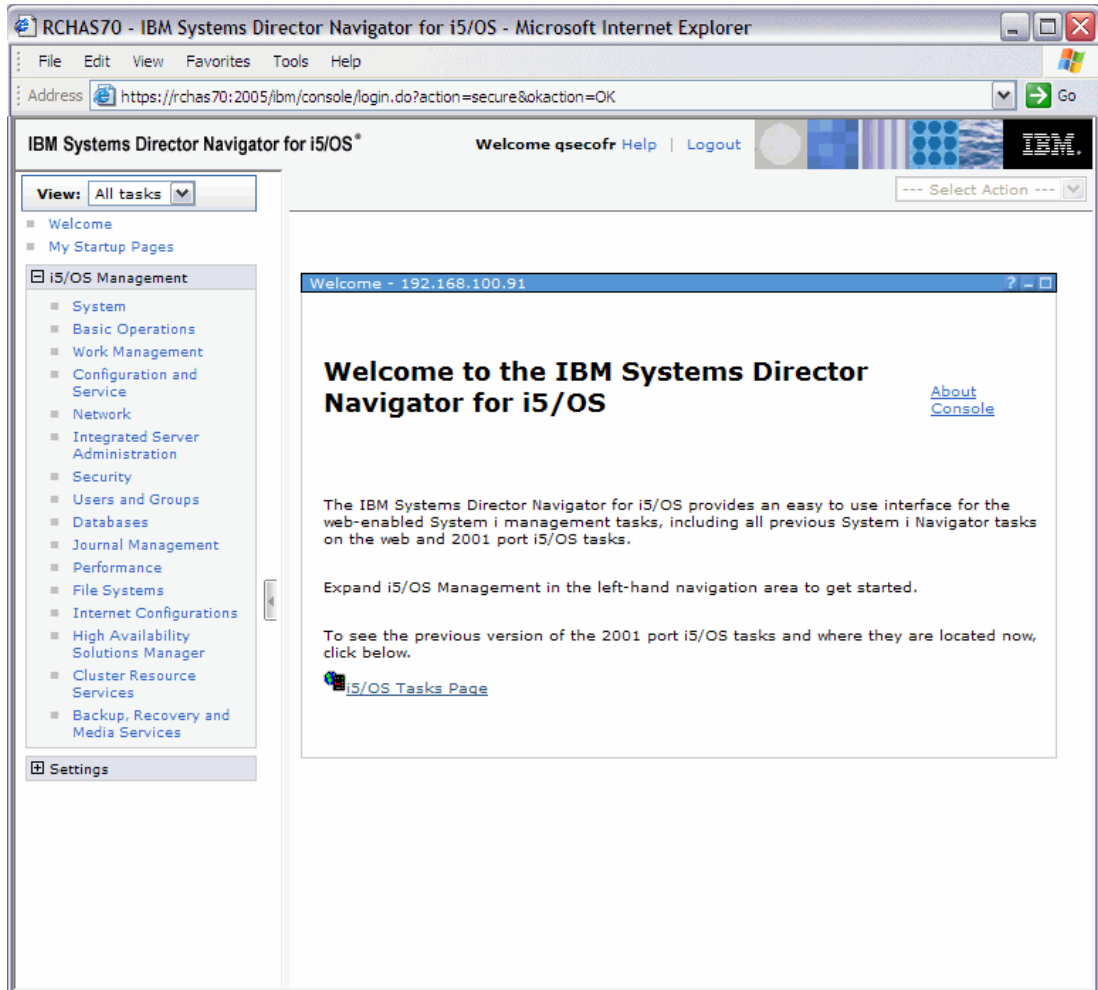


Figure 8-2 Network tasks option on the i5/OS Management menu

Next you click **Show All Network Tasks**, as shown in Figure 8-3.

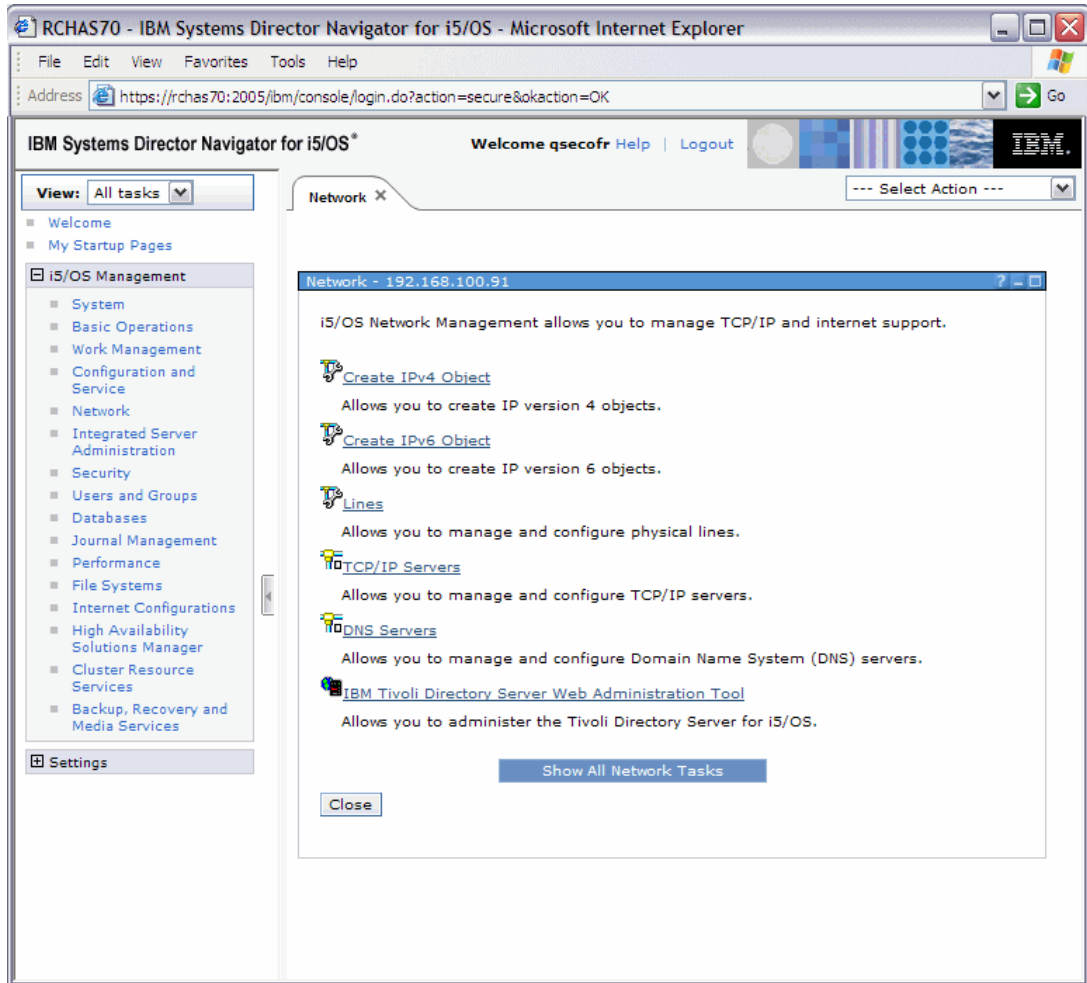


Figure 8-3 Show All Network Tasks

The next panel presented is a list containing all network tasks, as shown in Figure 8-4.

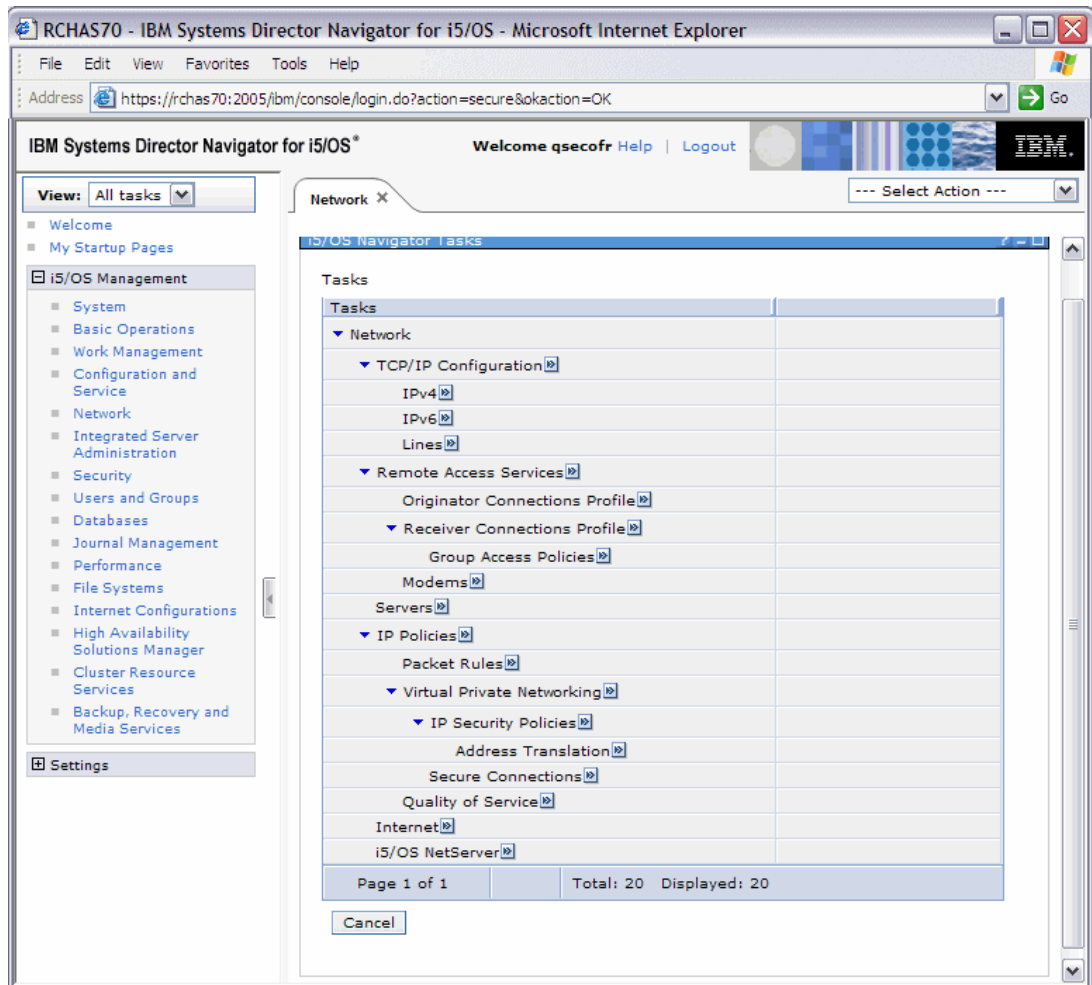


Figure 8-4 All network tasks

8.2 Virtual IP enhancements

Virtual IP enhancements that allow the TCP/IP routing and workload balancing capabilities are summarized here:

- ▶ Virtual IP address support has been extended to include IPv6 addresses.
- ▶ A Point-to-Point Protocol (PPP) interface or a Layer Two Tunneling Protocol (L2TP) interface can use a virtual IP address as the local IP address to provide fault tolerance for remote connections, as shown in Figure 8-5 on page 223.
- ▶ Because the virtual IP address is not bound to a single physical interface, the system never responds to an Address Resolution Protocol (ARP) request to the virtual IP address unless you enable proxy ARP for the virtual IP address. In other words, by enabling proxy ARP, a local interface can respond to the ARP requests on behalf of the virtual IP address. Otherwise, remote systems must have a route defined to reach the address. You can now configure virtual IP proxy ARP for a virtual IP interface while it is active.

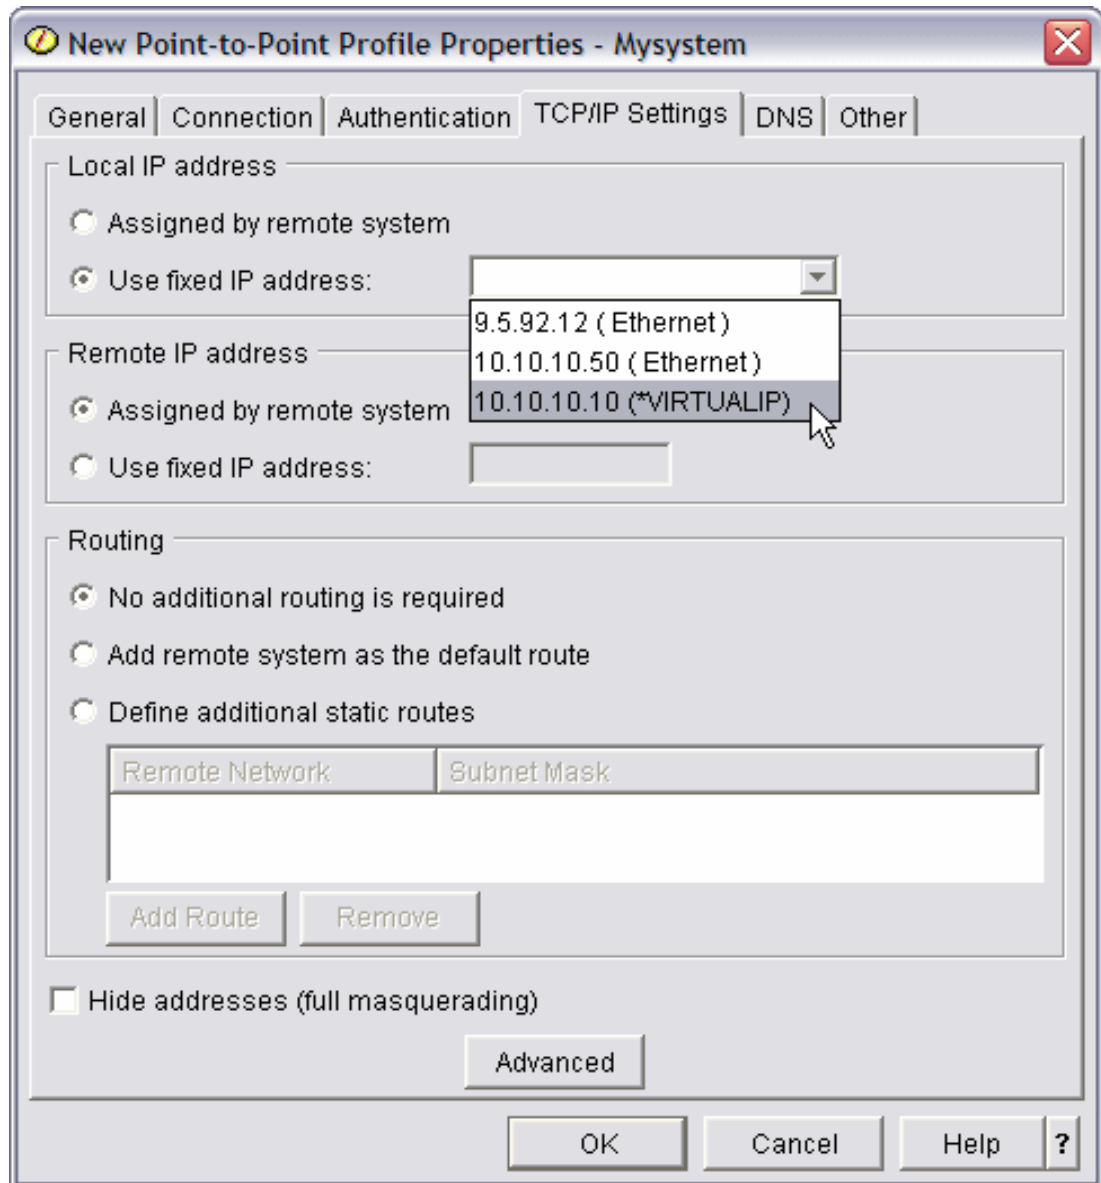


Figure 8-5 Point-to-Point Profile, TCP/IP settings

8.3 Open Shortest Path First routing protocol

Open Shortest Path First (OSPF) is a link-state routing protocol that was developed for IP networks and is based on the Shortest Path First (SPF) algorithm. OSPF is an Interior Gateway Protocol (IGP).

In an OSPF network, routers or systems within the same area maintain an identical link-state database that describes the topology of the area. Each router or system in the area generates its link-state database from the link-state advertisements (LSAs) that it receives from all the other routers or systems in the same area and the LSAs that it itself generates. An LSA is a packet that contains information about neighbors and path costs. Based on the link-state database, each router or system calculates a shortest-path spanning tree, with itself as the root, using the SPF algorithm.

Key advantages of OSPF

- ▶ Compared with distance-vector routing protocols such as the Routing Information Protocol (RIP), OSPF is more suitable for serving large, heterogeneous Internet networks. OSPF can quickly recalculate the routes when the network topology changes.
- ▶ With OSPF, you can divide an Autonomous System (AS) into areas and keep area topologies separate to decrease the OSPF routing traffic and the size of the link-state database of each area.
- ▶ OSPF provides equal-cost multipath routing. You can add duplicate routes to the TCP stack using different next hops.
- ▶ OSPF Hello protocol and link-state database exchange. After routers or systems in an OSPF network ensure that their interfaces are functional, they first send out Hello packets, using the Hello protocol over their OSPF interfaces, to discover neighbors. Neighbors are routers or systems that have interfaces to the common network. After that, neighboring routers or systems exchange their link-state databases to establish adjacencies.

The new IBM i Open Shortest Path First (OSPF) support includes a set of System i commands to configure your system as a participant in an OSPF networking environment.

For complete information about OSPF, refer to the IBM i Information Center at the following Web address and select the path **Networking** → **TCP/IP applications** → **Open Shortest Path First support**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

8.4 Domain Name System

8.4.1 Berkeley Internet Name Domain version 9

The Berkeley Internet Name Domain (BIND) version 9, introduced this release, provides several features to enhance the performance of your Domain Name System (DNS) server. For example, it supports name-to-address and address-to-name lookups in all currently defined forms of IPv6. It uses the view statement, which allows a single DNS instance to answer a query differently depending on where the query comes from, such as the Internet or an intranet. In addition, it uses journal files to hold dynamic updates for a zone.

The previous BIND 4.9.3 and BIND 8.2.5 are no longer supported, and need to be migrated to BIND 9.

8.4.2 New configuration commands

The following configuration commands were added to make it easier to manage those DNS configuration files on your system:

- ▶ Create Remote Name Daemon Control Configuration (CRTRNDCCFG)
This command is used to generate RNDCCFG configuration files. It is a convenient alternative to writing the rndc.conf file and its corresponding controls and key statements in the named.conf file.
- ▶ DNS Configuration Utility (CHKDNSCFG)
This command checks the syntax of the configuration file called named.conf. It does not provide the support to check the semantics of the configuration file.

► DNS Zone Utility (CHKDNSZNE)

This command checks the syntax and integrity of a zone data file. It is helpful to check zone data files before you add them into a DNS server.

8.4.3 New query and update utilities

The following query and update utilities were added to enhance the management capabilities of your DNS server:

Domain Information Groper

You can use the Domain Information Groper (DIG) query tool to retrieve DNS information about hosts, domains, and other DNS servers based on the response of a DNS server. You can also use it to verify that a DNS server is responding correctly before you configure your system to use it. See the example in Figure 8-6 about how the DIG command can be used to retrieve the IPv4 address of the domain `ibm.com` setting the Query Type to `*A` (IPv4 Address record) and Query Class `*IN` (the Internet class).

```
Start DIG Query (DIG)

Type choices, press Enter.

Query name . . . . . > ibm.com

Query type . . . . . > *A          *A, *AAAA, *ANY, *AXFR...
Query class . . . . . *IN        *IN, *CH, *HS, *ANY
Reverse lookup . . . . . *NO     *NO, *YES, *IP6INT
Domain name server . . . . . *CFG

Domain name server port . . . . 53      1-65535

Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys
```

Figure 8-6 Start DIG Query (DIG)

Start HOST Query

The Start HOST Query (HOST) command is used for DNS lookups. It converts domain names to IP addresses (either IPv4 or IPv6) and vice versa. The command parameters are shown in Figure 8-7 on page 226.

```

                                Start HOST Query (HOST)

Type choices, press Enter.

Query name . . . . . > www.ibm.com

Query type . . . . . *A           *A, *AAAA, *ANY, *AXFR...
Query class . . . . . *IN         *IN, *CH, *HS, *ANY
Domain name server . . . . . *CFG

                                                                Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys

```

Figure 8-7 Start HOST Query (HOST)

Dynamic Update Utility

The Dynamic Update Utility (NSUPDATE) command submits Dynamic DNS Update requests, as defined in Request for Comments (RFC) 2136, to a DNS server. This allows resource records to be added or removed from a zone while the DNS server is running. Thus, you do not need to update records by manually editing the zone file. A single update request can contain requests to add or remove more than one resource record, but the resource records that are dynamically added or removed with the NSUPDATE command should be in the same zone.

Remote Name Daemon Control

The Remote Name Daemon Control (RNDC) command allows a system administrator to control the operation of a name server. It reads a configuration file, called rndc.conf, to determine how to contact the name server and to determine what algorithm and key it should use. If no rndc.conf file is found, then, by default, an rndc-key._KID file that is created during installation is used, which automatically grants access through the loopback interface.

8.5 Netserver

Enhancements in Netserver include:

- ▶ NetServer adds support for IPv6 access in the file sharing and print sharing functions.
- ▶ Starting with IBM i 6.1, support for configuring and managing NetServer file and printer shares is available in the browser-based IBM Systems Director Navigator for System i interface.
- ▶ Improved performance of open and close.

- ▶ Additional characters allowed in file names.

8.6 Virtual Private Network

A Virtual Private Network (VPN) allows your company to securely extend its private intranet over the existing framework of a public network, such as the Internet. With VPN, your company can control network traffic while providing important security features such as authentication and data privacy.

You can now use IP version 6 to create a VPN with the following connection types: host-to-host, host-to-gateway, and gateway-to-gateway. VPN connections support stronger encryption and hashing and IP version 6 to address, range, subnet, host name. All VPN wizards were updated to accept the new IP version 6 ID types.

For more information about VPN, refer to *Security Guide for IBM i V6.1*, SG24-7680.

8.7 FTP client and server extended FTP commands

FTP command enhancements in IBM i 6.1 are:

Extended Passive (EPSV) and Extended Port (EPRT) FTP commands are enhanced in IBM i 6.1.

In extended passive mode, the FTP server operates exactly the same as in passive mode. However, it only transmits the port number (not broken into high and low bytes) and the client is to assume that it connects to the same IP address that was originally connected to. Extended passive mode was added by RFC 2428 in September 1998. The extended passive mode EPSV command requests that a server listen on a data port and wait for a connection. The EPSV command takes an optional argument. The response to this command includes only the TCP port number of the listening connection. The format of the response, however, is similar to the argument of the EPRT command. This allows the same parsing routines to be used for both commands. In addition, the format leaves a place holder for the network protocol and/or network address, which may be needed in the EPSV response in the future. The response code for entering passive mode using an extended address must be 229.

For all FTP transfers where the control and data connections are being established between the same two machines, the EPSV command should be used. Using the EPSV command benefits performance of transfers that traverse firewalls or Network Address Translators (NATs). RFC 1579 [Bel94] recommends using the passive command when behind firewalls since firewalls do not generally allow incoming connections (which are required when using the PORT (EPRT) command). Additionally, using EPSV as defined in RFC 2428 does not require NATs to change the network address in the traffic as it is forwarded. The NAT would have to change the address if the EPRT command is used. If the client issues an EPSV ALL command, NATs may be able to put the connection on a fast path through the translator, because the EPRT command will never be used and therefore, translation of the data portion of the segments will never be needed. Figure 8-8 on page 228 shows an FTP client session requesting the EPSV.

Extended passive mode is supported for both IPv4 and IPv6 networks.

It enables FTP to communicate more seamlessly over a Network Address Translator (NAT) when using a secure SSL control connection.

```
File Transfer Protocol

Previous FTP subcommands and messages:
Connecting to host MySystem at address 192.168.10.10 using port 21.
 220-QTCP at MySystem.
 220 Connection will close if idle more than 5 minutes.
> giancarlo
 331 Enter password.
 230 GIANCARLO logged on.
OS/400 is the remote operating system. The TCP/IP version is "IBMi6.1M0".
 250 Now using naming format "0".
 257 "QGPL" is current library.
> quote epsv
229 Entering Extended Passive Mode (|||35527|).

Enter an FTP subcommand.
===>

F3=Exit   F6=Print   F9=Retrieve
F17=Top    F18=Bottom  F21=CL command line
```

Figure 8-8 FTP Client session using EPSV

8.8 Mail support enhancements

The following topics highlight IBM i 6.1 enhancements in e-mail functions.

8.8.1 Simple Mail Transfer Protocol (SMTP)

Simple Mail Transfer Protocol (SMTP) is the protocol that enables IBM i to send and receive e-mail according to international standard protocols. SMTP is essentially the end-to-end delivery of mail from one mail server to another. There is a direct connection between an SMTP sender (the client) and the destination SMTP receiver (the server). The SMTP client keeps the mail at the sender until it transmits and copies it successfully to the SMTP receiver (server).

IBM i SMTP supports the distribution of notes, messages, and ASCII text documents. SMTP can support formats other than plain text by using the Multipurpose Internet Mail Extensions (MIME) protocol. MIME is the Internet standard for sending mail with headers that describe the contents of the mail messages to the receiving client. These messages can contain video, audio, or binary parts.

The IBM i Information Center contains extensive information on setting up SMTP, its usage, and configuration tips. Select **Expand Network** → **TCP/IP applications, protocols, servers** → **e-mail** at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

The IBM i 6.1 enhancements to the e-mail function include:

- ▶ Simple Mail Transfer Protocol (SMTP) support is extended to send e-mail through a secure connection and authenticate with a user ID and encrypted password using Secure Sockets Layer (SSL) and Transport Layer Security (TLS).
- ▶ Secure e-mail (S/MIME) support and SSL enablement of Simple Mail Transfer Protocol (SMTP) and Post Office Protocol (POP) servers, which enable secure authentication and secure download of e-mail. See Figure 17-1 on page 433 for the new encryption options in the SMTP server properties.
- ▶ A new API, QtmsCreateSendEmail(), builds a MIME document, signs and encrypts it as a S/MIME document, and sends it.
- ▶ The mailhub option enables mail to be forwarded to a single server to assist with the setup of outgoing transactions. You can now define, in the SMTP server attributes (CHGSMTPA), the Forwarding mailhub server (FWDHUBSVR) to specify a domain name of the SMTP server that this SMTP server should forward all non-local e-mails to, as shown in Figure 8-9.

```

                                Change SMTP Attributes (CHGSMTPA)

Type choices, press Enter.

Override reject connect list . . *NO           *SAME, *NO, *YES
Allow bare line feed . . . . . *YES          *SAME, *NO, *YES
Verify identification . . . . . *NO           *SAME, *NO, *YES
Allow authentication . . . . . *NONE         *SAME, *RELAY, *LCLRLY, *NONE
Verify MSF messages . . . . . *YES          *SAME, *YES, *NO
Verify from user . . . . . *ALL            *SAME, *ALL, *LIST, *NONE
Forwarding mailhub server . . . *NONE

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

```

Figure 8-9 SMTP server attributes

FWDHUBSVR(*NONE) means no mail will be forwarded to another SMTP server.

FWDHUBSVR (up to 255 characters) is a fully qualified mail exchanger domain name of the SMTP server that this IBM i partition's SMTP server should forward all non-local e-mails to. This makes this SMTP server behave like a mailhub. If an IP address is specified instead of a domain name, it will not be resolved as a mail exchanger.

8.8.2 Post Office Protocol (POP)

Starting with IBM i 6.1, you can allow Post Office Protocol (POP) to use encryption to secure POP data streams (including passwords). Encryption is provided with either the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS).

8.9 IBM Tivoli Directory Server for i5/OS

Enhancements to IBM Tivoli Directory Server for i5/OS (LDAP) include:

8.9.1 Replication conflict resolution

In a network with multiple master servers, the IBM Tivoli Directory Server has the ability to automatically detect and resolve conflicting changes so that directories on all servers remain consistent. When replication conflicts are detected, the conflicting change is reported in the server log and also recorded in a *lost and found* log file so that an administrator can recover any lost data.

For more information about IBM Tivoli Directory Server troubleshooting, see the IBM Information Center page at the following link and select the path **Networking** → **TCP/IP applications** → **IBM Tivoli Directory Server for i5/OS (LDAP)** → **Troubleshooting Directory Server**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

8.9.2 The ldapmodify command

Two new options have been added to the ldapmodify command:

- e errorfile Specifies the file to which rejected entries are written.
- n Specifies that the changes that would be made are preceded by an exclamation mark and printed to standard output.

8.9.3 Multi-threaded replication

You can replicate using multiple threads, improving overall throughput of replication.

8.9.4 Password encryption

IBM Tivoli Directory Server provides a configuration option for the encryption of user password data prior to storage in the directory. This encryption option can be used to prevent clear text password data from being accessed by regular directory users as well as administrative directory users.

8.9.5 The IBMAttributeTypes attribute

IBM Tivoli Directory Server 6.0 allows for the first 128 characters of an attribute to be used to create the table name.

8.9.6 Schema changes

You can increase the column size by schema modification. This allows you to increase the maximum length of attributes through schema modification by either using Web Administration or the `ldapmodify` utility.

8.9.7 Distributed directory

IBM Tivoli Directory Server has been made to be a distributed directory. Coupled with a proxy server, the distributed directory feature allows a cluster of directories to appear as one. The distributed directory feature along with the proxy server feature will enable directory deployments to hold millions of entries.

8.9.8 `ldapmodrdn`

The IBM Tivoli Directory Server supports `modifyDN` with the `newsuperior` attribute on a leaf node.

8.9.9 Using Trace TCP/IP Application to help find problems

You can use the Trace TCP/IP Application (`TRCTCPAPP`) command to trace an active server instance.

8.9.10 Read access to projected users

You can prohibit all search operations directed to the user `projected`.

8.9.11 Multiple server instances

You can have multiple directory servers on your IBM i system. Each server is known as an instance. If you were using the directory server on a previous release of IBM i, it will be migrated to an instance with the name `QUSRDIR`. You can create multiple instances of the directory server to service your applications.

8.9.12 Password policy

Administrative accounts can be locked due to excessive authentication failures. This feature applies only to remote client connections. The account is reset at server startup. A new attribute is defined for allowing administrative locking of an account.

8.9.13 Directory Server equivalence

The IBM i 6.1 Directory Server is equivalent to IBM Tivoli Directory Server Version 6.0.

8.9.14 Migrating from IBM i 5.4 or i5/OS V5R3 to IBM i 6.1

Use this information if you have a Directory Server running under IBM i 5.4 or i5/OS V5R3.

IBM i 6.1 introduces new functions and capabilities to Directory Server. These changes affect both the LDAP Directory Server and the graphical user interface (GUI) of System i Navigator.

To take advantage of the new GUI functions, you need to install System i Navigator on a PC that can communicate over TCP/IP to IBM i. System i Navigator is a component of System i Access for Windows. If you have an earlier version of System i Navigator installed, you should upgrade to IBM i 6.1.

IBM i 6.1 supports direct upgrades from IBM i 5.4 and i5/OS V5R3. The Directory Server is upgraded to IBM i 6.1 the first time the server is started. The LDAP directory data and the directory schema files are automatically migrated to conform to IBM i 6.1 formats.

When you upgrade to IBM i 6.1, you should be aware of the following migration issues:

- ▶ When you upgrade to IBM i 6.1 and start the, it automatically migrates your schema files to IBM i 6.1 and deletes the old schema files. However, if you have deleted or renamed the schema files, Directory Server cannot migrate them. You might receive an error or Directory Server might assume that the files have already been migrated.
- ▶ After you upgrade to IBM i 6.1, you should first start your server once to migrate existing data before importing new data. If you try to import data before starting the server once and you do not have enough authority, the import might fail. Directory Server migrates directory data to the IBM i 6.1 format the first time that you start the server or import an LDIF file. Plan to allow some time for this migration to complete.
- ▶ IBM i 6.1 introduces the ability to have multiple Directory Server instances on your IBM i system. If you were using the Directory Server before upgrading to IBM i 6.1, your Directory Server will be migrated to an instance. This includes moving the configuration and schema files from the /QIBM/UserData/OS400/DirSrv directory to the /QIBM/UserData/OS400/DirSrv/idsslapd-QUSRDIR directory. This is referred to as the default Directory Server instance and will be named the QUSRDIR instance. Also, two objects in the QUSRSYS library are moved to a new library, QUSRDIRCF. This migration will occur when the Directory Server is started for the first time after the upgrade to IBM i 6.1.
- ▶ Following migration, the LDAP Directory Server will automatically start when TCP/IP starts.

8.9.15 Migrating data from a version prior to i5/OS V5R3 to IBM i 6.1

Use this information if you have a Directory Server running under OS/400 V4R4, V4R5, or V5R1.

IBM i 5.4 does not support direct upgrades from i5/OS V4R4, V4R5, or V5R1.

Note that when you upgrade from OS/400 V4R4 to any later release, you should be aware of the following issues:

- ▶ OS/400 V4R4 and earlier releases of Directory Server did not take time zones into account when creating time stamp entries. Beginning with OS/400 V4R5, the time zone is used in all additions and modifications to the directory. Therefore, if you upgrade data from OS/400 V4R4 or earlier, Directory Server adjusts existing createtimestamp and modifytimestamp attributes to reflect the correct time zone. It does this by subtracting the time zone that is currently defined on the system from the time stamps that are stored in the directory. Note that if the current time zone is not the same time zone that was active when the entries were originally created or modified, the new time stamp values will not reflect the original time zone.
- ▶ If you are upgrading data from OS/400 V4R4 or earlier, be aware that the directory data will require approximately twice as much storage space as it required previously. This is because in OS/400 V4R4 or earlier versions, Directory Server supported only the IA5 character set and saved data in ccsid 37 (single byte format). Directory Server supports

the full ISO 10646 character set. After you upgrade, you should start your server once to migrate existing data before importing new data. If you try to import data before starting the server once and you do not have enough authority, the import might fail.

If you want to migrate these releases to IBM i5.4, you can follow either of the following procedures:

- ▶ Upgrading from OS/400 V4R4, V4R5, or V5R1 to an interim release:
You can migrate your Directory Server to upgrade to an interim release (OS/400 V5R2 or i5/OS V5R3), and then to IBM i 6.1.
- ▶ Saving the database library and installing IBM i 6.1
You can migrate your Directory Server by saving the database library that Directory Server uses in OS/400 V4R4 or V4R5 and then restoring it after installing IBM i 6.1.

8.10 Advanced program-to-program communications

To prevent a potential data integrity problem, advanced program-to-program communications (APPC) waits for an acknowledgement from the remote system after sending the DETACH signal to end the transaction. You can now configure the wait time using the QACRETRY and QACINTERVL data areas in the QGPL library.

You can use the QACRETRY data area to specify the number of retries to receive the acknowledgement, and use the QACINTERVL data area to specify the interval for each retry. To create the data areas:

- ▶ Create these two data areas and set the values by using the Create Data Area (CRTDTAARA) command.
 - Set the type of values that are contained in the two data areas to *DEC using the TYPE parameter.
 - Set the length for these two data areas to 2 using the LEN parameter.
 - Set the initial value for these two data areas using the VALUE parameter, with 99 as the maximum value.

For example, the following command creates the QACRETRY data area and sets the number of retries to 99:

```
CRTDTAARA DTAARA(QGPL/QACRETRY) TYPE(*DEC) LEN(2) VALUE(99)
```

If you do not create either or both of these two data areas, the default values are used:

- ▶ 8 for the number of retries.
- ▶ 15 seconds for the interval for each retry.

For either of the two data areas, if you create the data area but do not specify the initial value for it, the data area is initialized to a value of 0.

Set appropriate values for the two data areas of the wait time based on the specific situation:

- ▶ For network connections with slow SNA request-response unit (RU) from the remote application and a slow-speed switched line, set a relatively long wait time. By doing so, you can ensure that the data and the DETACH signal can flow to the remote system within the wait time.

For those remote applications that do not send acknowledgements when they are finished, set an appropriate wait time. By doing so, you can prevent the local system from waiting

longer than necessary before sending the UNBIND signal to complete the flow of the conversation between the two systems.



Printing

This chapter discusses the following topics:

- ▶ “Printing overview” on page 236
- ▶ “Dates and times in Coordinated Universal Time” on page 237
- ▶ “Adding a printer with System i Navigator” on page 239
- ▶ “Generating PDF files with IBM i” on page 245
- ▶ “Start Spool Reclaim command” on page 251
- ▶ “Updated open feedback area” on page 251

9.1 Printing overview

IBM i has powerful printing and output functions. You can present information using overlays, bar codes, graphics, images, and more. IBM i supports a variety of industrial-type printing and presentation solutions.

Choosing and implementing a printing and presentation solution requires you to be familiar with both your organization's requirements and resources, and the capabilities provided by IBM i.

IBM i has both Basic Printing and Advanced Function Presentation (AFP).

Note: AFP is an architecture-based system of hardware and software for creating, formatting, viewing, retrieving, printing, and distributing information using a wide variety of printer and display devices. AFP is the original, integrated data stream on IBM i for generating fully composed pages of data.

The following steps offer a high-level overview of the IBM i printing process:

- ▶ The printing process starts when an application program runs. The application program creates output data. The output data is based on the application program and information contained in the printer file.
- ▶ If print spooling is selected, the output data is placed in a spooled file and the spooled file is placed in an output queue. If direct printing is selected, the output data is sent directly to the printer.
- ▶ The destination of the output data is based on values stored in several printing elements, such as job description, user profile, workstation description, printer file and system values. Output queues are used to manage spooled files.
- ▶ Spooled files in output queues can be:
 - Printed
 - Kept as records
 - Used as input to other applications
 - Transferred to other output queues
 - Sent as e-mail
 - Used to create PDF files

Note: Spooled files can be received from other IBM i and non-IBM i systems.

- ▶ The printer writer program interacts between the output queue and the printer and can be used to convert the printer data stream.
- ▶ The printer writer program included in IBM i supports a variety of printer data streams. Print Services Facility™ for IBM i provides additional function that provides support for the Advanced Function Presentation (AFP) Intelligent Printer Data Stream (IPDS).
Each printer must have a printer device description. The printer device description contains a configuration description of the printer. Printers can be attached by a variety of attachment methods.
- ▶ A remote writer allows you to route spooled files from an output queue on your system to another system.

Figure 9-1 shows the complete printing process on IBM i.

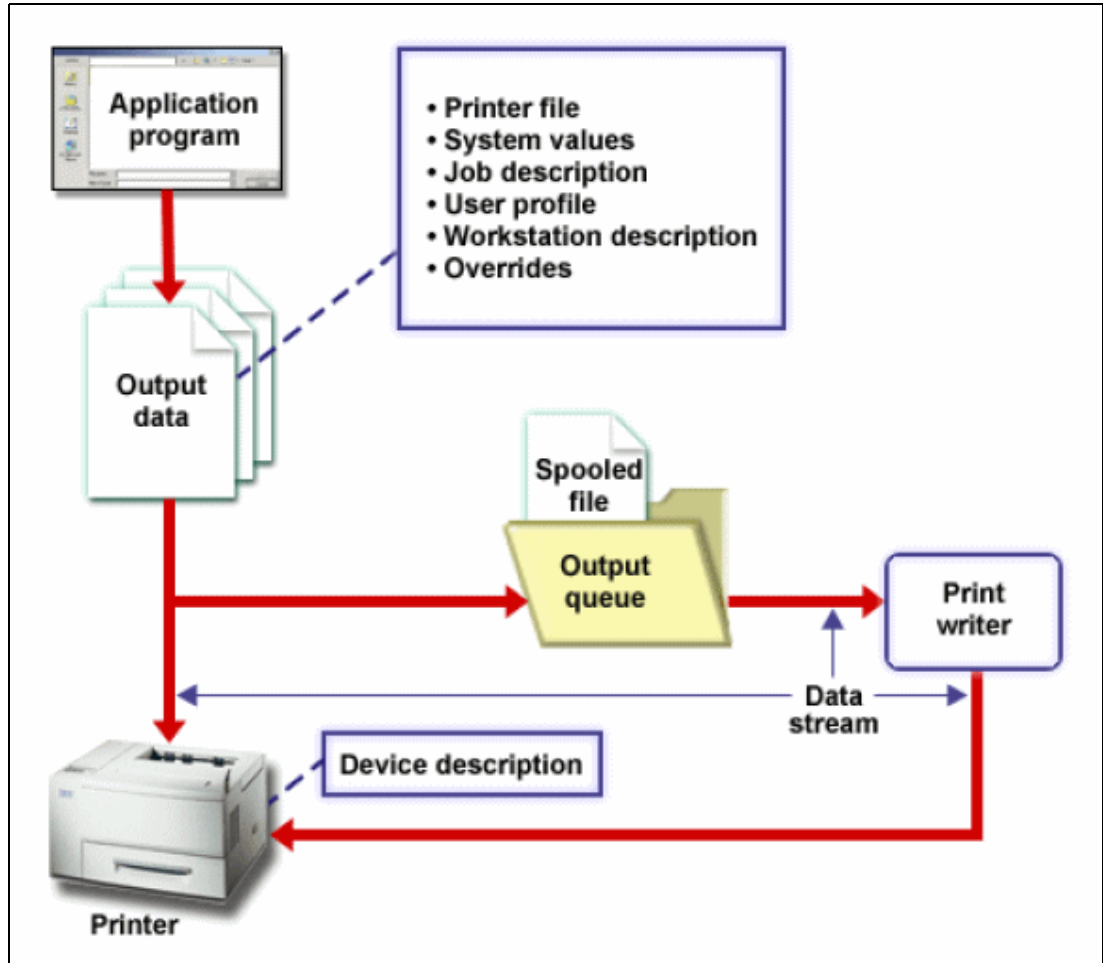


Figure 9-1 IBM i printing process

The following sections of this chapter introduce the new or significantly changed information for basic printing in IBM i 6.1:

- ▶ “Dates and times in Coordinated Universal Time” on page 237
- ▶ “Adding a printer with System i Navigator” on page 239
- ▶ “Generating PDF files with IBM i” on page 245
- ▶ “Start Spool Reclaim command” on page 251
- ▶ “Updated open feedback area” on page 251

9.2 Dates and times in Coordinated Universal Time

The dates and times used throughout the spooling function are now stored in Coordinated Universal Time (UTC). To support UTC, some changes are made to the data queue formats.

Data queues are a type of system object that you can create, to which one high-level language (HLL) procedure or program can send data, and from which another HLL procedure or program can receive data. The receiving program can be already waiting for the data, or can receive the data later.

There are two types of data queue support for spooled files:

9.2.1 Data queue support on output queues

Support is available to optionally associate a data queue with an output queue using the Create Output Queue (CRTOUTQ) or Change Output Queue (CHGOUTQ) command. Entries are logged in the data queue when spooled files are in ready (RDY) status on the output queue. A user program can determine when a spooled file is available on an output queue using the Receive Data Queue (QRCVDTAQ) API to receive information from a data queue.

The data queue must be created with a maximum message length (MAXLEN) parameter value of at least 128 bytes. The sequence (SEQ) parameter value should be *FIFO or *LIFO. The format of the CRTDTAQ command is:

```
CRTDTAQ DTAQ (<library name>/<data queue name>) MAXLEN(128) SEQ(*LIFO)
```

Then, on the DTAQ keyword of the Create Output Queue (CRTOUTQ) command or the Change Output Queue (CHGOUTQ) command, specify the data queue that is created with the CRTDTAQ command.

After a data queue is associated with an output queue, any spooled file that is placed on the output queue in a ready status causes an entry to be placed on the data queue.

9.2.2 Environment variable data queue support

Using the Add Environment Variable (ADDENVVAR) or Change Environment Variable (CHGENVVAR) command, you can associate a data queue with a job or with the system. As spooled files are created, the entries are logged in the data queue. Using the Receive Data Queue (QRCVDTAQ) API to receive information from the data queue, a user program can determine when a spooled file has been created by the job or by the system.

Using the CL command ADDENVVAR, and specifying a fully qualified data queue name for the environment variable QIBM_NOTIFY_CRTSPLF, you can associate a data queue with a job or with the system.

Important: The data queue must be created with a record length of at least 144 bytes. It must also have a public authority of *USE, or you need to grant the QSPL user profile *USE private authority to the data queue. You must ensure that the containing library has a public authority of *EXECUTE, or you need to grant the QSPL user profile *EXECUTE private authority to the library. The format of the CRTDTAQ command is:

```
CRTDTAQ DTAQ (<library name>/<data queue name>) MAXLEN(144) AUT(*USE)
```

As of IBM i 6.1, if you want the data queue information to have dates and times in Coordinated Universal Time (UTC), you can use the same environment variable with a different value. Replace the *DTAQ with *DTA2 in the ADDENVVAR command above. This value causes a type 03 data queue entry to be placed in the data queue instead of a type 02 entry. The command format is as follows:

```
ADDENVVAR ENVVAR(QIBM_NOTIFY_CRTSPLF)  
VALUE('*DTA2 <library name>/<data queue name>') LEVEL(*JOB | *sys)
```

The length of the type 03 data queue entry is 200 bytes. Use the following command format to create a type 03 data queue:

```
CRTDTAQ DTAQ (<library name>/<data queue name>) MAXLEN(200) AUT(*USE)
```

Once a data queue is associated with a job or the system, any spooled file created by the job or system will automatically have an entry placed in the data queue.

Table 9-1 shows the new and enhanced record type 01 data queue entry formats.

Table 9-1 Updated record type 01 data queue entry formats

Decimal Offset	Hex Offset	Type	Description
80	50	CHAR(7)	Spooled file creation date Identifies the date on which the spooled file was created in CYYMMDD format (local system time).
87	57	CHAR(1)	Reserved
88	58	CHAR(6)	Spooled file creation time Identifies the time that the spooled file was created in HHMMSS format (local system time).
94	5E	CHAR(7)	Spooled file creation date in UTC Identifies the date on which the spooled file was created in CYYMMDD format.
101	65	CHAR(1)	Reserved
102	66	CHAR(6)	Spooled file creation time in UTC Identifies the time that the spooled file was created in HHMMSS format.
108	6C	CHAR(20)	Reserved

Table 9-2 shows the new and enhanced record type 02 data entry formats.

Table 9-2 Updated record type 02 data queue entry formats

Decimal Offset	Hex Offset	Type	Description
130	82	CHAR(7)	Creation date Identifies the date on which the spooled file was created in CYYMMDD format (local system time).
137	89	CHAR(6)	Creation time Identifies the time that the spooled file was created in HHMMSS format (local system time).
143	8F	CHAR(1)	Reserved

For information about the new record type 03 data queue entry format, refer to Information Center for IBM i at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzalu/rzalurectype03.htm>

9.3 Adding a printer with System i Navigator

A new Add Printer Wizard was added in System i Navigator for IBM i 6.1. The printer wizard is used for adding and configuring IBM i printers. The Add Printer Wizard in System i Navigator

is also used to create a printer device description and its associated objects. This is an alternative to using CL commands to add and configure printers.

The IBM Systems Director Navigator for i has a similar wizard interface under **Basic Operations** → **Add a network printer**. It also provides printer administration functions. For more information about IBM Systems Director Navigator for i, refer to Chapter 16, “IBM Systems Director Navigator for i” on page 403.

The following steps show the process of using the System i Navigator for adding and configuring an IBM i printer attached to the network through TCIP/IP:

1. In System i Navigator, expand **My Connections** → **System name** → **Basic Operations**.
2. Right-click **Printers**.
3. On the menu, select **Add printer**; see Figure 9-2.

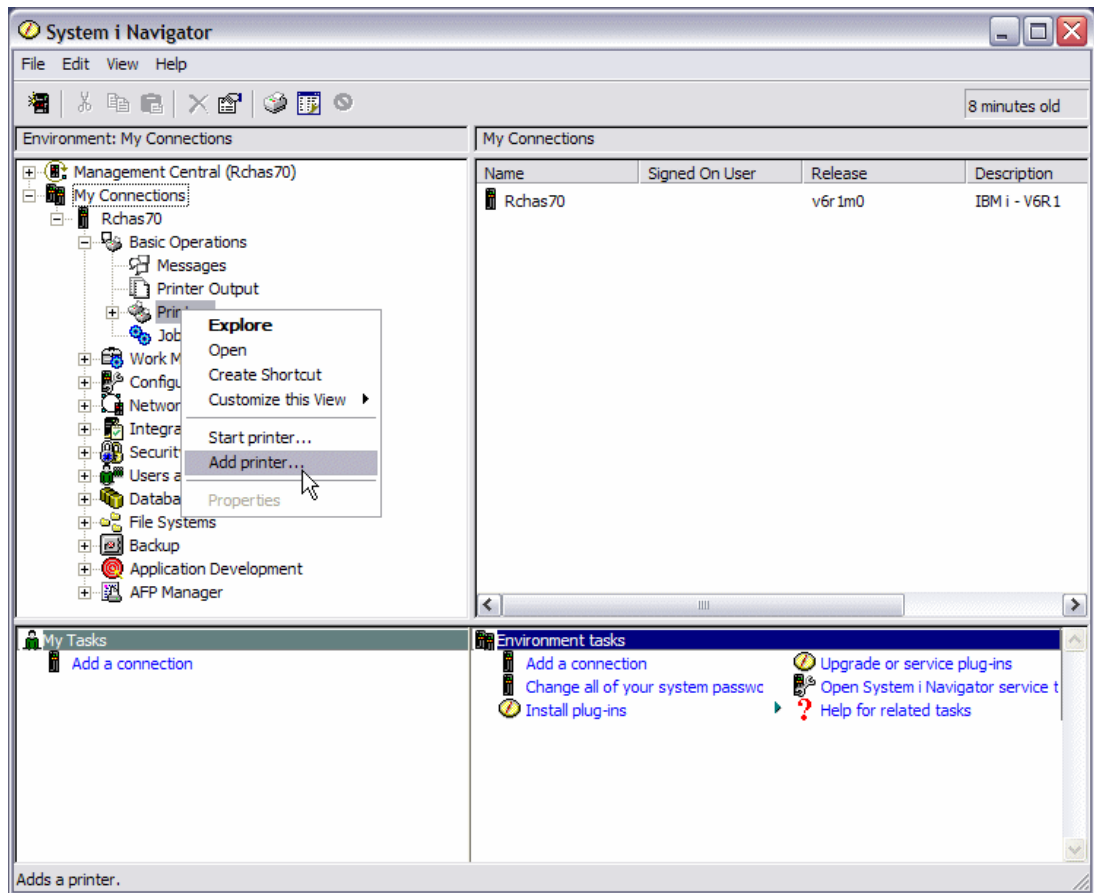


Figure 9-2 Add printer from context menu

4. The Printer Configuration Wizard is launched. Click **Next** to continue.

5. Choose the location of your printer: DNS, IP Address or URL, as shown in Figure 9-3, and click **Next**.

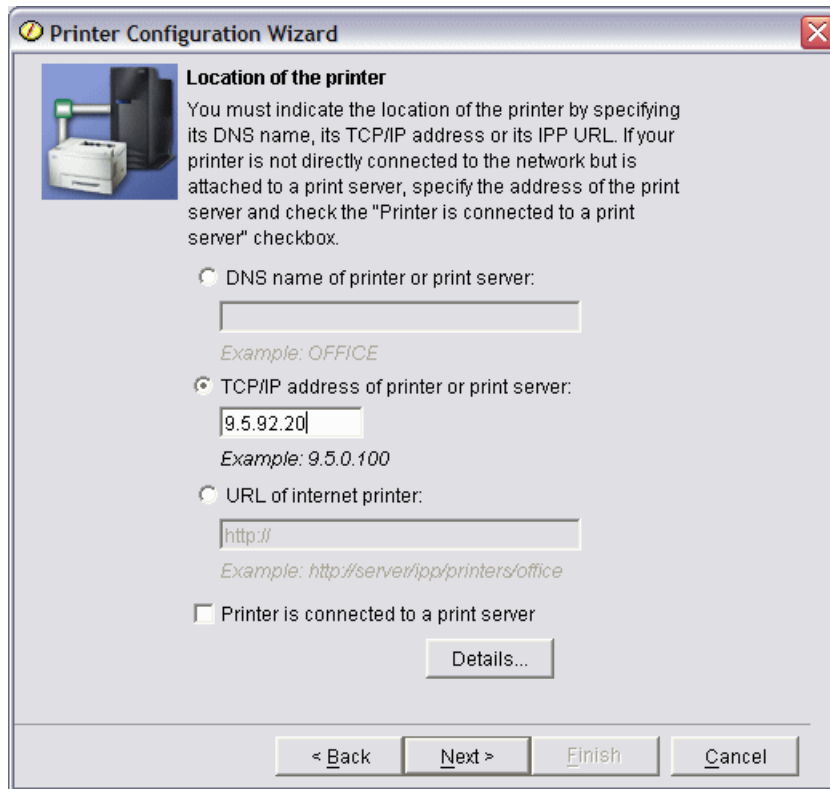


Figure 9-3 Setting location of the printer

Note: Check the *Printer is connected to a print server* check box, if your printer is connected to the network through a print server.

6. If “Printer is connected to a print server” was checked, then the printer server manufacturer and model should be selected, as shown in Figure 9-4.

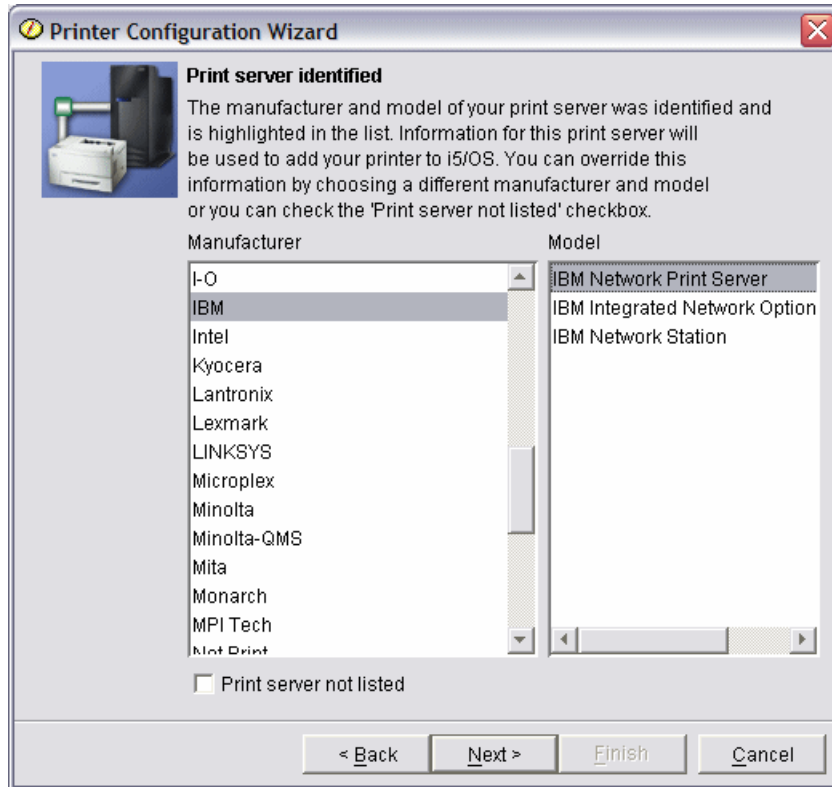


Figure 9-4 Print server manufacturer and model

Physical port number should be selected, if the print server has more than one physical port and the printer is not connected to port 1 of the print server, as shown in Figure 9-5, and click **Next**.

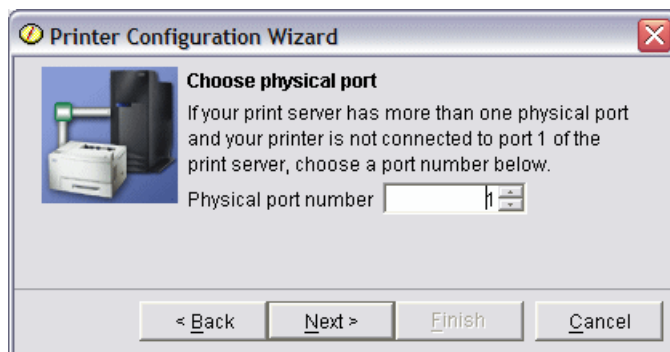


Figure 9-5 Choose physical port number

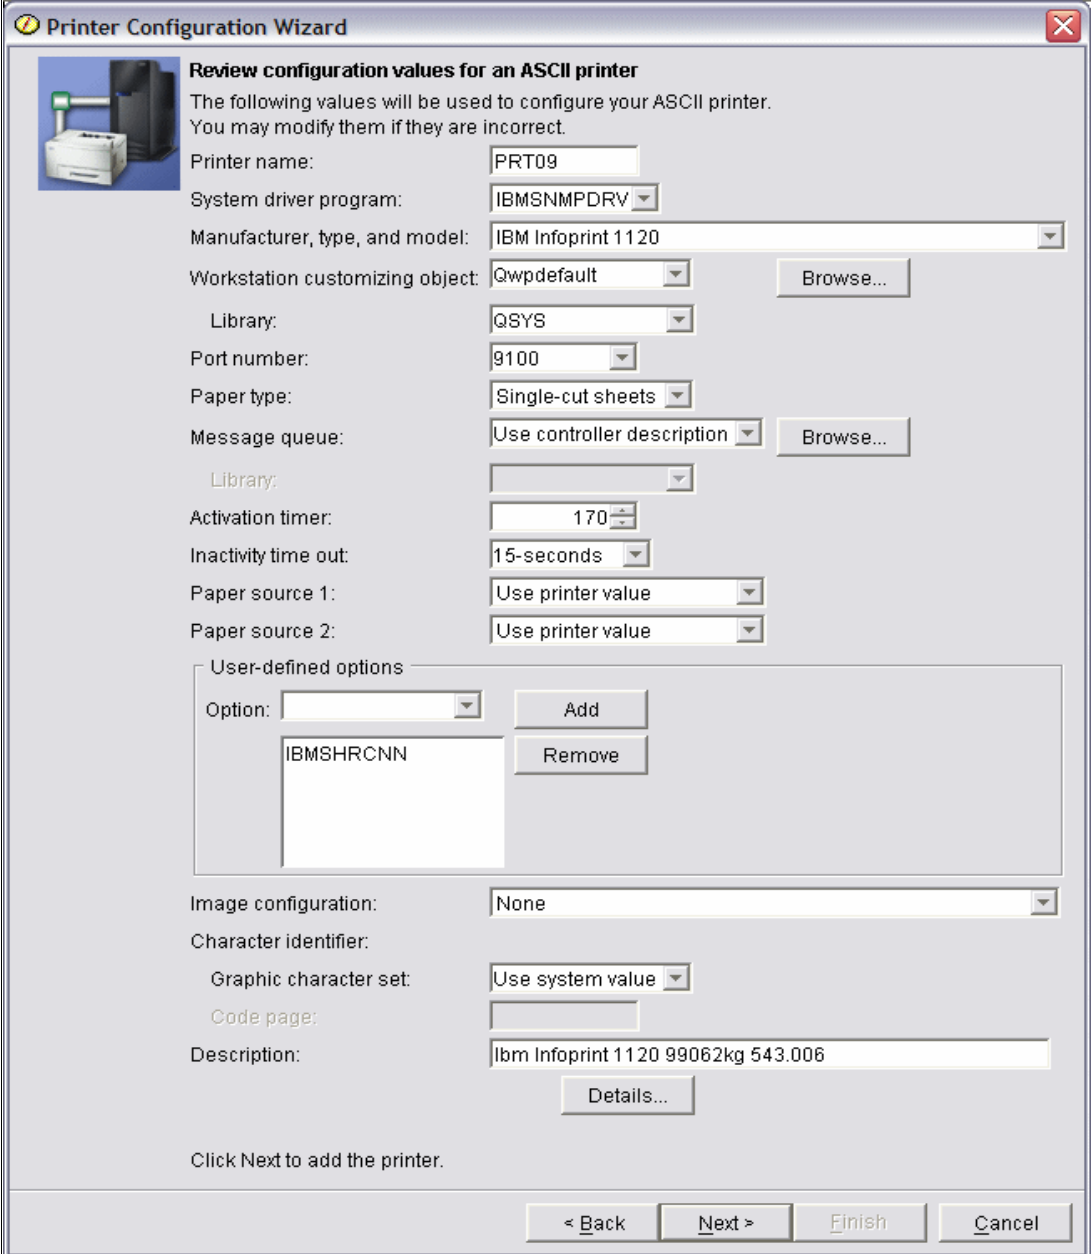
7. Select the printer manufacturer and model from the list, and click **Next**. If the printer is not listed check the *Printer not listed* check box.
8. Choose whether the printer should be configured as IPDS capable or not, and click **Next**.

Note: Intelligent Printer Data Stream (IPDS) printers use an all-points-addressable data stream that allows users to position text, images, and graphics at any defined point on a printed page.

The Print Services Facility (PSF) product is required when IPDS is used with a network-attached printer.

Accept or change the default values for printer configuration and click **Next**.

9. Review and change the printer configuration as shown in Figure 9-6, then click **Next**.



Printer Configuration Wizard

Review configuration values for an ASCII printer

The following values will be used to configure your ASCII printer. You may modify them if they are incorrect.

Printer name: PRT09

System driver program: IBMSNMPDRV

Manufacturer, type, and model: IBM Infoprint 1120

Workstation customizing object: Qwpdefault

Library: QSYS

Port number: 9100

Paper type: Single-cut sheets

Message queue: Use controller description

Library:

Activation timer: 170

Inactivity time out: 15-seconds

Paper source 1: Use printer value

Paper source 2: Use printer value

User-defined options

Option:

IBMSHRCNN

Image configuration: None

Character identifier:

Graphic character set: Use system value

Code page:

Description: lbm Infoprint 1120 99062kg 543.006

Click Next to add the printer.

< Back Next > Finish Cancel

Figure 9-6 Review and change printer configurations

Following are definitions for some of the configuration settings:

► System driver program

Specifies the name of a system-defined and provided driver program, which provides the capability to send IBM i printer output to a printer attached over a TCP/IP network.

Possible values include:

- IBMSNMPDRV - An IBM SNMP printer driver program is used.
- IBMPJLDRV - An IBM network printer driver program is used.
- HPPJLDRV - An HP-compatible printer driver program is used.

► Workstation customizing object specifies the qualified name of a work station customizing object to be used when printing. This parameter is ignored when Advanced Function Printing (AFP) is used.

– Port number

Specifies the port number of the printer. Valid values range from 0 through 65535. Some printer devices may further restrict the valid value range.

– Message queue

Displays the message queue and library that receive messages sent to the printer when the writer is started.

– Activation timer

Specifies the amount of time (in seconds) to wait for the printer to respond to the activation request from the IBM i printer driver. If the printer does not respond within this time, it is considered not available.

– Inactivity Time out

Specifies the amount of time the IBM i printer driver will wait before it closes the connection to the printer when there are no other IBM i spooled files to be printed. This will free the printer, allowing it to receive spooled files from another source.

– User-defined options

Specifies, for spooled output only, one or more user-defined options to be used by user applications or user-specified programs that process spooled files. A maximum of four user-defined options can be specified.

– Image configuration

Specifies the image configuration object for this printer. An image configuration object provides transform services for a variety of image and print data stream formats.

– Character identifier

Specifies, for printers, the character identifier (graphic character set and code page) indicating which font is used to print the job and file separator pages when no font is explicitly specified. This parameter is used only when Advanced function printing (AFP) or host print transform is used.

Note: This parameter is used only when Advanced function printing (AFP) or host print transform is used.

10. Click **Next**, then **Finish**. The new printer is now listed in the System i Navigator printers list, as shown in Figure 9-7 on page 245.

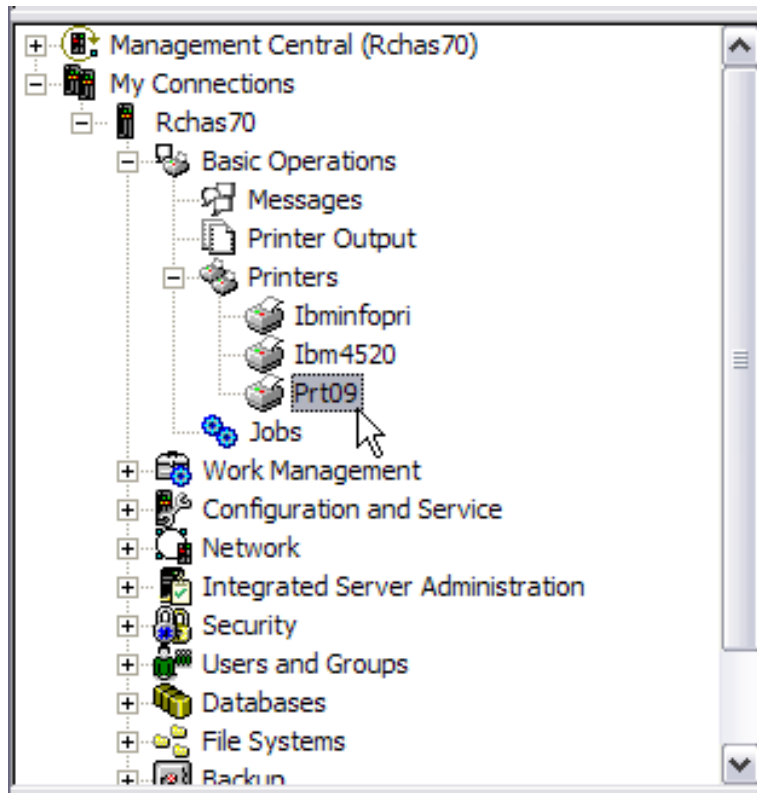


Figure 9-7 New printer in printers list

9.4 Generating PDF files with IBM i

You can generate Adobe® Portable Document Format (PDF) documents directly from IBM i applications with only a minor change or override to your existing printer files.

9.4.1 New print transform support

There is a new licensed program available with IBM i 6.1: IBM Transform Services for i5/OS 5761TS1 - option 1 provides AFP to PDF transform. It delivers Adobe Portable Document Format (PDF) output support to IBM i applications:

- ▶ Document format controlled by existing printer file parameters and DDS keywords
- ▶ Text data preserved in standard encoding
- ▶ PDF output optionally written to Integrated File System (IFS) directory to make it readily accessible to client applications

Additional output options available with printer output:

- ▶ Write AFP output data to stream files in IFS instead of spooled files
- ▶ Output formats supported by the Host Print Transform (HPT)
 - Tagged Image File Format (TIFF)
 - Graphic Interchange Format (GIF)

9.4.2 New print transform prerequisites

The following products and options must be installed to provide PDF output support:

- ▶ 5761-SS1, Option 03: Extended Base Directory Support (IFS)
- ▶ 5761-SS1, Option 33: Portable Application Solutions Environment (PASE)
- ▶ 5761-TS1: IBM Transform Services for i5/OS (new in IBM i 6.1)
 - *BASE and option 1: AFP to PDF Transform

9.4.3 Workstation Customizing object (WSCST) parameter

You can use the WSCST parameter to transform the printer file output to final form before writing it to a stream file. This parameter is only valid when the print file has the following:

- ▶ The `to stream file (TOSTMF)` parameter is not set to `*NONE`.
- ▶ The `device type (DEVTYPE)` parameter is set to `*AFPDS`.

The WSCST object specified must have an attribute of `TRANSFORM` or `CTXFORM`. If the WSCST object has an attribute of `TRANSFORM`, the host print transform function is used to convert the data before writing it to the stream file.

If the WSCST object has an attribute of `CTXFORM`, the Portable Application Solutions Environment (PASE) and the IBM Transform Services for i5/OS licensed program (5761-TS1) must be installed. Three `CTXFORM` WSCST objects are included with IBM i, which will transform the output data to the Adobe Portable Document Format (PDF). These objects require option 01 (AFP to PDF Transform) of the 5761-TS1 product to be installed. They differ in how they map AFP font references to PDF font references.

5761-TS1 includes three special IBM-provided customization objects:

- ▶ **QCTXPDF**

This is the default PDF WSCST object that is used when the special value `*PDF` is specified. It includes tags to map all of the CJK World Type fonts included in IBM i option 43 to corresponding registered CMaps and character collections. This allows ideographic characters to be rendered without embedding these very large fonts within the document. The smaller World Type fonts are not mapped by this object and will be embedded in the output document to allow non-Latin1 SBCS languages to be presented.

- ▶ **QCTXPDFWT**

This object contains tags to map all of the World Type fonts included with IBM i option 43 to corresponding Type 1 fonts or registered CMaps and character collections. This generally results in the smallest possible PDF file size, but only reliably presents Latin and CJK languages.

- ▶ **QCTXPDFMBD**

This is the minimal PDF WSCST object included. It omits all font tags, so that the default font mapping is used for all font references. The default behavior is to map all SBCS AFP font references to standard PDF fonts, map all DBCS AFP fonts to common CID fonts, and embed all TrueType font references.

9.4.4 Customizing WSCST object of *CTXFORM transform type

The way the PDF transform maps fonts is by customizing your WSCST object of the `*CTXFORM` transform type.

As with the previous WSCST object classes, you can retrieve the source from shipped WSCST objects using the Retrieve WSCST source (RTVWSCST) command, and generate customized versions with the Create WSCST (CRTWSCST) command. Because of the need to support open character strings for many parameters, the source physical file used for WSCST source of this class is encoded in UTF-8 (CCSID 1208). Attempts to use source physical files of a different encoding with either command will fail, and message CPF5D33 or CPF5D34 will be issued.

WSCST Device Class (WSCST) tag

WSCST objects of this device class are identified by specifying the value CTXFORM in the WSCST tag. This must be the first tag specified in the WSCST source member. The syntax is:

```
:WSCST DEVCLASS = CTXFORM.
```

This value indicates that the WSCST object defines configuration information for a transform from the IBM Transform Services for i5/OS licensed program (5761-TS1). This product and the appropriate product option must be installed to use WSCST objects of this type.

9.4.5 Create printer files

The Create Printer File (CRTPRTF) command creates a printer device file from the information specified on this command, and optionally from the data description specifications (DDS) contained in a source file.

A printer device file is used to send records to a printer device. The printer device file identifies the printer device used and the spooling requirements; it does not contain data.

In 6.1 two new parameters were added to the CRTPRTF command, as shown in Figure 9-8 on page 248.

The new parameters are:

► To Stream File (TOSTMF)

Specifies the directory or stream file where the output data is to be written. All directories in the path name must exist.

Note: This parameter is only valid when the Workstation customizing object (WSCST) parameter is not *NONE, or the Printer device type (DEVTYPE) parameter is *AFPDS.

The TOSTMF parameter has the following options:

– *NONE

The output is written to a spooled file if the SPOOL parameter is *YES, or directly to the printer device (DEV) if the SPOOL parameter is *NO.

– path-name

Specify the path name for the directory or stream file where you want the output data to be written.

Note: Printer files must be explicitly closed before a job using the file ends. By the time the Workstation Print data management module is called during job end processing, access to IFS has been lost. If that happens, a stream file will have incomplete data or no data at all.

► Workstation Customizing Object (WSCST)

Specifies the workstation customizing object to use to transform the printer file output to final form before writing it to a stream file. This parameter is only valid when the To stream file (TOSTMF) parameter is not *NONE. For more information about the WSCST parameter, refer to 9.4.3, “Workstation Customizing object (WSCST) parameter” on page 246.

```

                                Create Printer File (CRTPRTF)

Type choices, press Enter.

User Defined Option . . . . . *NONE          Character value, *NONE
      + for more values
User Defined Data . . . . . *NONE

User Defined Object:
  Object . . . . . *NONE          Name, *NONE
  Library . . . . .              Name, *LIBL, *CURLIB
  Object type . . . . .          *DTAARA, *DTAQ, *FILE...
To stream file . . . . . *NONE

Workstation customizing object *NONE          Name, *NONE, *PDF
  Library . . . . .              Name, *LIBL, *CURLIB
  Maximum file wait time . . . . . *IMMED      Seconds, *IMMED, *CLS
  Share open data path . . . . . *NO          *NO, *YES
  
```

Figure 9-8 New parameters for the CRTPRTF command

Note: TOSTMF and WSCST parameters are available for the following printer file CL commands:

- CRTPRTF
- CHGPRTF
- OVRPRTF

Use the following steps to create a print file for PDF:

1. Create the print file:

```

CRTPRTF FILE(QGPL/PRTPDF) DEVTYPE(*AFPDS)
TEXT('Print file for PDF output') EXPDATE(*DAYS) DAYS(21)
TOSTMF('/printing') WSCST(*PDF)
  
```

2. Create an IFS folder, for example /printing:

Using the IBM i command Add Link (ADDLNK) or the System i Navigator as shown in Figure 9-9 on page 249:

File Systems → Integrated File System → Right Click on File System name → New Folder

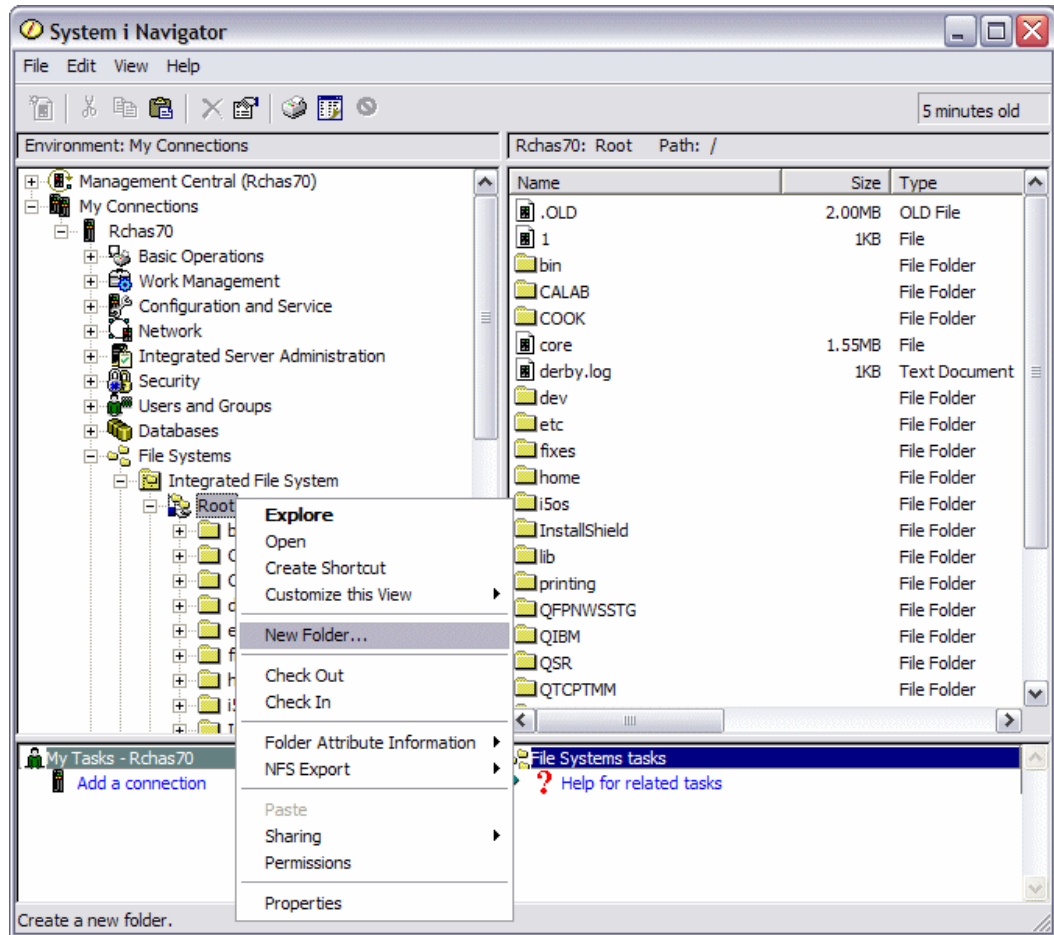


Figure 9-9 Create an IFS folder

3. Override the printer file:
OVRPRTF FILE(QPDSPLIB) TOFILE(QGPL/PRTPDF)
4. Display a library:
DSPLIB LIB(PFREXP) ASPDEV(*) OUTPUT(*PRINT)

9.4.6 PDF output considerations

Document format considerations

The following considerations apply to the document format:

All text data remains searchable with standard encoding:

- ▶ ANSI Windows for SBCS data
- ▶ Registered PDF CMaps for DBCS data
- ▶ ToUnicode CMaps for all data with TrueType font reference

By default, most fonts are not embedded in the document:

- ▶ WSCST objects can be modified to embed Type1 fonts.
- ▶ WSCST objects can be modified not to embed TrueType fonts.

- ▶ QCTXPDF maps CJK TrueType fonts shipped in option 43.
- ▶ QCTXWT maps all TrueType fonts shipped in option 43.

Page size specified on the printer file is used for PDF output:

- ▶ Computer Output Reduction (COR) not applied.
- ▶ Most client print drivers can *scale-to-fit* printed output.

Resource considerations

The following considerations apply to the resources:

AFP Library Resources (IBM i objects) support includes:

- ▶ Overlays
- ▶ Page segments
- ▶ Form Definitions

AFP IFS Resources (AFPRSC keywords) format supports:

- ▶ JFIF (JPEG File Interchange Format -JPEG)
- ▶ TIFF (Tag Image File Format)
- ▶ GIF (Graphics Interchange Format)

Limitations

The following limitations apply:

- ▶ Some bar code types currently not supported in the BARCODE keyword.
- ▶ AFP font resources not transformed:
 - User-customized fonts
 - Non-Latin SBCS encoding
- ▶ CHRSIZE keyword not supported.
- ▶ Constant Back overlays not supported.
- ▶ UNISCRIP keyword not supported.

Customization considerations

The following considerations apply to customization:

- ▶ Source file must be created with CCSID 1208 (UTF-8):
`CRTSRCPF FILE(WSCSTSRC) CCSID(1208)`
- ▶ Use the RTVWSCST command to retrieve source from existing WSCST objects:
`RTVWSCST DEVTYPE(*CTXFORM) MFRTYPMDL(*PDF) SRCMBR(MYWSCST) SRCFILE(WSCSTSRC)`
- ▶ Use the CRTWSCST command to create a WSCST object:
`CRTWSCST WSCST(MYWSCST) SRCFILE(WSCSTSRC)`
- ▶ Workstation customization functions can be used to customize the functions of ASCII printers that use the host print transform function. Using the workstation customization functions, the following customizations can be done:
 - Customize the functional characteristics of a supported ASCII printer.
 - Customize the functional characteristics and specify all necessary parameters required to support a normally unsupported ASCII printer.

For more information about Customizing ASCII printers that use the host print transform function, refer to IBM i operating system Information Center for IBM i 6.1 at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzalu/p2ch09.htm?tocNode=int_175666

9.5 Start Spool Reclaim command

The Start Spool Reclaim (STRSPLRCL) command is used to fix output queues and spooled files that are left in unrecoverable states.

If a writer job ends abnormally or an incomplete update has occurred, the output queue or spooled files associated with the writer job could remain in various statuses (such as WTR, PRT, PND, and MSGW).

The STRSPLRCL command does not wait for the spool reclaim function to complete. You can use this command to fix output queues and spooled files in SYSBASE, in an independent auxiliary storage pool (IASP), or in both.

If an asterisk (*) is specified for the ASP group (ASPGRP) parameter and the current thread has an ASP group in its name space, two separate requests to reclaim output queues could be sent. If a matching output queue is found in *SYSBAS, a request is sent to the spool maintenance job in *SYSBAS. If a matching output queue is found in the specified ASP group, a request is sent to the spooled maintenance job for that ASP group. When the reclaim function has completed, message CPC3309 is sent to the QHST and QSYSOPR message queues by each spooled maintenance job.

If *ALL/*ALL is specified for the output queue (OUTQ) parameter, and *SYSBAS or * is specified for the ASP group (ASPGRP) parameter, additional cleanup is done on job queues and completed jobs. If a job is in OUTQ status but does not own any active spooled files, the job will be removed from the system.

Important: Spool control (*SPLCTL) special authority is needed to use this command. If a specific ASP group is specified for the ASPGRP parameter, the user must have use (*USE) authority to all ASP device descriptions in the ASP group and the status of ASP group must be AVAILABLE.

9.6 Updated open feedback area

The open data path (ODP) contains information about the merged file attributes and information returned by input or output operations. The ODP only exists while the file is open.

The open feedback area is the part of the open data path (ODP) that contains general information about the file after it has been opened. It also contains file-specific information, depending on the file type, plus information about each device defined for the file. This information is set during open processing and can be updated as other operations are performed.

In 6.1 the open feedback area for the printer has been updated with additional sections.

Table 9-3 shows the updated information in 6.1.

Table 9-3 ODP Changes in 6.1

Offset	Data type	Length in bytes	Contents
0	Character	2	Open data path (ODP) type: ▶ DS Printer file not being spooled. Display, tape, intersystem communications function (ICF), save, or diskette file not being spooled. ▶ SP Printer or diskette file being spooled or inline data file. ▶ ST Printer file data is being written to an IFS streamed file.
137	Binary	2	CCSID - Coded Character Set Identifier
139	Character	1	Reserved
140	Binary	4	Offset to the extensions list. Number of bytes from the start of the open feedback area to the extensions offset table. The open feedback extensions table occurs after the device name list in the ODP.
144	Binary	2	Number of devices defined for this ODP. For printers this always has a value of 1.
146	Character		Device name of definition list.

The printer open feedback area has the following parts:

▶ Device definition list

The device definition list part of the open feedback area is an array structure. Each entry in the array contains information about each device or communications session attached to the file.

The number of entries in this array is determined by the number at offset 146 of the open feedback area. The device definition list begins at offset 148 of the open feedback area. The offsets shown for it are from the start of the device definition list rather than the start of the open feedback area.

▶ Open feedback extensions list

The open feedback extension list contains an array of offsets to new sections of the open feedback area.

Currently, only the first open feedback extension list entry is defined.

Table 9-4 shows the open feedback extensions list entries.

Table 9-4 Open feedback extensions list entries

Offset	Data type	Length in bytes	Contents
0	Binary	4	Length of the open feedback extensions list

Offset	Data type	Length in bytes	Contents
4	Binary	4	Offset from the beginning of the open feedback area to the stream file name extension. This field is only valid when the open data path (ODP) type is ST.
8	Binary		Reserved

► Stream file name open feedback extension:

The stream file name open feedback extension is an optional part of the open feedback area that describes the output path.

The stream file name is always expressed as an absolute path name, even if a relative path was specified when the file was opened. For example, the stream file name might be:

`‘/home/username/mydir/myfile.pdf’`

Table 9-5 shows the steam file name entries.

Table 9-5 Steam file name entries

Offset	Data type	Length in bytes	Contents
0	Binary	2	CCSID of the stream file path name
2	Binary	4	Length in bytes of the path name
6	Binary	2	Offset from the beginning of the stream file name open feedback extension to the path string
	Char		Absolute path name of the stream file



Integration with BladeCenter and System x

This chapter discusses the following topics:

- ▶ “IBM i Support for Power Blades” on page 256
- ▶ “Changes to iSCSI-attached integrated servers” on page 258
- ▶ “Changes to integrated Windows servers” on page 259
- ▶ “VMware ESX Server support” on page 261
- ▶ “Changes to integrated Linux servers” on page 261
- ▶ “IXS and IXA-attached integrated Windows servers” on page 262
- ▶ “Linux running on IXS and IXA-attached hardware” on page 263

10.1 IBM i Support for Power Blades

This section provides an overview for the available Power Blades supported by IBM i.

10.1.1 BladeCenter JS12, JS22, JS23, and JS43

IBM i 6.1 and later support the BladeCenter JS12, JS22, JS23, and JS43 Power Blades. The Blades are supported in both the BladeCenter H Chassis and the BladeCenter S Chassis.

10.1.2 IBM client to VIOS on supported BladeCenter systems

IBM i supported BladeCenters use PowerVM Standard Edition, which includes:

- ▶ PowerVM Hypervisor
- ▶ Micro-Partitioning
- ▶ Virtual I/O Server (VIOS) with Integrated Virtualization Manager (IVM)
- ▶ Dynamic Logical Partitioning (LPAR)

The IBM i operating system uses the VIOS partition for access to BladeCenter resources and attached storage devices. The VIOS partition owns the Fibre Channel, Ethernet, and SAS adapters on the blade and virtualizes the resources to the IBM i partition.

For additional information, see IBM i integration with BladeCenter and System x at:

<http://www.ibm.com/systems/i/advantages/integratedserver/library.html>

10.1.3 IBM BladeCenter JS12

The JS12 Power blade is a 2-core blade server based on the POWER6 processor. Its hardware is very similar to that of the JS22, with several important differences:

- ▶ There is a single POWER6 socket with two processor cores.
- ▶ The processor cores operate at 3.8 GHz instead of the JS22's 4 GHz.
- ▶ Two integrated SAS drives are supported on the blade.
- ▶ Four additional memory DIMM slots are supported, for a total of eight slots.

The JS12 includes the same SP, embedded Gigabit Ethernet (IVE) ports, CFFv and CFFh I/O expansion slots, and embedded SAS and USB controllers. The JS12 is capable of supporting twice the number of memory DIMMs because of the new, shorter DIMM design, which allows the DIMMs to be plugged in vertically. IBM i on the JS12 is supported in BladeCenter H and BladeCenter S.

10.1.4 IBM BladeCenter JS22

The JS22 Power blade is a 4-core blade server based on the POWER6 processor. The JS22 fits in a standard IBM BladeCenter chassis and has an integrated Service Processor (SP), two Gigabit Ethernet ports, SAS and USB controllers, and a SAS disk drive. The embedded Ethernet ports are Integrated Virtual Ethernet (IVE) ports, also present on other POWER6-based servers.

Additional I/O is provided by CFFh and CFFv expansion cards, which allow connections to external storage and tape via switches in the BladeCenter chassis. IBM i on the JS22 is supported in BladeCenter H and BladeCenter S.

10.1.5 IBM BladeCenter JS23 Express

The IBM BladeCenter JS23 Express is a single-wide blade server based on the POWER6 processor. The JS23 contains two sockets and four POWER6 cores, which use the IBM enhanced 65-nm lithography and operate at 4.2 GHz. The JS23 also includes 32 MB of shared L3 cache per socket and 4 MB of dedicated L2 cache per core. Up to 8 DDR2 memory DIMMs are supported for a maximum of 64 GB. The blade server includes two embedded 1-Gb Ethernet ports, onboard SAS and USB controllers, EnergyScale™ power management and an FSP-1 service processor. A maximum of one onboard Serial-attached SCSI (SAS) or Solid® State Disk (SSD) drive is supported. The JS23 supports two types of I/O expansion adapters: Combination Form Factor – Horizontal (CFFh) and the new Combination I/O Form Factor – Vertical (CIOv). For IBM i, the JS23 is supported in IBM BladeCenter H and BladeCenter S.

10.1.6 10.1.6. IBM BladeCenter JS43 Express

The IBM BladeCenter JS43 Express is a double-wide blade server based on the POWER6 processor. The JS43 consists of the JS23 main blade unit and a Symmetric Multiprocessing (SMP) expansion unit, which occupies an adjacent blade bay in an IBM BladeCenter. The JS43 has the following hardware characteristics:

- ▶ Four sockets and eight POWER6 cores operating at 4.2 GHz
- ▶ 32 MB of L3 cache per socket and 4 MB of L2 cache per core
- ▶ Up to 16 DDR2 memory DIMMs for a maximum of 128 GB
- ▶ Four embedded 1-Gb Ethernet ports
- ▶ Onboard SAS and USB controllers, EnergyScale power management and FSP-1 service processor on the main blade unit
- ▶ Up to two SAS or SSD onboard drives
- ▶ Up to two CFFh and two CIOv I/O expansion adapters

For IBM i, the JS43 is supported in BladeCenter H and BladeCenter S.

10.1.7 Overview of backup and restore for IBM i on Power blade

With the April 2009 announcements, IBM greatly simplified the backup and restore process for IBM i on a Power blade, making it similar to that on IBM Power rack servers. VIOS can now virtualize the TS2240 SAS-attached tape drive to IBM i, as VIOS has previously done with disk, optical and network devices. This new functionality enables an administrator to perform backups to SAS-attached LTO4 tape directly from IBM i, using native save and restore commands or BRMS.

The virtualized tape drive is not physically attached to IBM i. The TS2240 is attached to a SAS CFFv or CIOv adapter in VIOS through a SAS I/O module in the BladeCenter. Once the tape drive is recognized by VIOS, it is assigned an rmtX device name, using rmt0 for the first one. The rmtX device is virtualized to IBM i by mapping it to a separate Virtual SCSI server adapter, which is in turn connected to a Virtual SCSI client adapter in the IBM i partition. Once available in IBM i, the tape drive is recognized as a 3580 model 004 tape device and is

assigned a TAPxx device description. This TAPxx device can then be used for all types of backup operations using native IBM i save and restore commands or BRMS.

Note: We strongly recommend that both BladeCenter and IBM i experts be involved in getting IBM i up and running as a partition on a BladeCenter configuration. IBM Systems Lab Services and Training is available to assist with both solution design and implementation at:

<https://www.ibm.com/systems/services/labservices/platforms/>

You should review Chapter 14 of the *IBM i 6.1 System i Performance Capabilities Reference* manual for performance tips regarding IBM i on Power blades. The same chapter of the manual also contains performance information for IBM i hosting IBM i and VIOS hosting IBM i on Power servers.

As part of your workload estimation process you should also obtain and use the following:

- ▶ IBM Systems Workload Estimator (WLE) Web-based sizing tool to size the processor and memory requirements for IBM i on Power blades.
- ▶ IBM Disk Magic to size the external disk component of performance. Make sure to select IBM i as the host operating system for the storage subsystem and not AIX or VIOS. While IBM i on Power blades does not attach directly to storage, the tool will correctly assume VIOS is involved.

Note: Disk Magic is available for IBM and IBM Business Partner use. To get a copy of Disk Magic, go to the System Sales or Partnerworld Web site and Search on Disk Magic.

Refer to the following for further technical and implementation resources for IBM i on Blades:

- ▶ Overview of IBM i on Blades:
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ Detailed description of functions and requirements for IBM i on Blades:
http://www.ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf
- ▶ Services voucher for IBM i on Power blades:
<http://www.ibm.com/systems/power/hardware/vouchers/index.html>

10.2 Changes to iSCSI-attached integrated servers

This section describes some of the changes to iSCSI-attached integrated servers.

10.2.1 Configuring hardware resource names

The hardware resource name is now configured by specifying the Network Server Host Port resource name, which is in the form CMN*nn* by default.

Note: For IBM i 5.4, the resource names were configured in the form LIN*nn*. If you are upgrading from IBM i 5.4 to IBM i 6.1, Network Server Host Adapter (NWSH) device descriptions are not automatically reconfigured. Configure the NWSH to point to the new resource name before you use it.

For information about finding the resource name refer to the IBM i Information Center at the following Web address and select the path:

System i integration → iSCSI attached System x and blade systems → Configuring i5/OS → Preparing for the integrated server operating system installation → Creating an NWSH object for each new System i iSCSI HBA port.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

10.2.2 Shared memory pools

Shared data memory pools are supported on iSCSI-attached integrated servers. Use this function to isolate iSCSI I/O operations from other i5/OS I/O operations.

Note: If you are upgrading from IBM i 5.4, shared data pools replace the private memory pool. Configure shared data memory pools for the integrated servers.

See the following IBM i Information Center path for details on shared data pools:

System i integration → iSCSI attached System x and blade systems → Planning for integrated servers → i5/OS memory requirements.

10.2.3 Challenge Handshake Authentication Protocol

iSCSI-attached integrated servers support both target and initiator Challenge Handshake Authentication Protocol (CHAP).

See the following IBM i Information Center path for details on configuring CHAP:

System i integration → iSCSI attached System x and blade systems → Managing integrated server environments → Managing the iSCSI network for integrated servers → Configuring security between i5/OS and integrated servers → Configuring CHAP for integrated servers.

10.2.4 iSCSI direct connect

iSCSI direct connect allows the target and initiator iSCSI HBAs to connect without a network switch.

10.3 Changes to integrated Windows servers

Windows is supported on BladeCenter and System x models that are attached to the IBM i platform through the iSCSI host bus adapter (iSCSI HBA), on System x models that are attached to the IBM i platform through the Integrated xSeries® Adapter (IXA) and on the Integrated xSeries Server (IXS).

The following Microsoft Windows Server operating system versions are supported on various iSCSI configurations:

- ▶ iSCSI
 - 64-bit x64 versions (full and core installations):
 - Windows Server 2008 Standard

- Windows Server 2008 Enterprise
- Windows Server 2008 Standard without Hyper-V
- Windows Server 2008 Enterprise without Hyper-V
- Windows Web Server 2008
- 32-bit x86 versions:
 - Windows Server 2003 R2 Standard
 - Windows Server 2003 R2 Enterprise
 - Windows Server 2003 Standard plus Service Pack 1 (SP1)
 - Windows Server 2003 Enterprise plus Service Pack 1 (SP1)
 - Windows Server 2003 Web plus Service Pack 1 (SP1)
- ▶ IXA and IXS
 - 32-bit x86 versions:
 - Windows Server 2003 R2 Standard
 - Windows Server 2003 R2 Enterprise
 - Windows Server 2003 Standard
 - Windows Server 2003 Enterprise
 - Windows Server 2003 Web
 - Windows 2000 Server
 - Windows 2000 Advanced server

Windows Server 2008 x64 editions are supported on qualified iSCSI-attached integrated servers. See the following Web link for a detailed list of BladeCenter and System x models supported with iSCSI:

<http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/servermodels/>

10.3.1 Backing up active servers

Support for backing up active iSCSI-attached integrated Windows servers is added.

For more information on iSCSI backups see the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Installing, configuring, and managing Windows → Backing up and recovering integrated Windows servers → Using i5/OS to back up disks for active integrated Windows servers.

Use the Disable User Profile (DSBUSRPRF) value on the Install Windows Server command or the Network Server Description to specify that user profiles will not be disabled on the Windows operating system when they are disabled on the i5/OS operating system. For more information about users and groups refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Integrated server concepts → user and group concepts.

A new value for time synchronization is supported. Specify None to ensure that the integrated server time is never synchronized with the i5/OS time. For more information about time synchronization refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Configuring i5/OS → Configuring time synchronization for integrated servers.

10.4 VMware ESX Server support

IBM i integration with BladeCenter and System x enables businesses to operate heterogeneous environments that include Intel®-compatible servers running VMware ESX Server, in addition to core business applications on the IBM i platform.

VMware ESX Server is the host server for abstracting processor, memory, storage and networking resources of a physical server into multiple virtual machines. The ESX Server is part of the VMware Infrastructure.

VMware ESX Server is supported on iSCSI-attached integrated server hardware.

IBM i 5.4 or IBM i 6.1 is required in order to run VMware ESX Server on BladeCenter blade or System x models that are attached to the IBM i platform through an iSCSI host bus adapter (iSCSI HBA).

For more information about VMware installation, configuration and management refer to the following sources:

- ▶ IBM i IBM i Information Center path:
System i integration → iSCSI attached System x and blade systems → Installing, configuring, and managing VMware ESX Server.
- ▶ *VMware VI3 on BladeCenter and System x Integrated with System i*, SG24-7408.

10.5 Changes to integrated Linux servers

IBM i 5.4 or IBM i 6.1 is required in order to run Linux on BladeCenter blade or System x models that are attached to the IBM i platform through an iSCSI host bus adapter (iSCSI HBA).

The following enhancements are included in IBM i 6.1:

10.5.1 Supported versions of Linux

SUSE Enterprise Linux Server 10 for AMD64 and Intel EM64T (SLES 10) and Red Hat Enterprise Linux 5 for x86-64 (RHEL 5) are supported on iSCSI-attached integrated Linux servers.

10.5.2 File-level backup

iSCSI-attached Linux servers now support file-level backup from the i5/OS operating system.

For more information about backing up and recovering, refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Installing, configuring, and managing Linux → Backing up and recovering integrated Linux servers.

10.5.3 Backing up active Linux servers

Support is added for backing up active Linux servers.

For more information about backing up storage spaces, refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Installing, configuring, and managing Linux → Backing up and recovering integrated Linux servers → Using i5/OS to back up disks for integrated Linux servers → Backing up storage spaces for active integrated Linux servers.

10.5.4 Time synchronization

A new value for time synchronization is supported. Specify None to ensure that the integrated server time is never synchronized with the IBM i time.

For more information about time synchronization refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Configuring i5/OS → Configuring time synchronization for integrated servers.

10.6 IXS and IXA-attached integrated Windows servers

10.6.1 Disabling user profiles

Use the Disable User Profile (DSBUSRPRF) value on the Install Windows Server command or the Network Server Description to specify that user profiles will not be disabled on the Windows operating system when they are disabled on the IBM i operating system.

For further details on the Disable User Profile command see the IBM i Information Center path:

System i integration → IXS or IXA-attached integrated Windows server_σ → Integrated server concept_σ → User and group concepts for integrated Windows servers.

10.6.2 Time synchronization

A new value for time synchronization is supported. Specify None to ensure that the integrated server time is never synchronized with the IBM i time.

For more information about time synchronization refer to the following IBM i Information Center path:

System i integration → iSCSI attached System x and blade systems → Configuring i5/OS → Configuring time synchronization for integrated servers.

10.6.3 Windows Server installation advisor

The Windows Server Installation Advisor no longer supports IXS and IXA-attached integrated servers. Use the Installation worksheet for IBM i parameters to select parameters for the Install Window Server (INSWNTSVR) command.

The installation worksheet can be found by following the IBM i Information Center path:

System i integration → IXS or IXA-attached integrated Windows servers → Installing and configuring IXS and IXA-attached integrated Windows servers → Planning for the installation of Windows server → Installation worksheet for i5/OS parameters.

10.7 Linux running on IXS and IXA-attached hardware

At IBM i 6.1, integrated Linux servers are supported only on iSCSI-attached hardware. Linux installations are not supported on IXS or IXA hardware.

Note: The operating system will allow migration from IBM i 5.4 to IBM i 6.1 for Linux servers running on IXS or IXS-attached hardware. However, the IXS/IXA integrated server function for Linux will be limited and will not be serviced in IBM i 6.1. For more information, see the *Linux on integrated servers* Web site at:

<http://www.ibm.com/systems/i/advantages/integratedserver/linux/>



IBM Advanced Job Scheduler for i5/OS

This chapter discusses the following topics:

- ▶ “IBM Advanced Job Scheduler for i5/OS” on page 266
- ▶ “Advanced Job Scheduler Web interface” on page 267
- ▶ “Error notification” on page 271
- ▶ “Output queue monitoring” on page 271
- ▶ “Command lists” on page 272
- ▶ “Multiple scheduling environments” on page 273
- ▶ “Report distribution” on page 274
- ▶ “Log” on page 274
- ▶ “Scheduling and running jobs” on page 275

11.1 IBM Advanced Job Scheduler for i5/OS

The IBM Advanced Job Scheduler for i5/OS (5761-JS1) licensed program is a powerful scheduler that allows unattended job processing 24 hours a day, 7 days a week. This scheduling tool provides more calendar features and offers greater control over scheduled events than the Management Central or the IBM i schedulers. You can also view job completion history and manage notification of a job's status.

If you want to schedule jobs on several systems in your network, the product must be installed on each of your systems. If you want to use the Advanced Job Scheduler (AJS) in System i Navigator (and in Management Central), then you must install the client plug-in from a system that has the Advanced Job Scheduler installed.

However, it is not necessary to install the Advanced Job Scheduler licensed program on each endpoint system in your Management Central network. When you install the Advanced Job Scheduler on the central system, jobs or tasks that you define on an endpoint system gather job information that is needed from the central system. You must set up all job definition information on the central system.

There are five interfaces to Advanced Job Scheduler functions:

- ▶ A 5250 workstation interface
 - To get to the main product menu on a 5250 workstation, type GO JS on a command line.
- ▶ CL commands that can be run in a batch environment
- ▶ A graphical interface through System i Navigator
- ▶ A wireless interface through Advanced Job Scheduler for Wireless
- ▶ A Web interface through System i Navigator for the Web
- ▶ IBM Systems Director Navigator for i

If systems in your network have the Advanced Job Scheduler installed locally, you can schedule tasks outside of the Management Central network. Under My Connections in System i Navigator, you have access to the Advanced Job Scheduler on that local system when you expand Work Management, as shown in Figure 11-1 on page 267.

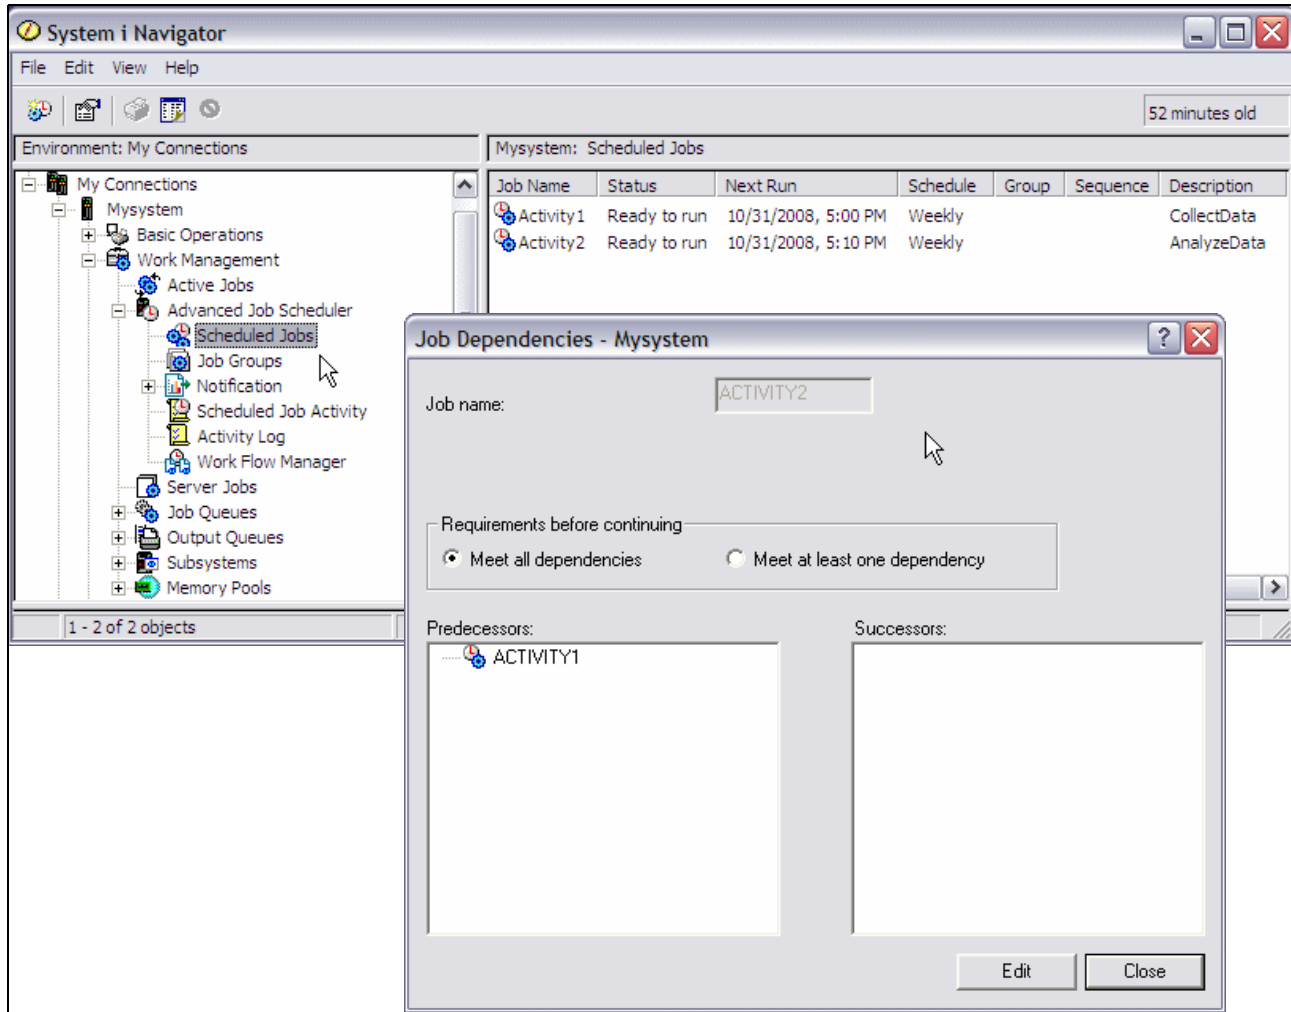


Figure 11-1 Advanced Job Scheduler

11.2 Advanced Job Scheduler Web interface

Using IBM Systems Director Navigator for i, you can remotely access scheduled jobs, job groups, scheduled job activity, and activity log information for AJS. All of these functions can be used without installing System i Access for Windows. Everything you need to access these tasks is already installed on your IBM i partition. When listing scheduled jobs and scheduled job activity, any column can be used for sorting or filtering. This is great for selecting jobs that begin with certain letters and then sorting by the next scheduled date and time. While you cannot make changes to jobs using the IBM Systems Director Navigator for i interface, there are many ways that you can manage your scheduled jobs.

Figure 11-2 on page 268 shows the main Advanced Job Scheduler functions under IBM Systems Director Navigator for i, after selecting **Work Management** → **Show All Work Management Tasks**.

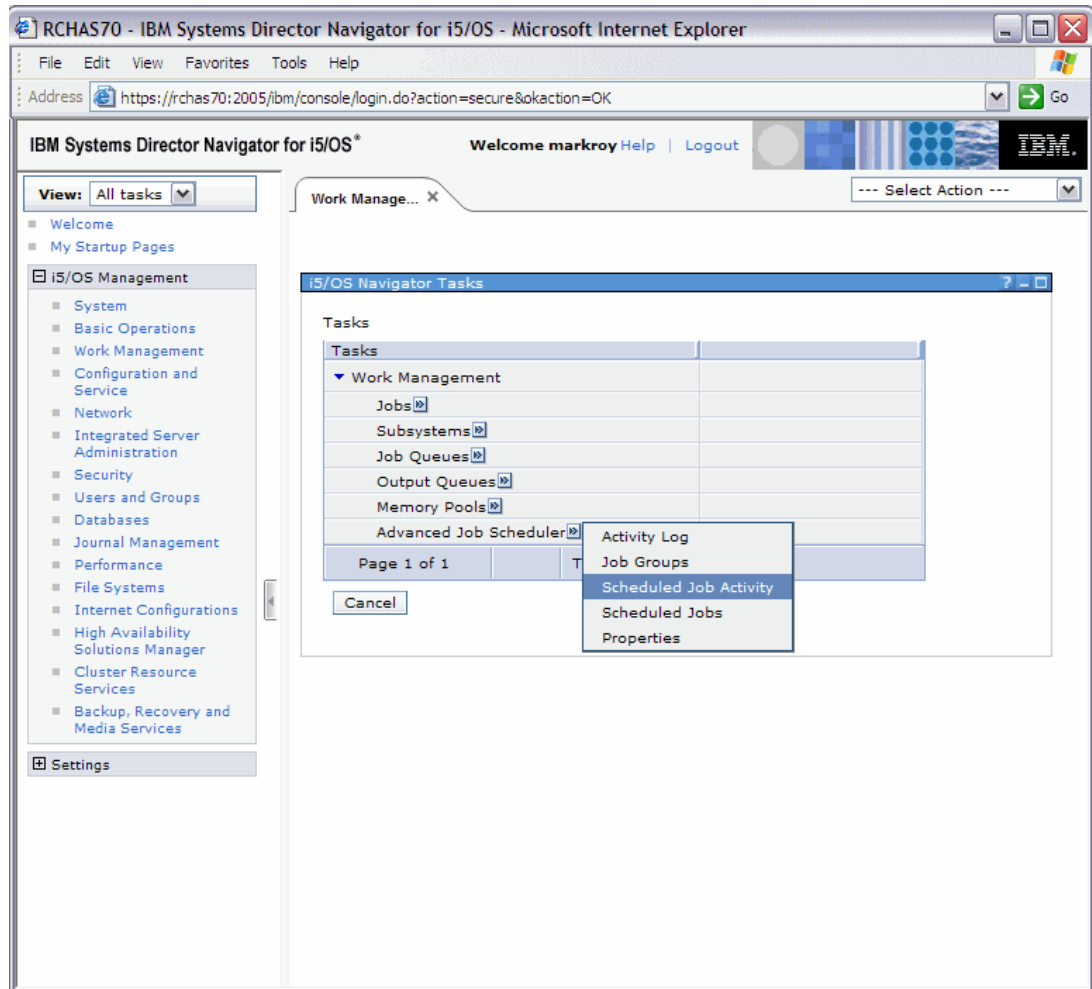


Figure 11-2 IBM Systems Director Navigator for i Advanced Job Scheduler interface

You can do the following for a scheduled job:

- ▶ **Properties**
This displays all the attributes that make up a scheduled job. You cannot make changes to jobs this release.
- ▶ **Hold/Release Schedule**
With this option you can hold or release a scheduled job.
- ▶ **Activity**
This lists each time the selected job has run.
- ▶ **Activity Log**
This displays the activity log for the selected job.
- ▶ **Run**
This allows you to run the job immediately. The completion of the job will not affect any dependencies. All resource dependencies must be met in order for the job to be submitted. A confirmation is required.

- ▶ Skip Next Run

Using this option will skip the scheduled run of the job. Using this option multiple times will continue to skip the next run. To reset the job back to the original schedule, use the Hold and Release options.

- ▶ Status

This displays the status information of a job while it is running.

- ▶ Job Dependencies

This provides a list of the predecessor and successor jobs for the selected job.

- ▶ Delete

This deletes the selected job. A confirmation is required.

Figure 11-3 shows the actions you can take in the Advanced Job Scheduler using the IBM Systems Director Navigator for i interface.

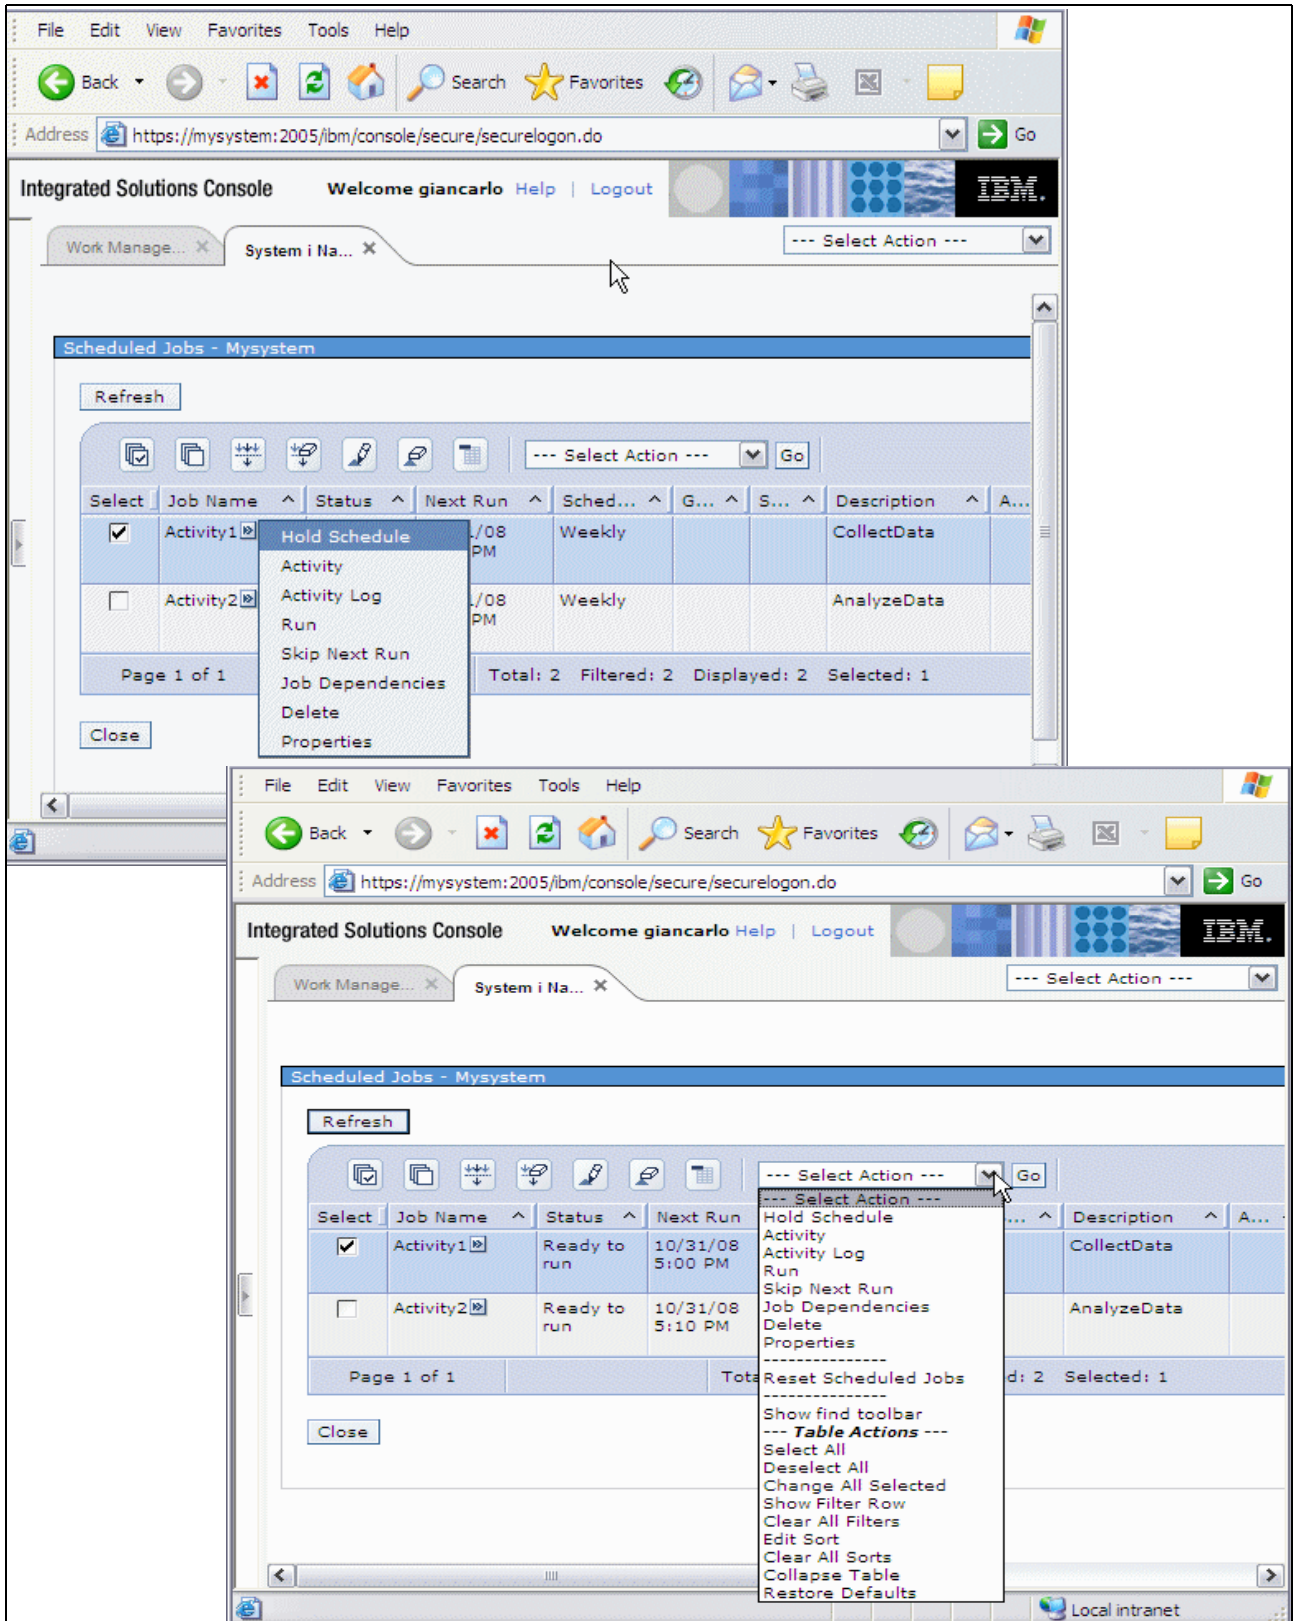


Figure 11-3 Web interface

Subsequent sections of this chapter describe the new functions and enhancements for IBM Advanced Job Scheduler for i5/OS.

Note: The new functions and enhancements are not available in the IBM Systems Director Navigator for i interface.

11.3 Error notification

AJS is enhanced with the ability to send a notification message when an inquiry message is sent to a scheduled job. These types of messages cause the job to wait for a reply or end immediately because of a default reply. For example, assume a job receives an inquiry message that a file has reached its maximum size. The notification recipient would receive an e-mail message with one of the following:

Use error message	The text of the error message is sent.
Send job log	The job log of the job will be embedded in the body of the message.
Message text	The user-defined text entered in the Message field is sent.

You can access the Error notification fields on the Notification tab of a scheduled job.

Note: The notification will occur the next time the job monitor wakes up, which is normally no more than five minutes. Error notification does not occur for remote jobs.

11.4 Output queue monitoring

You can now set up the Output Queue Monitor feature to monitor for any spooled file that is created in a specific output queue. Each Output Queue Monitor is defined with an output queue and a Report Distribution List. After a spooled file is created on the output queue, the Report Distribution List associated with the monitor is used to determine how the spooled file is to be distributed. A spooled file must be created with a Ready status in order for the monitor to attempt distribution. Also holding, then releasing, a spooled file that is in a monitored output queue will start the distribution function again. You can access Output Queue Monitors lists by selecting **Advanced Job Scheduler** → **Notification** → **Output Queue Monitors** from within System i Navigator, as shown in Figure 11-4 on page 272.

Note: In order for an Output Queue Monitor to distribute spooled files, the job monitor for the QUSRIJS data library (scheduling environment) must be active, as well as the job monitor for the data library that contains the Output Queue Monitor.

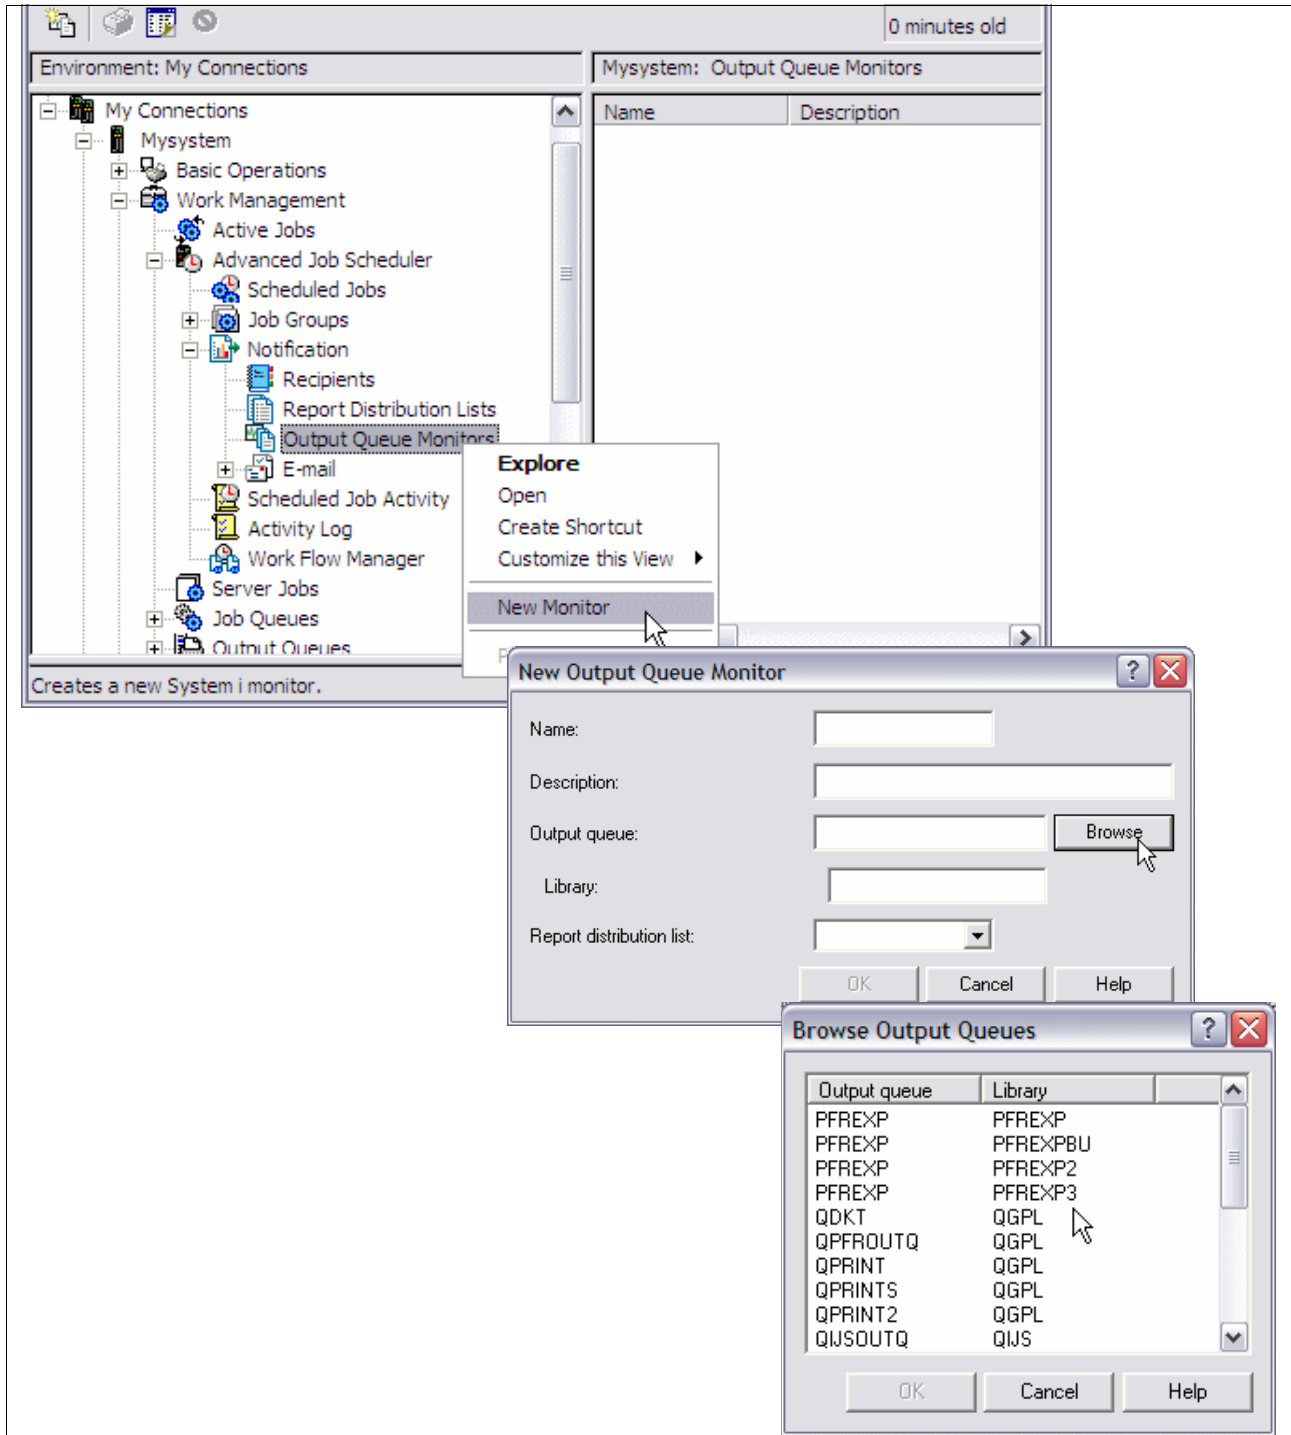


Figure 11-4 Create a new Output Queue Monitor

11.5 Command lists

There are several new features in the command list function of a scheduled job.

Specific commands in the command list can be bypassed. This can be helpful in case a command is temporarily not needed but you do not want to remove it from the list. From the

properties of a scheduled job using the System i Navigator for Windows, you can de-select the check box that enables a command for processing. This will cause the command to be bypassed the next time the job is run.

The maximum command length is increased from 512 to 4096.

Commands can run on non-IBM i operating systems, such as Windows. They can be processed on a single system or a group of systems. When processing commands on a group of systems, each command must complete normally on each system within the group before the next command is processed. The Run Remote Command (RUNRMTCMD) command is used to issue the commands on the remote systems. The RUNRMTCMD command allows IBM i users to run a command on a remote system that is running the target portion of this function. The target portion of this function can be a rexecd (remote executing daemon). The Incoming Remote Command (IRC) service of the IBM System i Access for Windows is recommended on the remote systems to process the commands.

Important: Commands that are processed on non-IBM i operating systems may report that the command ended normally even though the command did not complete normally. This can occur when the remote executing daemon, rexecd, being used does not return the appropriate error code information to the RUNRMTCMD. The process will produce a spooled file with information about the process that may help to determine whether the command completed normally. The spooled file cannot be used by the job scheduler to determine the completion of the job because of translation issues and inconsistencies.

In order to take advantage of this new feature, it may be helpful to run through a quick test outside of the Advanced Job Scheduler. For information about how to set up and execute a test, along with information about setting up AJS to run remote commands on other operating systems, refer to *Advanced Job Scheduler Command List Changes* Web page at:

http://www.ibm.com/systems/i/software/jscheduler/command_list.html

11.6 Multiple scheduling environments

With *multiple scheduling environments* you can set up any number of scheduling environments running at the same time. This is extremely helpful for having a test and production set of scheduled jobs or when replicating live data from a production system to a backup system. Refer to *Setting up multiple scheduling environments* in the IBM i Information Center, selecting the following path: **Systems management** → **Work Management** → **Managing work** → **Managing job scheduling** → **Advanced Job Scheduler** → **Scheduling jobs with Advanced Job Scheduler** → **Setting up the Advanced Job Scheduler** → **Setting up multiple scheduling environments**.

When adding a new environment to AJS, you can select to create the data library. This will create a new clean data library. This option may not be necessary when restoring or replicating data from another system. If this feature is not used, the data library must exist before adding the environment to AJS.

There are two new Command Language (CL) commands that are helpful during the switch-over from a production system to a backup system. The Change Data Library using JS (CHGDLJS) command allows you to change the monitor designator for a data library. The monitor designator determines whether a data library (scheduling environment) can be monitored. A Scheduling Environment must be monitored in order for scheduled jobs to submit automatically. The Set Data Library using JS (SETDLJS) command allows you to change the data library that is associated with a user profile. The data library associated with

a user is used when the user accesses the product. A user can only access one data library at a time.

The Work with Jobs using JS (WRKJOBJS) and Work with History using JS (WRKHSTJS) commands have a new data library parameter in order to access a different environment. The data library can also be changed from the main General Properties tab for AJS by using the drop-down arrow to select a different library. When a new data library is specified or selected, it will stay in effect for that user until a different data library is specified or selected.

A new field was added to the definition of a data library to determine whether to auto-start the monitor after an IPL. This is helpful when you want to be in control of when a scheduling environment is to be started.

The job monitor for a scheduling environment can have its own name. The default name is QIJSSCD.

11.7 Report distribution

The report distribution function for AJS allows you to automatically distribute spooled files to one or more e-mail addresses as a PDF or TXT attachment. It can also distribute spooled files to one or more output queues. Report distribution has been enhanced with the following:

- ▶ Report distribution has the capability to turn off producing the report distribution report. This is a new check box found on the Properties of the Notification folder.
- ▶ Several replacement variables have been added to the subject and message of an e-mail with regard to the spooled file that is being distributed. The value of the variable is replaced when the e-mail is produced. These variables will only be replaced during report distribution. To see a complete list of replacement variables, with an explanation of the value that will be replaced, refer to the Advanced Job Scheduler Report Distribution Web page at:

http://www.ibm.com/systems/i/software/jscheduler/report_distribution.html

11.8 Log

The following are the enhancements to the display format of the Display Log for Job Scheduler (DSPLOGJS) command:

- ▶ When the detail of a message is needed, F4 is now used instead of F1. F1 is used to display Help text.
- ▶ The date and time of the log entry is now part of the list.
- ▶ The F11 key provides an alternative view that includes the message ID of the entry and the job that produced the entry.
- ▶ Milliseconds have been added in order to more accurately sequence the entries.
- ▶ The number of message IDs that can be specified on the DSPLOGJS command has been increased to 10.

11.9 Scheduling and running jobs

11.9.1 Scheduled jobs properties

There are two new scheduled jobs properties that affect how jobs are reset. You can access these by right-clicking **Scheduled Jobs** and selecting **Properties**.

The new “Reset held jobs” check box is used to determine whether held jobs should have a next schedule date and time even though they are held. In the past, held jobs had a scheduled date and time. When the scheduled date and time had been reached it would be logged that the job did not run because it was held. Then the job was set with the next scheduled date and time. This meant that the job monitor might wake up more often than necessary. If the “Reset held jobs” box is not checked, the held job will not contain a scheduled date and time.

Note: This will affect the next time a held job has reached its scheduled date and time.

Another new field is the “Reset after change” check box. In the past, clicking **OK** after displaying the properties of a scheduled job would result in the schedule date and time being reset (recalculated), even though the schedule had not been changed. Now if the “Reset after change” check box is not checked, the schedule of a job has to change before the job is reset.

Figure 11-5 shows the new options in the scheduled jobs properties.

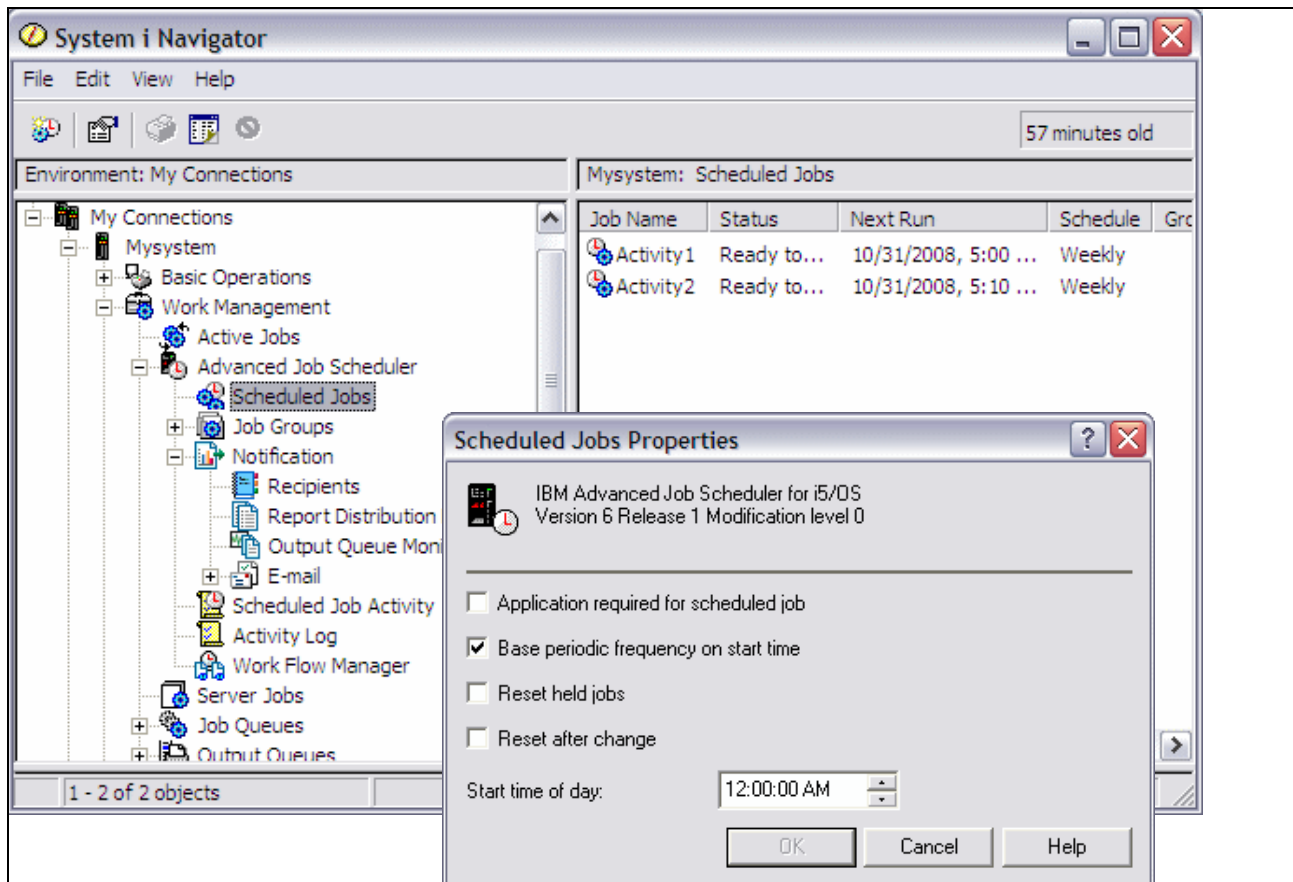


Figure 11-5 Scheduled Jobs Properties

11.9.2 Day of week option for subordinate jobs

On the System i Navigator interface, the schedule of a subordinate job in a group was enhanced to allow days of the week to be selected. This is so that you can define a group to run every day of the week but certain jobs within the group to only run on one or more days of the week.

11.9.3 Start Group using Job Scheduler (STRGRPJS)

The Start Group using JS (STRGRPJS) command has a new parameter called Based on day of week. For instance, if *MON is specified, the only group jobs that will run after the initial job will be those that contain Monday in the list of days of the week to process.

11.9.4 Print Schedule using Job Scheduler (PRTSCDJS)

The Application parameter was added to the Print Schedule using JS (PRTSCDJS) command. This enables you to select a specific application for reporting the schedule.



Web serving

This chapter provides an overview of the changes made to the HTTP Server for IBM i (5761-DG1) product in IBM i 6.1.

- ▶ “Apache Software Foundation Tomcat servlet engine” on page 278
- ▶ “Large file support added” on page 278
- ▶ “IPv6 support” on page 278
- ▶ “New module mod_proxy_balancer” on page 279
- ▶ “New module mod_filter” on page 280
- ▶ “Integrated Web application server” on page 280
- ▶ “New Web services server” on page 282

For more detailed information about these changes, refer to the IBM i Information Center and select the path **Networking** → **HTTP Server** → **What’s new for V6R1**.

12.1 Apache Software Foundation Tomcat servlet engine

The Apache Software Foundation (AFS®) Jakarta Tomcat servlet engine is no longer supported or shipped with the operating system. It has been replaced with the integrated Web application server. For more information about the integrated Web application server, refer to 12.6, “Integrated Web application server” on page 280.

12.2 Large file support added

Apache 2.2.x now contains support for files larger than 2 GB. Support for handling greater than 2 GB request bodies has also been added.

12.3 IPv6 support

The HTTP Server for IBM i supports running over either IPv4 or IPv6.

IP version 6 (IPv6) is supported by the underlying Apache Portable Runtime (APR) library. Apache gets IPv6 listening sockets by default. Additionally, the Listen, NameVirtualHost, VirtualHost, and other directives support IPv6 numeric address strings. For example, Listen [2001:db8::1]:8080 will identify their IP numeric address parameters.

Fast Response Cache Accelerator (FRCA) and Proxy support for IPv6 are also available.

12.4 New module mod_proxy_balancer

The proxy balancer module has been added to provide load balancing services for the main proxy module (mod_proxy), shown in Figure 12-1. The load balancer enables requests to be shared among workers by two methods, request counting and weighted traffic counting. Request counting just counts the number of requests and distributes requests across workers until they have each served an equal number of requests.

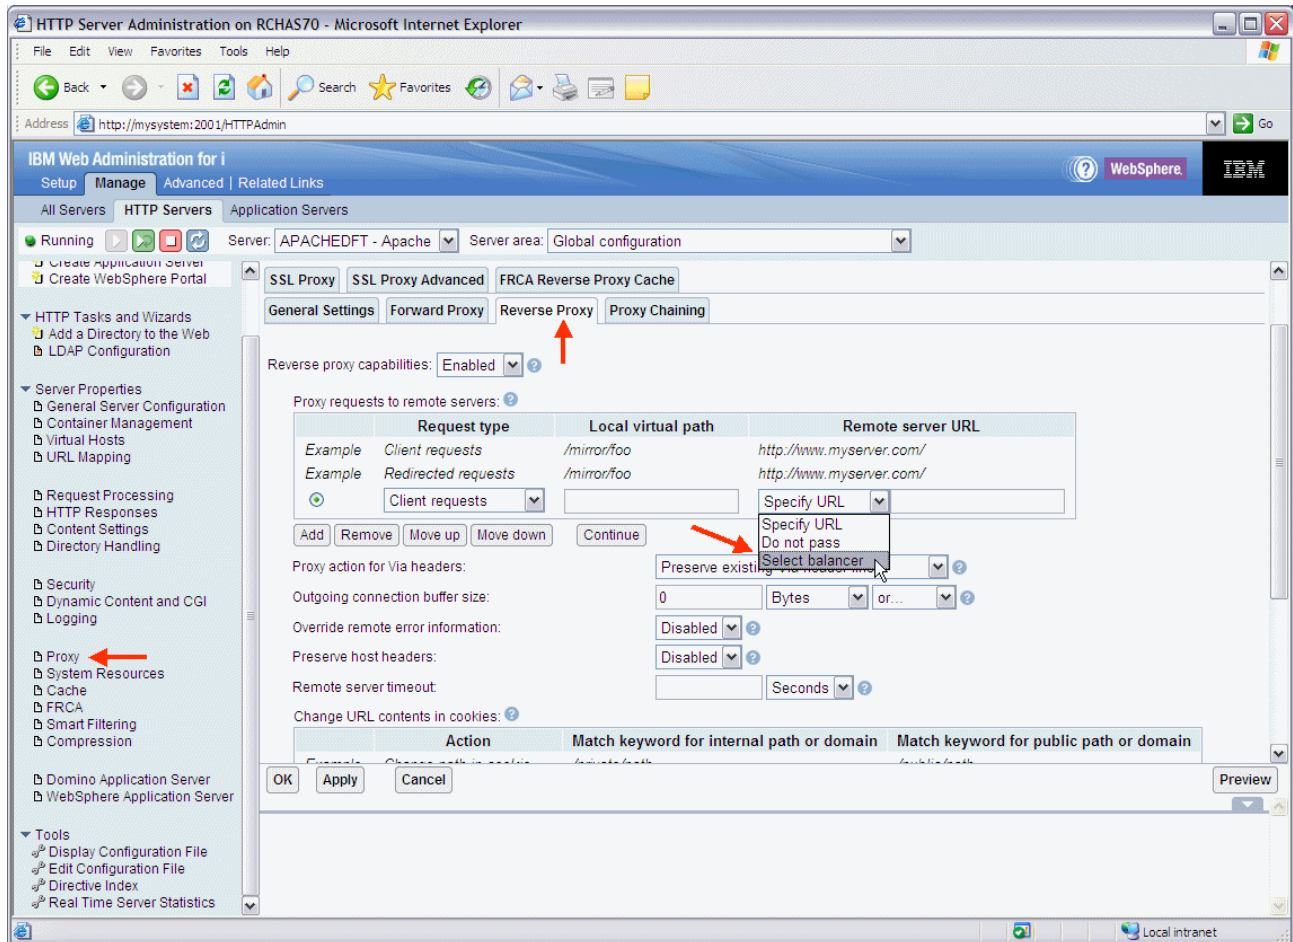


Figure 12-1 Proxy balancer

12.5 New module mod_filter

The filter module provides context-sensitive configuration of output content filters. This module gives server administrators a great deal of flexibility in configuring the filter chain. In fact, filters can be inserted based on any Request Header, Response Header, or Environment Variable, as shown in figure Figure 12-2. The filter module has been extended to allow filters to be executed based on conditional criteria. This changes the previous implementation model under which documents were merely filtered unconditionally according to the configuration of the AddOutputFilter directive or the minor flexibility offered by AddOutputFilterByType.

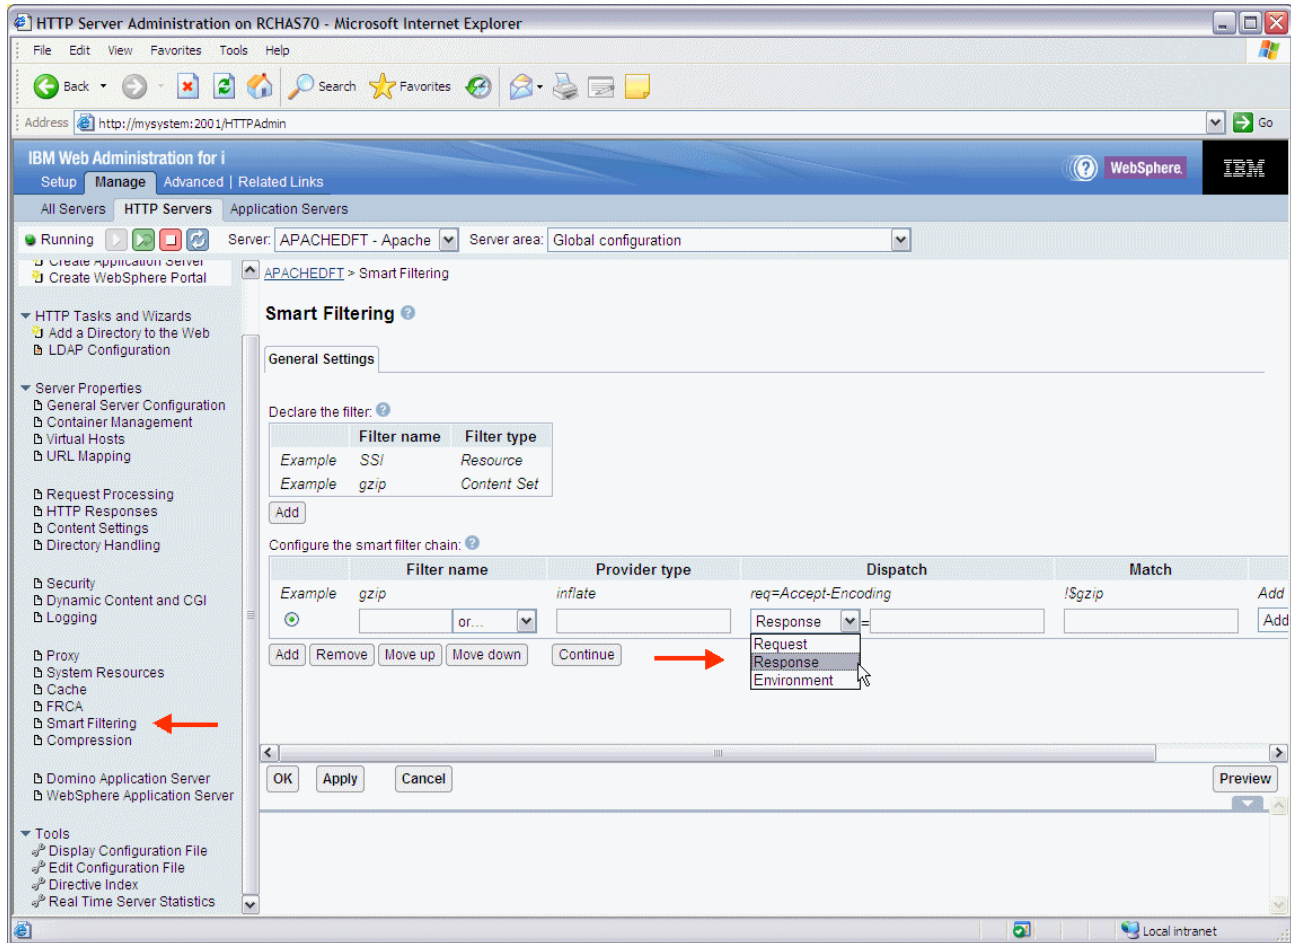


Figure 12-2 Filter module

12.6 Integrated Web application server

The IBM i integrated Web application server provides a minimal footprint, is easy to configure and includes a secure infrastructure for hosting servlet-based Web applications. Instances of this server can be created and managed through the IBM Web Administration for i. This Web container is the first step in getting your programs on the Web. It is a great place to prototype applications or run applications that do not require the scalability of WebSphere Application Server.

One of the main differences between the integrated Web application server and WebSphere Application Server is the format of the applications. Applications that run in WebSphere Application Server are in the EAR or WAR format. Applications for the integrated Web application server need to be in the Web Application Bundle (WAB) format. The WAB architecture consist of the jar files for the application and the manifest file that details all the class files that are used by the application. A WAB file is typically created as a Java Archive (JAR) file with the .jar file extension. Using the IBM Web Administration for i interface to deploy an application that is in a WAR file will deploy without issue. The interface will convert the WAR file into a WAB file before it is deployed to the Web server.

IBM i provides one instance already defined and started as part of the ADMIN HTTP server, as shown in Figure 12-3.

```

                                Work with Subsystem Jobs                                RCHAS70
                                                                                          11/05/08 17:41:30
Subsystem . . . . . : QHTTPSVR

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

Opt  Job          User          Type          -----Status-----  Function
-----
ADMIN  QLWISVR       QLWISVR       BATCHI        ACTIVE                  JVM-com.ibm.lw
ADMIN  QTMHHTTP      QTMHHTTP      BATCH         ACTIVE                  PGM-QZHBMAIN
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRLOG
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRHTTP
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRCGI
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRCGI
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRCGI
ADMIN  QTMHHTTP      QTMHHTTP      BATCHI        ACTIVE                  PGM-QZSRCGI
ADMIN2 QLWISVR       QLWISVR       BATCHI        ACTIVE                  JVM-com.ibm.lw
ADMIN3 QLWISVR       QLWISVR       BATCHI        ACTIVE                  JVM-com.ibm.lw
ADMIN4 QLWISVR       QLWISVR       BATCHI        ACTIVE                  JVM-com.ibm.lw

```

Figure 12-3 HTTP Admin server jobs

The “lightweight application server” architecture server is built on the OSGI/Eclipse platform and shares services and components from multiple organizations across IBM. This server shares a close affinity to the Lotus® Expeditor product and can be considered a “customized” configuration of it.

The IBM Integrated Web Application Server for i Web site includes a Frequently Asked Questions (FAQ) page. This page answers frequently asked questions about the Integrated Web Application Server for i. The FAQ page answers questions regarding the features of the product as well as comparisons to WebSphere Application Server, WebSphere Application Server Network Deployment (ND), and Apache Software Foundation (ASF) Jakarta Tomcat servlet engine. The *IBM Integrated Web Application Server for i Frequently Asked Questions (FAQ)* can be accessed at:

<http://www.ibm.com/systems/i/software/ias/faq.html>

12.7 New Web services server

The Web services server provides a convenient way to externalize existing programs running on IBM i, such as RPG and COBOL programs, as Web services. Web services clients can then interact with these IBM i program-based services from the Internet using industry standard Web service communication protocols such as Simple Object Access Protocol (SOAP). Instances of this server can be created and managed through the IBM Web Administration for i5/OS interface.

From the service provider perspective, the key element to externalizing a RPG or COBOL program as a Web Service is the new Web services server. This server is an integrated Web application server with an embedded Web services engine. The key element in allowing Integrated Language Environment® (ILE) programs, including those written in RPG and Cobol, to invoke a Web service, is the Web services client for ILE, which enables the generation of C and C++ Web service client proxies or stubs that can be invoked by ILE programs.



Web enablers

This chapter discusses the following topics:

- ▶ “Integration with the Integrated Web application server” on page 284
- ▶ “Web Performance Management tools” on page 285
- ▶ “WebSphere Application Server Express V7.0” on page 305
- ▶ “WebSphere Application Server performance data” on page 312
- ▶ “Simplified use of SSL with HTTP Administration server” on page 312

13.1 Integration with the Integrated Web application server

System i Navigator tasks on the Web and System i Access for Web are enabled to use the Web application server integrated with IBM i.

This Web application server is started when the IBM-supplied HTTP Administration server, IBM Web Administration for IBM i, is started. This HTTP Administration server is the one started by the Start TCP Server (STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)) command.

This support of the new integrated Web application server gives you the ability to be productive using the IBM i Access for Web functions and the System i Navigator tasks on the Web functions without having to install and configure a WebSphere Application Server product or an Apache server on IBM i.

Note: The integrated Web application server is also available under IBM i 5.4 with a cumulative PTF level made available in 4Q 2007. It is documented in the IBM i 5.4 Information Center. This publication, however, focuses on the IBM i 6.1 capabilities.

This application server supports user-defined instances that can have some basic Web applications deployed to it. IBM i 5.4 and 6.1 provide an already defined instance of an integrated Web application server that runs under the Administration HTTP server (*ADMIN). The following IBM-supplied Web applications use or can use this IBM *ADMIN integrated Web application server for IBM i 6.1:

- ▶ IBM Systems Director Navigator for i

For more information on this, refer to Chapter 16, “IBM Systems Director Navigator for i” on page 403.

- ▶ System i Navigator tasks on the Web

For more information, refer to 17.3, “System i Access for Web” on page 442.

Note, this support was included in IBM i 5.4 and remains available in IBM i 6.1. However, if you are not already using these functions on IBM i 6.1, we recommend using IBM Systems Director Navigator for i.

- ▶ System i Access for Web

For more information, refer to 17.3, “System i Access for Web” on page 442 and the IBM i Information Center, selecting the path **Connecting to System i** → **Choose your connection** → **iSeries Access** → **iSeries Access for Web**.

As previously mentioned, the integrated Web application server runs as a job under the IBM i-supplied HTTP Administration server (*ADMIN). We used the Windows-based System i Navigator path **Work Management** → **Subsystems** → **Active subsystems** → **qhttpsvr** to identify this job under the Function column, as shown in Figure 13-1.

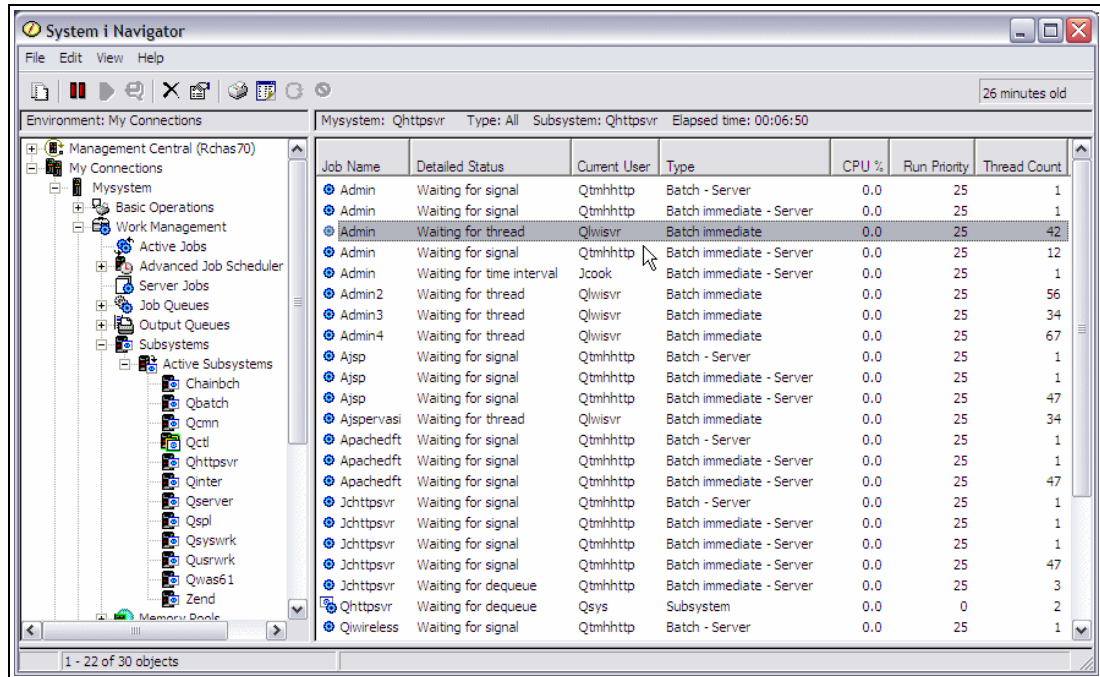


Figure 13-1 Manage all HTTP servers showing the integrated Web application servers

For more information about IBM i 6.1 enhancements for System i Access for Web, refer to 17.3, “System i Access for Web” on page 442.

13.2 Web Performance Management tools

Two new Web Performance Management tools are available:

- ▶ Web Performance Monitor
- ▶ Web Performance Advisor

Note: 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 IBM i version
- ▶ The following PTF groups for the corresponding IBM i 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

WebSphere Application Server 6.1.x, any Portal Server, and any Workplace releases are not supported.

Both of these tools are accessed from the IBM i-provided default HTTP Administration server, commonly referred to as the IBM Web Administration for IBM i. This product has been

available for many IBM i releases and continues to be enhanced with each new IBM i release and frequently within a release cycle.

This Web Administration for i server can be started using either of the following methods:

- ▶ By the Start TCP/IP Server command STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN).
- ▶ By the System i 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` or `http://hostname:2001/HTTPAdmin`

In IBM i 6.1, the browser typically first brings up the IBM Systems Director Navigator for i signon window. After signing on you can click **i5/OS Tasks Page** to get to that page.

Then:

1. On the i5/OS Tasks page, select the **IBM Web Administration for i** 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 **Manage Details** 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 13-2, we have selected **JCAPPSVR**. You see the links to the Web Performance Advisor and Web Performance Monitor.

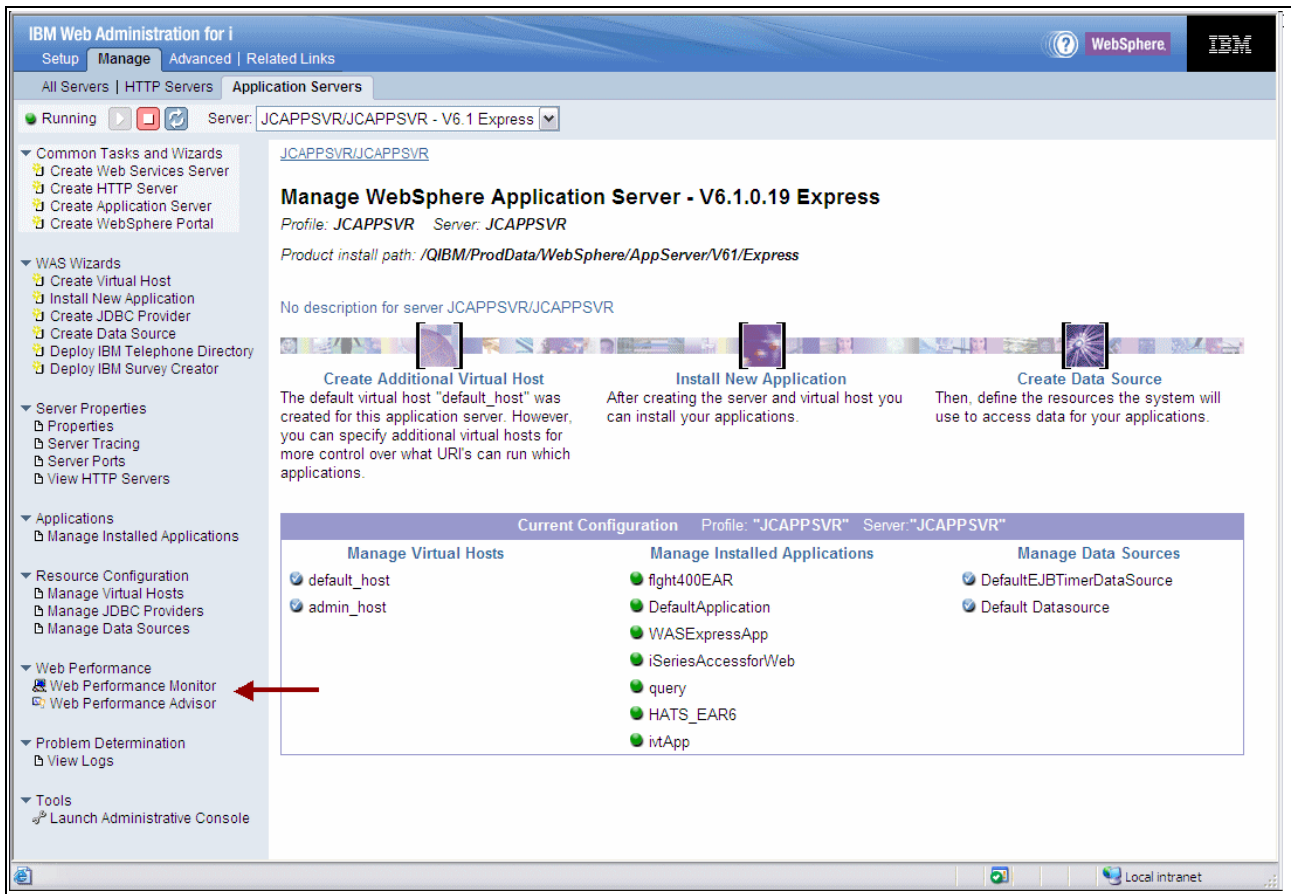


Figure 13-2 Accessing the Web Performance Monitor and Performance Advisor

The following sections provide additional overview information about these two Performance Management tools.

13.2.1 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
- ▶ DB2 for i

ARM is a standard of The Open Group, see

<http://www.opengroup.org>

It is 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 systems (partitions) and operating system environments, whereas the Web Performance Monitor monitors a single IBM i 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 subportions of a transaction, such as within the HTTP server, in the Web applications, and in 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 in 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 System i Navigator tasks on the Web.

Perform the following to run the Web Performance Monitor:

1. Access the Web Administration for i from your browser and select a specific applications server. For that server, select **Manage Details**. 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 13-3.

In our example we use the HTTP server JCHTTPSVR and WebSphere Application Server profile/server JCAPPSVR.

2. Click **Start** and processing begins. An hourglass is shown in front of the current item and items are checked off as they are completed. See Figure 13-3.

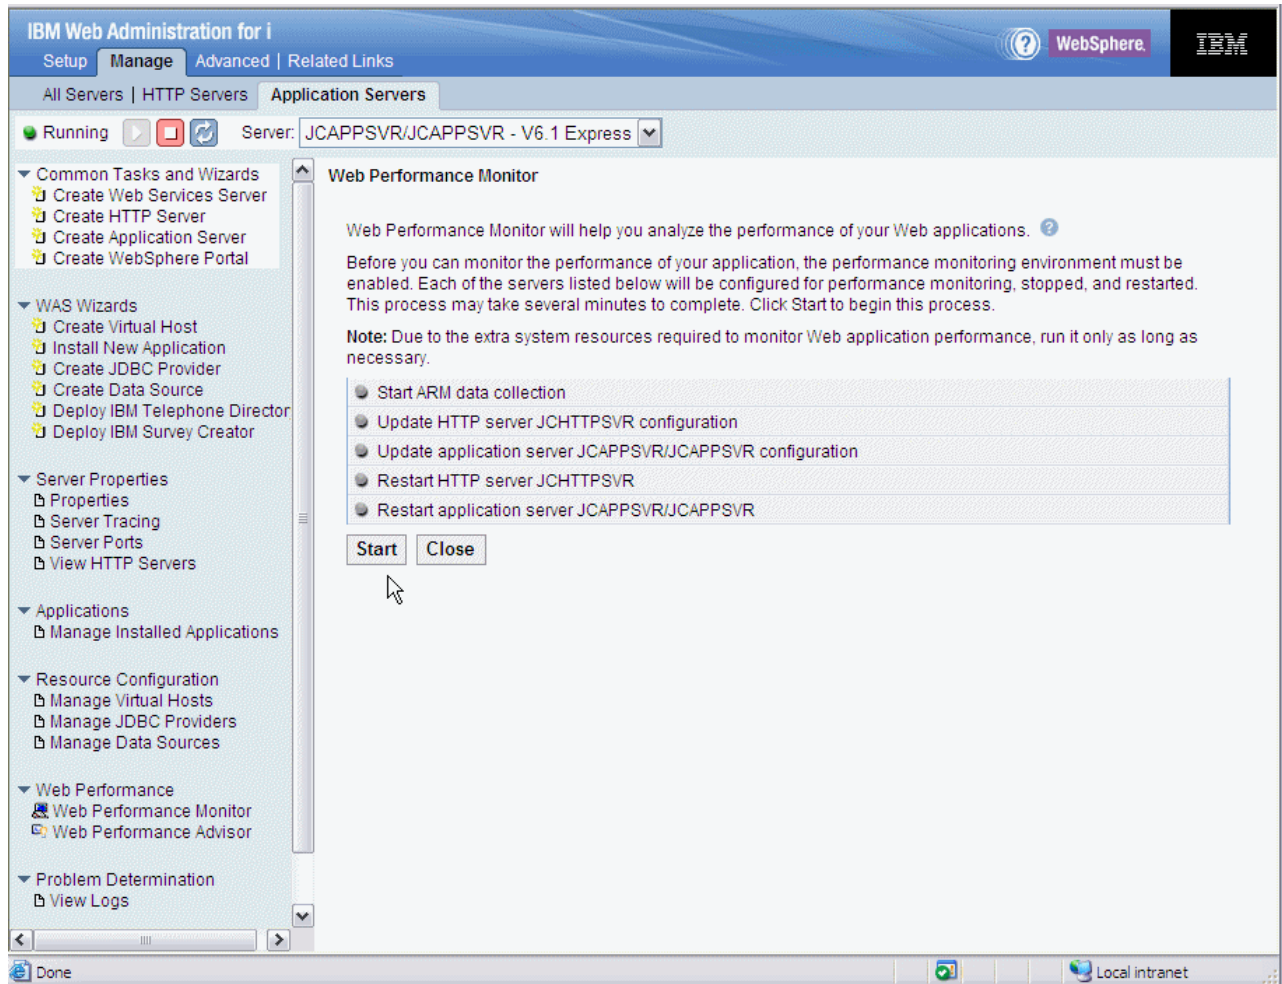


Figure 13-3 Web Performance Monitor - start example

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 IBM i partition. Notice that it shows a WRKACTJOB-like screen, as shown in Figure 13-4.

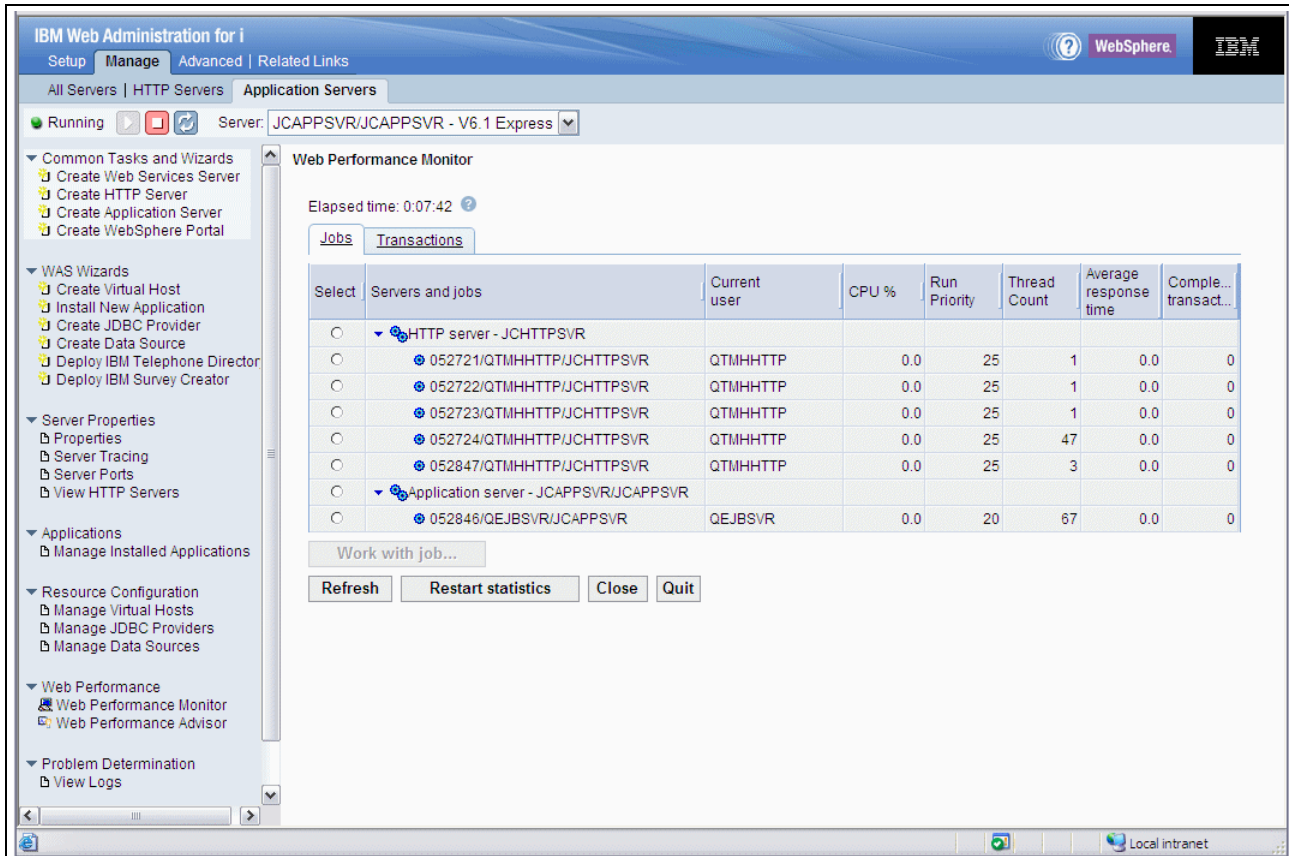


Figure 13-4 Web Performance Monitor - jobs example

If you select a job, the **Work with job** button is activated. If you then select that button, you are launched to the System i Navigator tasks on the Web application, provided that you have started the WebSphere Application Server. See Figure 13-5.

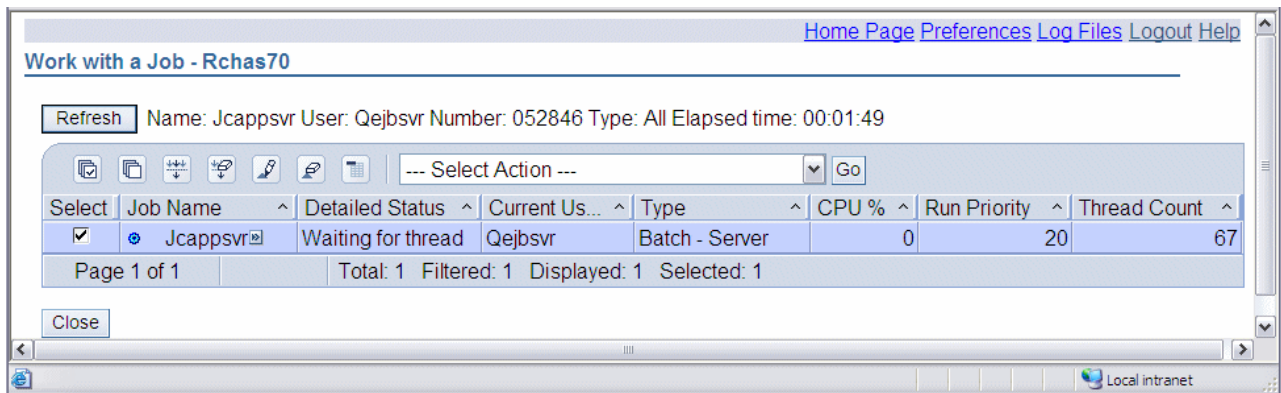


Figure 13-5 Work with job

You can also show a transactions view of the performance information by clicking the **Transactions** tab. In that view you can specify preferences and user filtering to control what

is shown on the panels. For example, show only the activity for selected users, as shown in Figure 13-6.

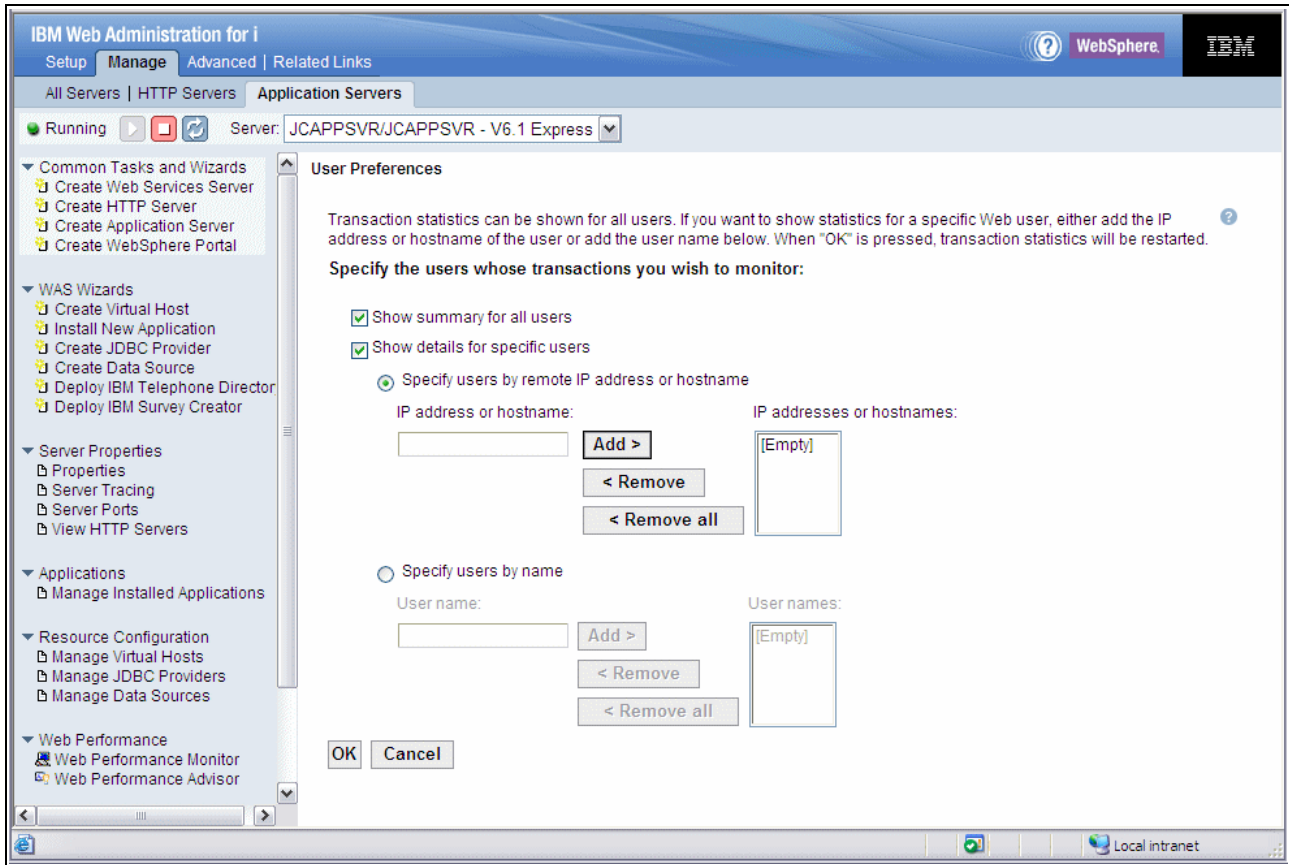


Figure 13-6 User Preferences

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.

To end the Web Performance Monitor, click **Quit** on the Jobs (shown in Figure 13-4 on page 290) or the Transactions window. This brings up the window shown in Figure 13-7. Click **Stop**.

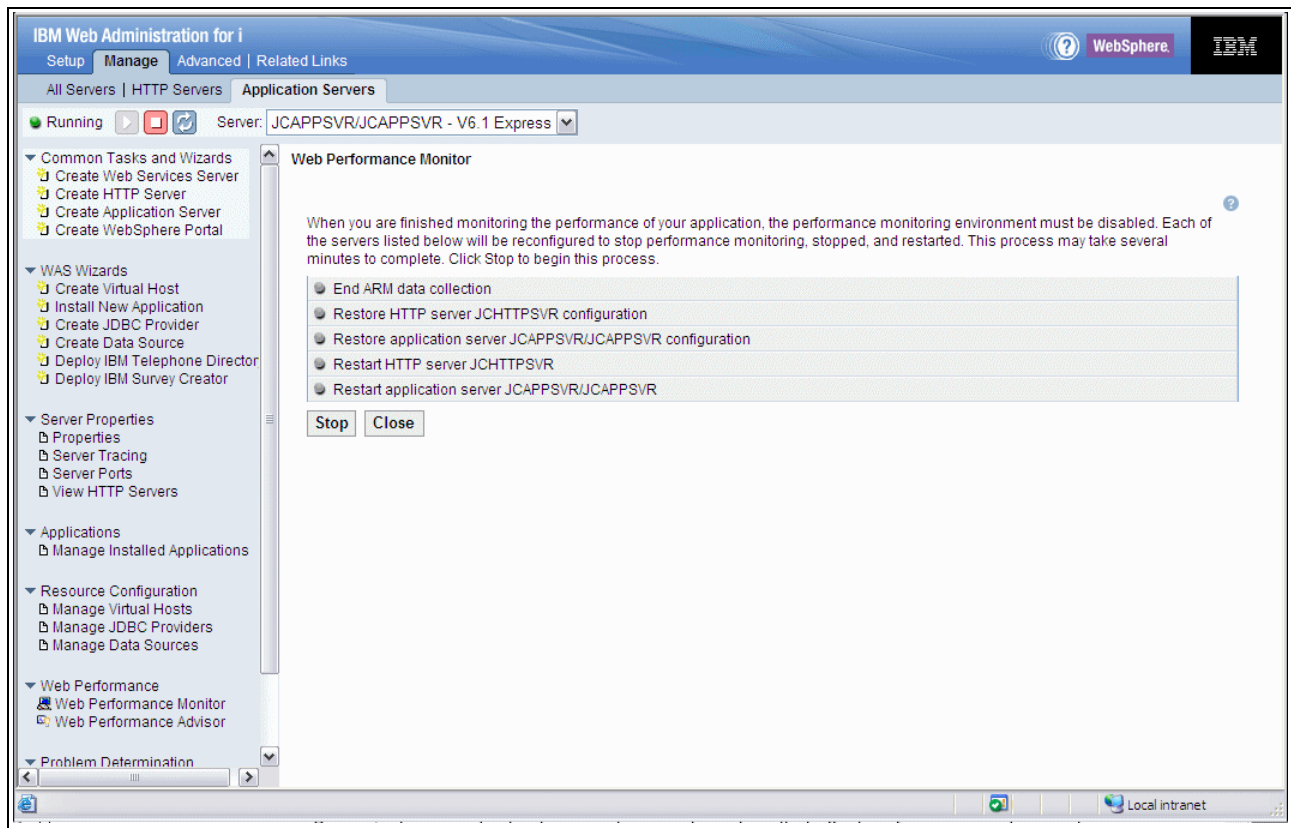


Figure 13-7 Ending the Web Performance Monitor

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.

13.2.2 Web Performance Advisor

The Web Performance Advisor is unique to IBM i, 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:

- ▶ Analyzing performance attributes

Does a quick analysis of your Web performance-related IBM i 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, 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.

- ▶ **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 to 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 i management of application servers interface to select the specific WebSphere Application Server. Select **Manage Details** and then click the **Web Performance Advisor** link (shown in Figure 13-2 on page 287).

You then get the initial advisor window shown in Figure 13-8 on page 294.

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 seen by the text and exclamation point icons in our example, shown in Figure 13-8. Our example shows WebSphere Application Server JCAPPSVR and HTTP server JCHTTPSVR.

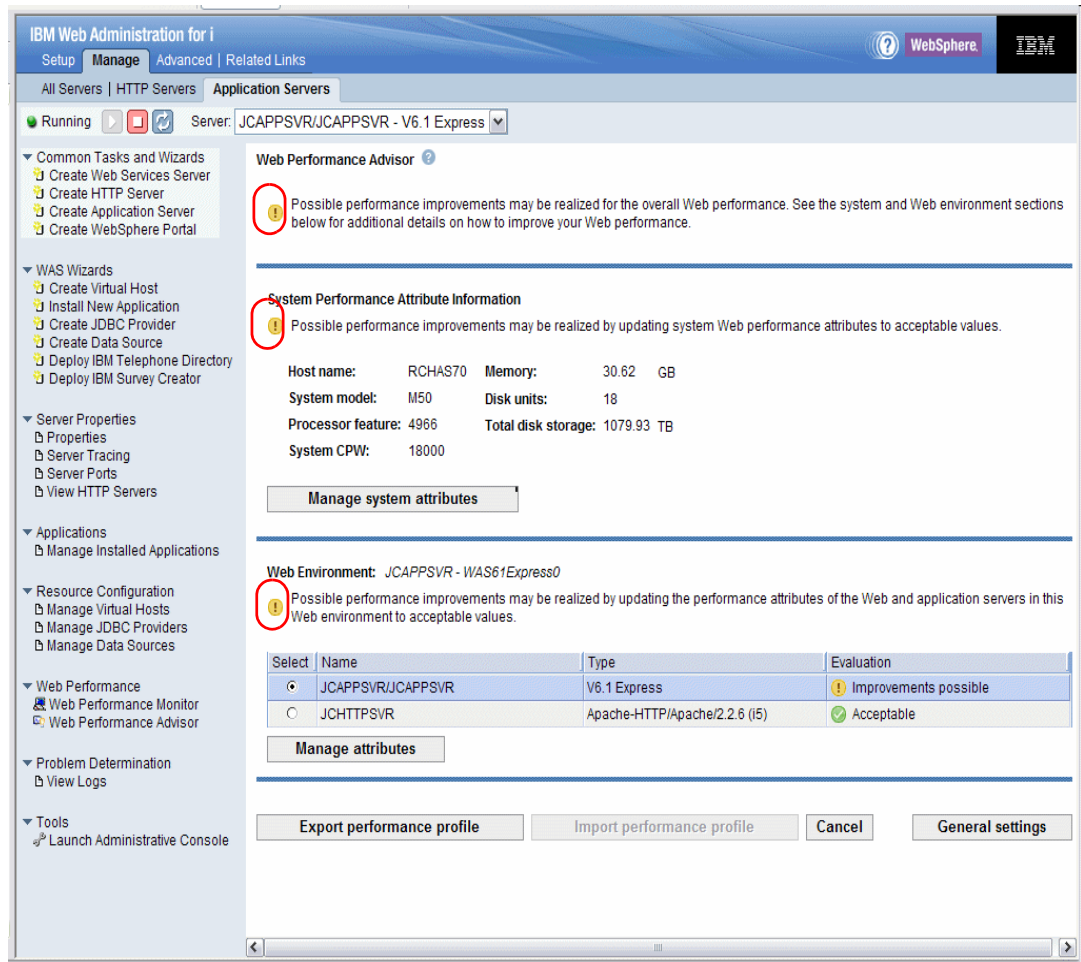


Figure 13-8 Web Performance Advisor - initial window

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, shown in Figure 13-8, 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.

Rchas70 is an IBM i partition. The host name, memory size, disk units, and total disk storage are for the Rchas70 partition. Other values are system-wide. Processor feature 4966 indicates that this is a POWER6 9409-M50 with processor speed 4.2 GHz and four processors active. The value 4966 corresponds to the contents of system value QPRCFEAT.

For POWER6 QPRCFEAT values, refer to *IBM Power 570 and Power 595 (POWER6) System Builder*, REDP-4439. An appendix lists QPRCFEAT values for all POWER6 models supporting IBM i 5.4 or 6.1 or later.

For older models, refer to *IBM System i Overview: Models 515, 525, 550, 570, 595, and More*, REDP-5052.

Figure 13-9 shows an example of System Resources tab information with *Advise* links you can review and consider making changes to the associated parameter. You also see a tab for performance settings, and also the PTF Groups tab and the Web PTFs tab.

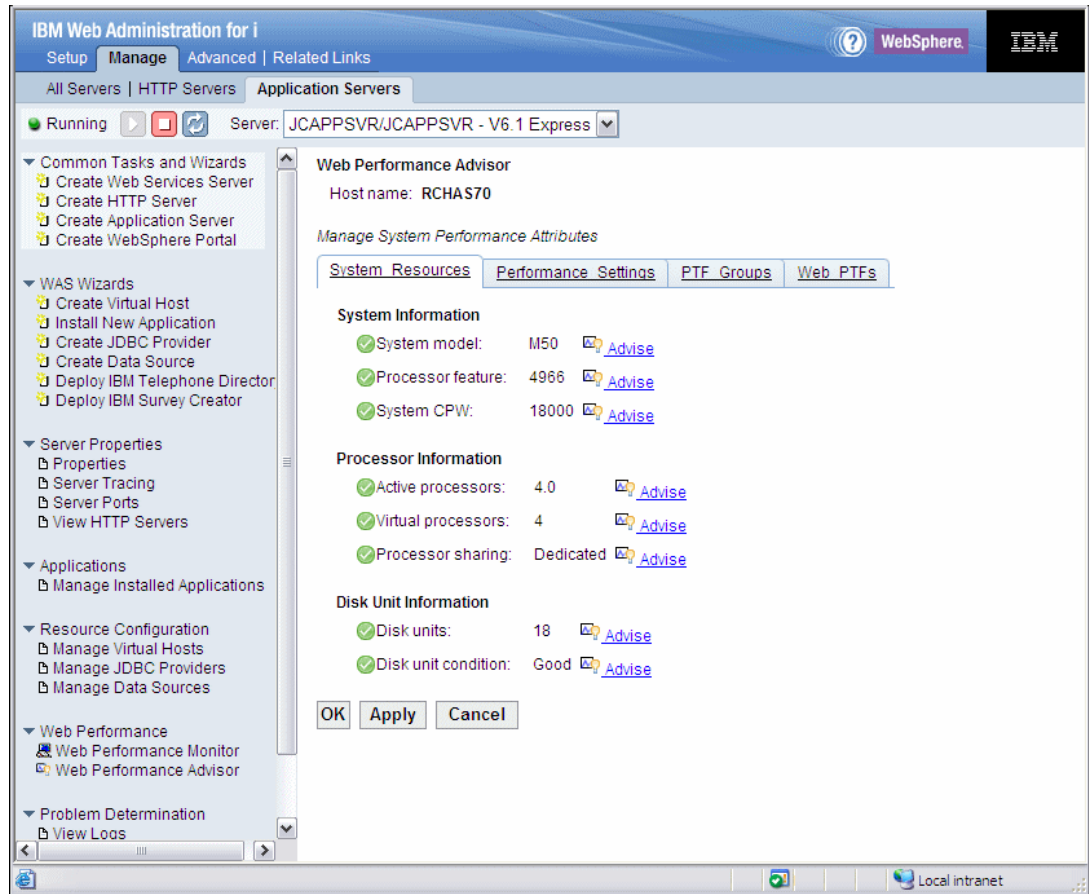


Figure 13-9 Web Performance Advisor - manage system attributes example

The PTF tabs provide a summary level list of PTFs, shown in Figure 13-10.

The screenshot displays the Web Performance Advisor interface for host RCHAS70. It features a navigation menu with tabs for System Resources, Performance Settings, PTF Groups, and Web PTFs. The main content area is divided into two sections:

The level and status of PTF Groups on this system

PTF Group	Level	Status	Description
SF99610	8190	Installed	CUMULATIVE PTF PACKAGE C8190610
SF99610	8288	Installed	CUMULATIVE PTF PACKAGE C8288610
SF99609	20	Installed	GROUP HIPER
SF99609	22	Installed	GROUP HIPER
SF99608	2	Installed	GROUP SECURITY
SF99601	5	Installed	DB2 UDB FOR ISERIES
SF99601	6	Installed	DB2 UDB FOR ISERIES
SF99562	3	Installed	JAVA

PTFs for Web related products on this system

PTF Information	Status
▶ 5733W61--WebSphere Application Server V6.1 for i5/OS	
▶ 5733W61--WebSphere Application Server V6.1 Express	
▶ 5761DG1--IBM HTTP Server for i5/OS	
▶ 5761JV1--IBM Developer Kit for Java	
▶ 5761JV1--Java Developer Kit 1.4	
▶ 5761JV1--Java Developer Kit 5.0	
▶ 5761JV1--J2SE 5.0 32 bit	
▶ 5761JV1--J2SE 5.0 64 bit	
▶ 5761JV1--Java SE Development Kit 6	
▶ 5761JV1--Java SE 6 32 bit	
▶ 5761JV1--Java SE 6 64 bit	
▶ 5761999--Licensed Internal Code	
▶ 5761SS1--i5/OS	
▶ 5761SS1--Extended Base Support	
▶ 5761SS1--Extended Base Directory Support	
▶ 5761SS1--Host Servers	
▶ 5761SS1--System Openness Includes	
▶ 5761SS1--Integrated Server Support	
▶ 5761SS1--Domain Name System	
▶ 5761SS1--Portable App Solutions Environment	
▶ 5761SS1--Digital Certificate Manager	
▶ 5761SS1--CCA Cryptographic Service Provider	
▶ 5761SS1--International Components for Unicode	

Figure 13-10 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 13-11 shows examples of the Advise windows for several of the partition's system resource parameters.

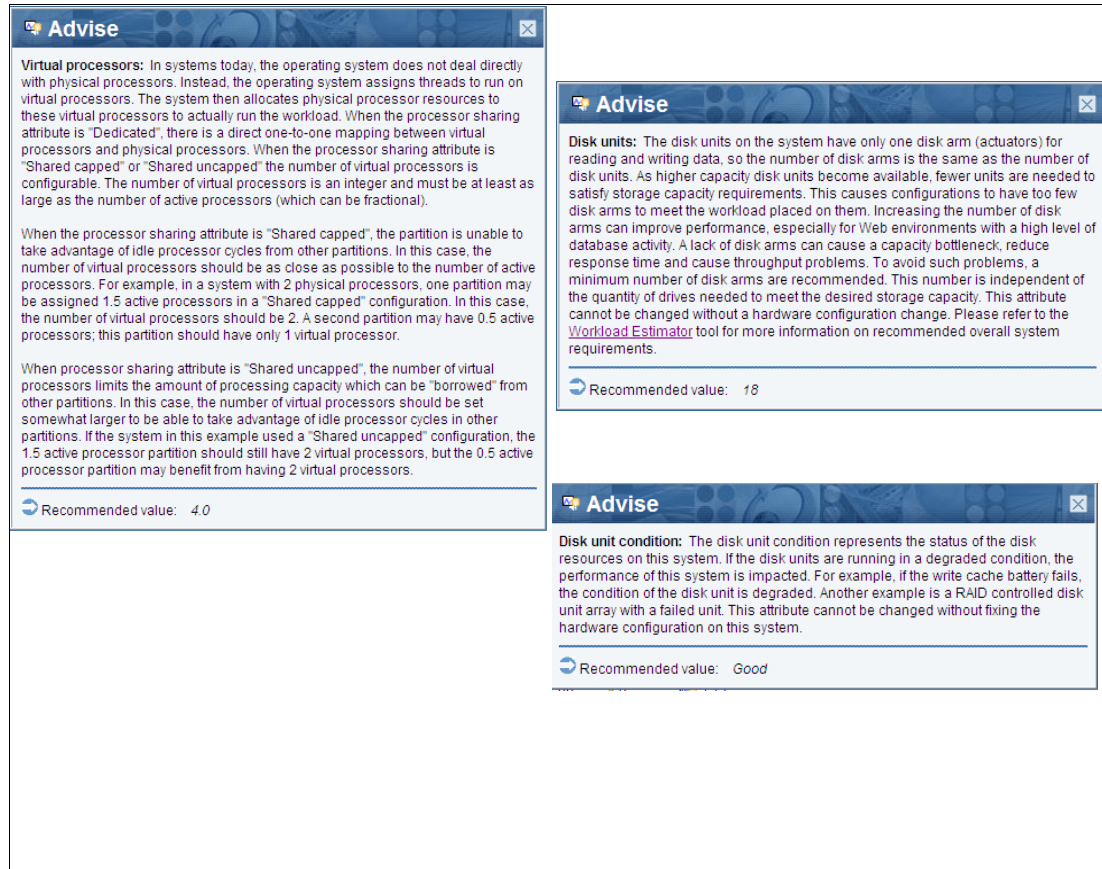


Figure 13-11 Example of system resources advice

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 M50 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 the 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 13-12 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.

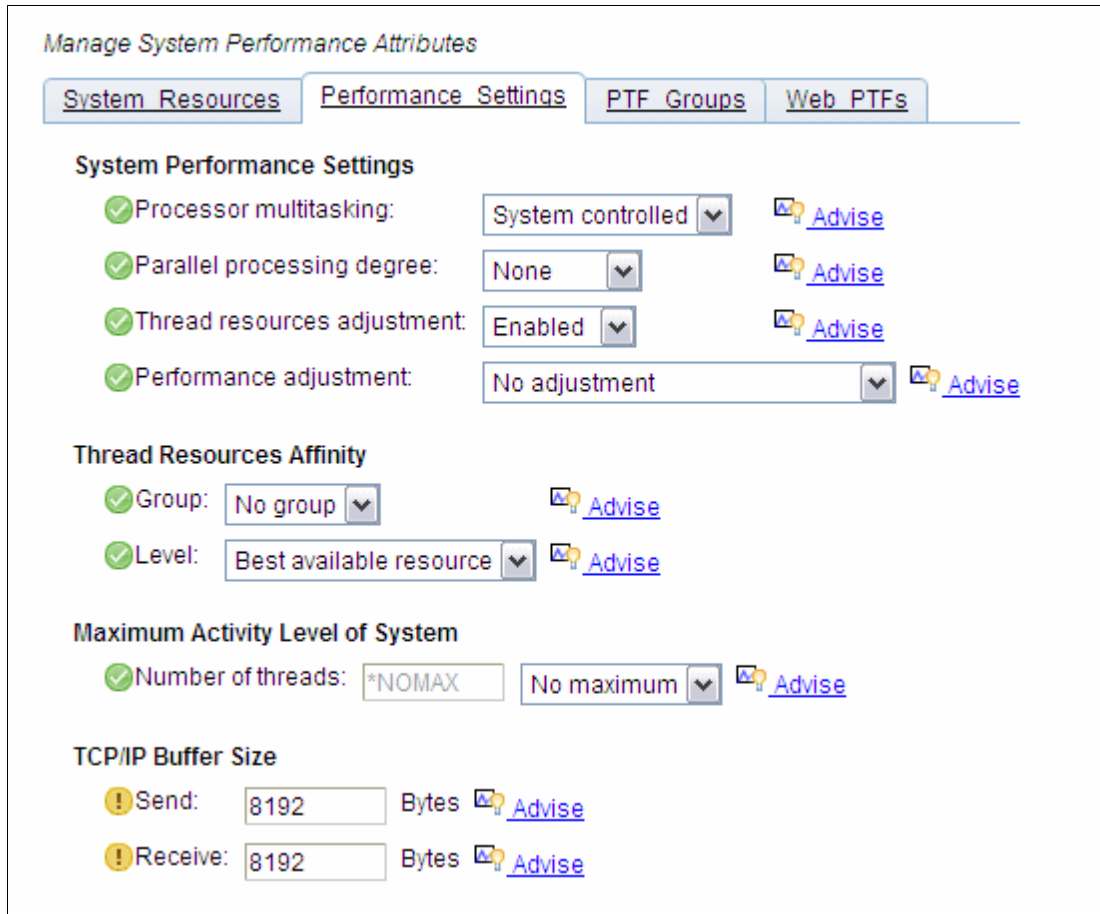


Figure 13-12 Web advisor - Performance Settings list

We chose to get the advise information about the TCP/IP send buffer size values, as shown in Figure 13-13.

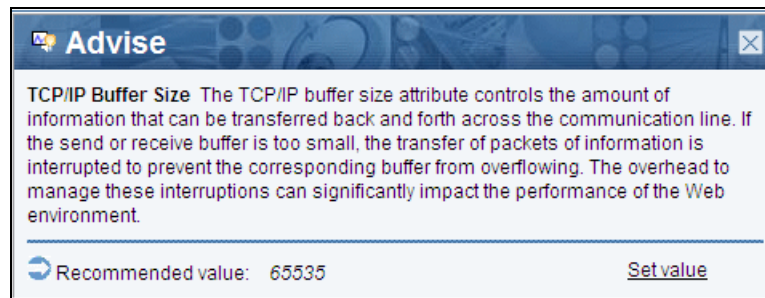


Figure 13-13 Performance Settings Advise examples: TCP/IP buffer

Review the Advise window text. Note that the 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.

The Web Environment section of the Web Performance Advisor (Figure 13-8 on page 294) allows you to select either your Web application server or HTTP server.

The Web Environment section 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 review the advise information, and, depending on the parameter, make the changes. Advise information is provided for each parameter presented to you.

Using the window shown in Figure 13-8 on page 294, we first selected **JCHTTPSVR** and then clicked **Manage attributes** to get the window shown in Figure 13-14.

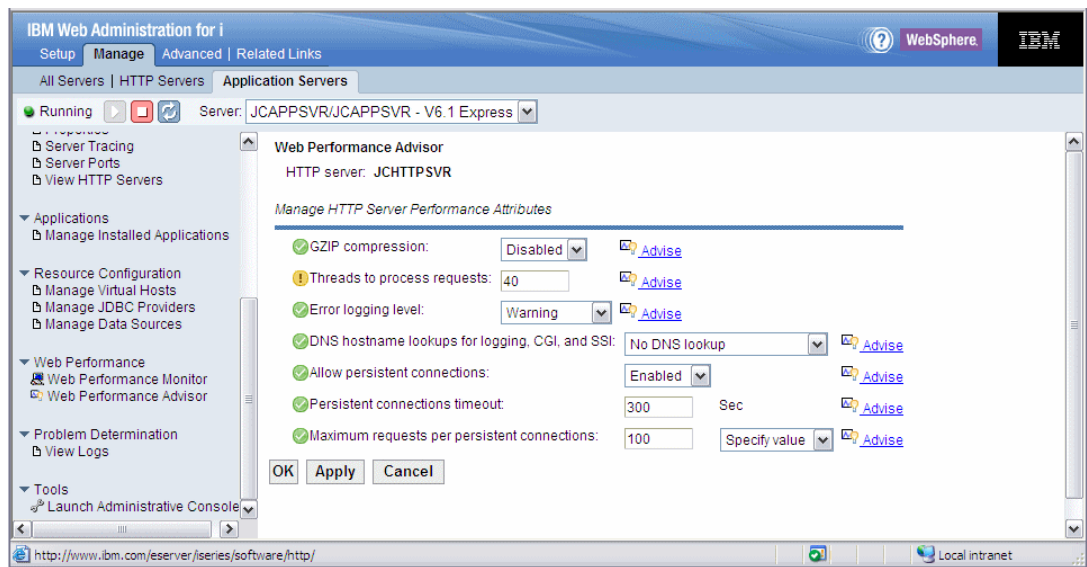


Figure 13-14 Web advisor - HTTP Server Web environment settings

Reviewing the Advise link for each parameter can provide useful information for those new to HTTP performance considerations and also provide helpful reminders to those who are familiar with, but not experts on, HTTP performance considerations.

In our example we show the Advise window (in Figure 13-15 on page 300) for the “Threads to process requests” parameter highlighted by the exclamation point icon (!). We clicked **Set value** in the lower right of the Advise window to change to the recommended 55 value, then we clicked **Apply**.

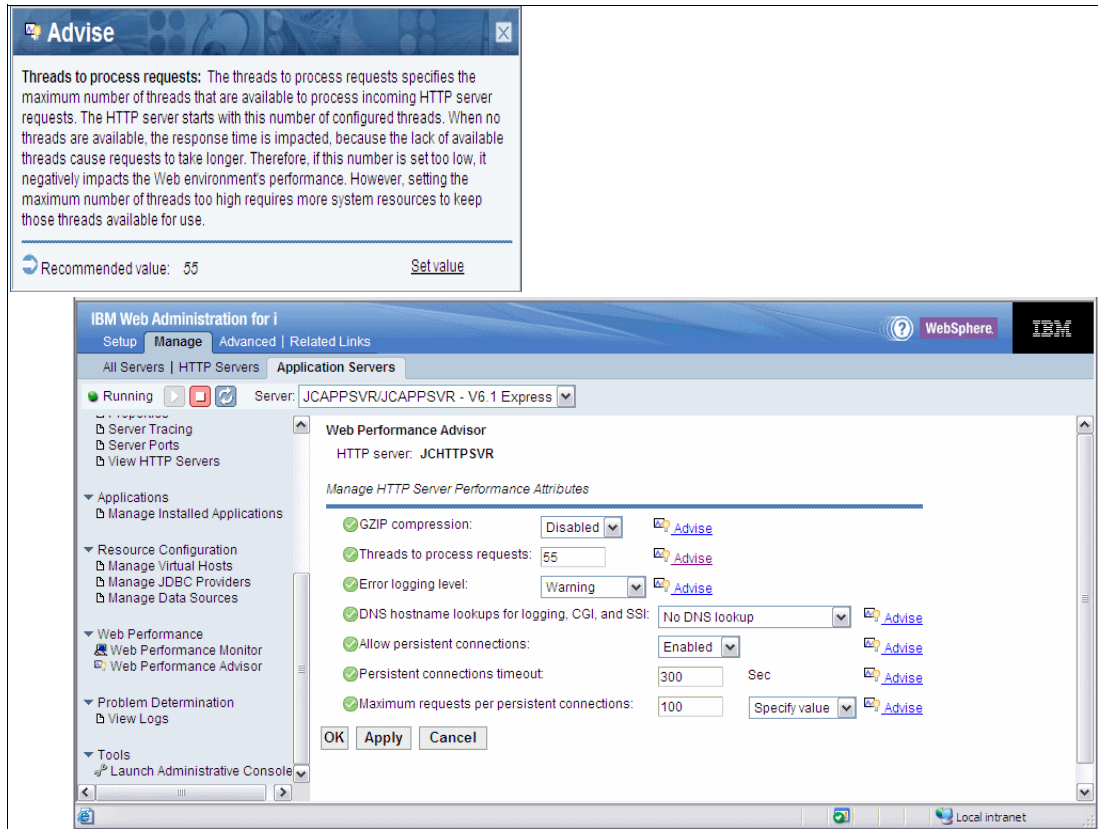


Figure 13-15 Advise window for threads to process

Now, using the window shown in Figure 13-8 on page 294, we selected **JCAPPSVR** and then clicked **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 13-8 on page 294. Click **Manage attributes** to initially open the JVM Settings tab window, as shown in Figure 13-16.

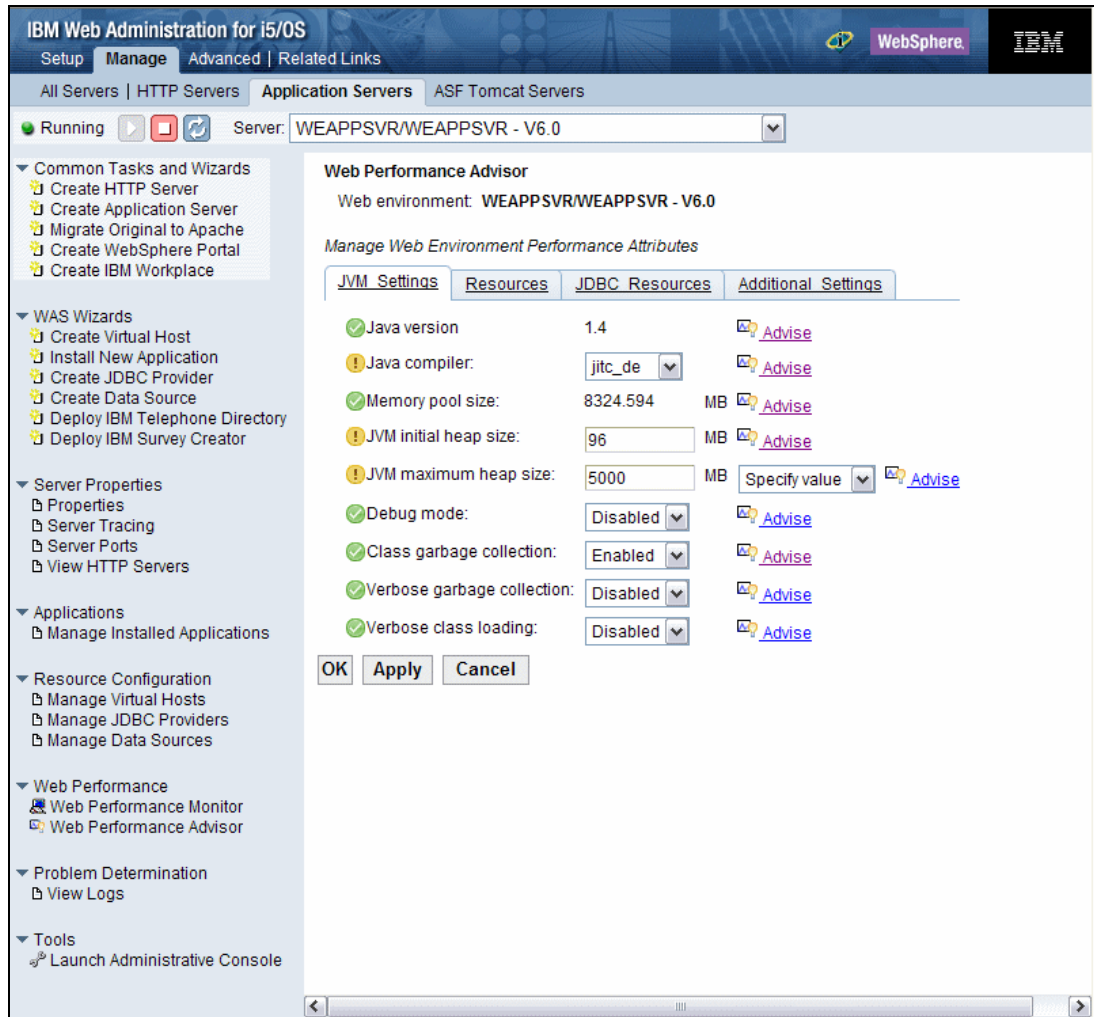


Figure 13-16 Web environment JVM settings

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 13-17 shows the Advise windows for each of these settings.



Figure 13-17 JVM settings advise examples

You can review the text in all of the Advise windows shown.

One of the important WebSphere Application Server parameters is the Java compiler one. The Direct Execution setting, `jltc_de`, was originally the fastest performing compiler option under IBM i. However, in the last few years the standardization of Java compilers across IBM i has resulted in the Just In Time (JIT) setting (`jltc`) being the best performing choice under IBM i.

We decided to make some of the advised changes. Figure 13-18 shows all of the changes we selected to make, before applying them.

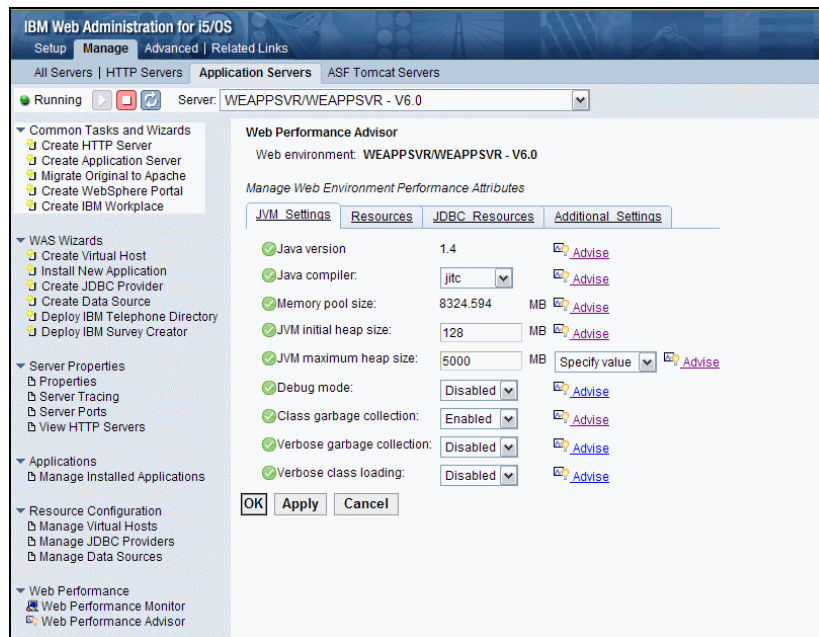


Figure 13-18 JVM settings after completing all changes to be applied

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 in 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 having to create your own applications.

Click the **Resources** tab to get the window shown in Figure 13-19.

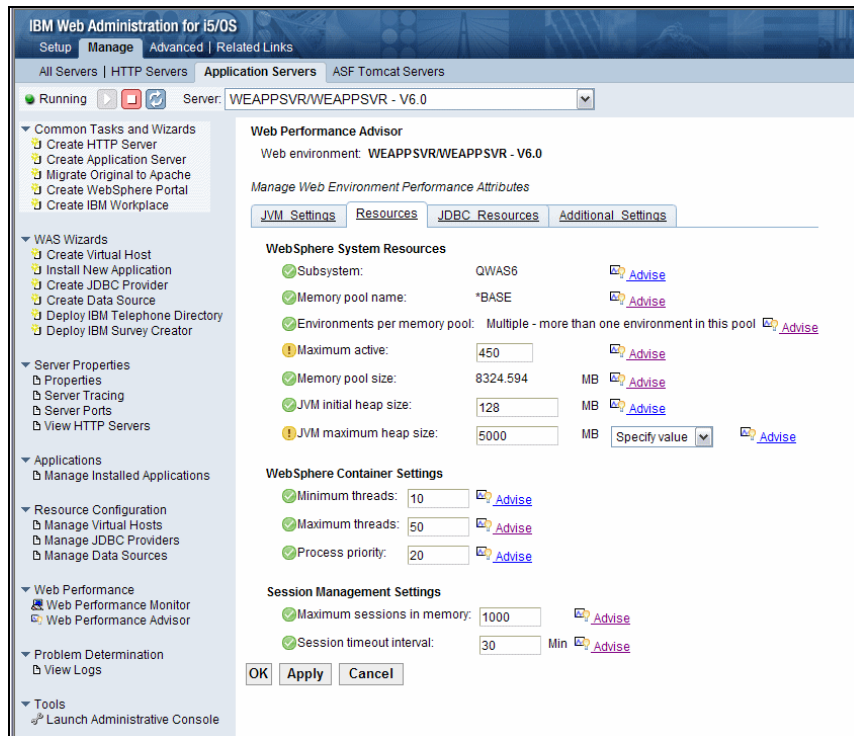


Figure 13-19 Advisor - Web Environment Resource

The resource settings cover IBM i-specific settings for the default subsystem in which WebSphere Application Server 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 13-20.

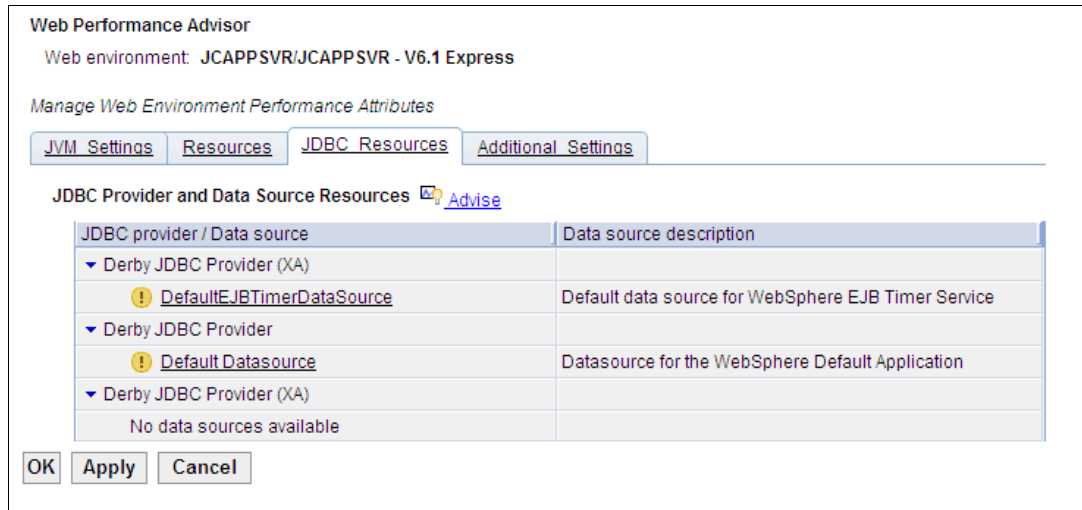


Figure 13-20 Advisor - Web environment - JDBC Resources

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 IBM i there are two commonly used IBM i 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 run in one of the database server jobs included under IBM i with the QSQSRVR job name prefix. This interface is unique to IBM i.

Remote functions 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 JCAPPSVR. Remember that the name itself does not identify which local or remote SQL interface under IBM i is being used.

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.

Last, but not least, we show a panel example of the Additional Settings tab window in Figure 13-21 on page 305.

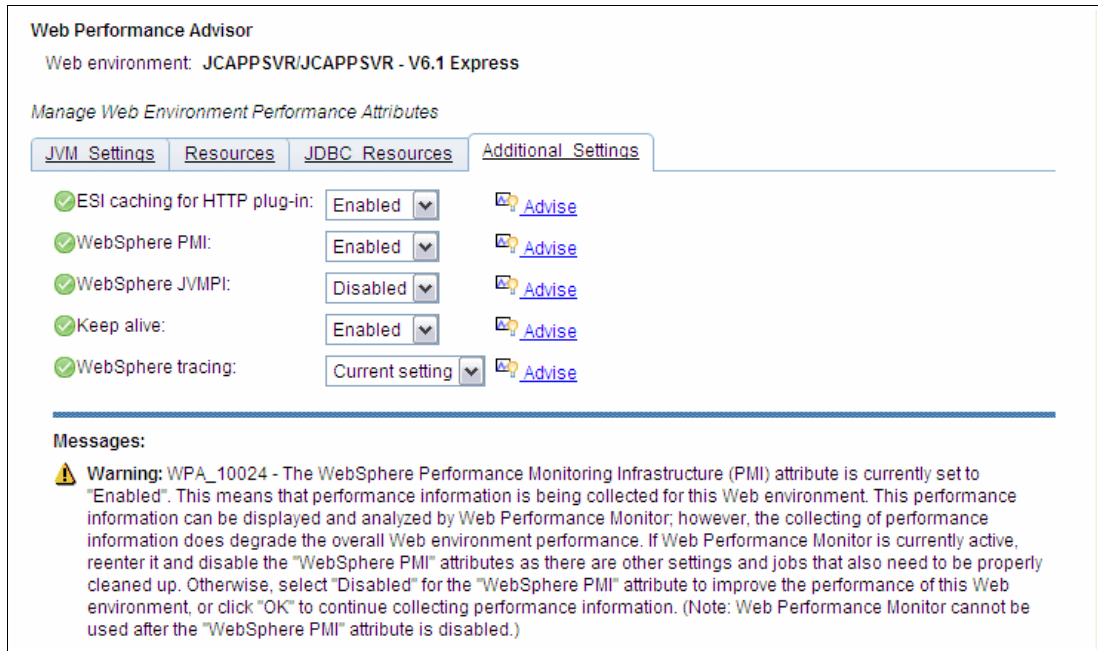


Figure 13-21 Advisor - Additional Settings

Additional Settings contains miscellaneous performance-related parameters. Note the parameters set to enabled with no exclamation points shown. However, as indicated under Messages, note that the WebSphere Application Server profile JCAPPSVR does have the WebSphere Performance Monitoring Infrastructure (PMI) enabled to collect performance information.

13.3 WebSphere Application Server Express V7.0

WebSphere Application Server Express V7.0 enhancements include:

- ▶ Java EE 5 certification, including Java Development Kit (JDK) 6.0, Web services enhancements, EJB 3.0 support, and Java Persistence API (JPA), which offer simplified programming models for building reusable persistent objects.
- ▶ Expanded Web services support in many areas, including support for JAX-WS, SOAP 1.2, MTOM, XOP, WS-ReliableMessaging, WS-Trust, WS-SecureConversation, and Kerberos Token Profile.
- ▶ Web 2.0 support extending SOA by connecting external Web services, internal SOA services, and Java Platform Enterprise Edition (JEE) objects into highly interactive Web application interfaces.
- ▶ New security management and auditing improvements with new tools and capabilities.
- ▶ New management and administration options to help reduce cost and increase flexibility. New OSGi-based runtime provisioning selects only the needed functionality for memory and space improvements.

13.3.1 Standards

WebSphere Application Server is fully enhanced and integrated with the Java EE 5 standards, including Web services and EJB 3.0 functions that were previously available as

feature packs in V6.1. V7.0 incorporates all the optional functionality of V6.1 into the base code. It continues to deliver new visionary standards through optionally consumable feature packs for WebSphere Application Server V7.0 to maximize flexibility and continue its long-running leadership in setting industry standards.

Java EE 5 certification

V7.0 builds on the EJB 3.0 and Web services feature packs, previously made available for V6.1, to deliver a fully certified Java EE 5 implementation. Java EE 5 delivers significant productivity improvements with the following:

- ▶ Support for EJB 3.0
- ▶ Exploitation of the annotation support added in Java SE V5 to reduce the number of artifacts required to build business applications
- ▶ Integration with the latest Web services standards
- ▶ Enhanced support for new Web applications

V7.0 delivers support beyond the Java EE 5 specification with enhancements in EJB deployment, automatic default bindings and reference resolution, and a new XML-based binding file format.

Java SE V6 support

It adds the ability to invoke the Java compiler from within the Java Virtual Machine (JVM) and includes scripts that enable you to access APIs within the JVM. Numerous components were added or upgraded in the areas of JMX, Java Web Start, JDBC 4.0, and new Web services.

EJB 3.0

It makes the programming model simpler. Programming is done using Plain Old Java Objects (POJO). This enables unit testing of business logic using a JVM instead of needing a full application server. Developers are free to focus on business logic in their applications instead of the ultimate execution environment. When you are ready to run in the EJB container, you can annotate the POJO with markers that tell the container what services to apply to the bean. In WebSphere, the EJB container dynamically applies those services at runtime, thus vastly simplifying the entire deployment process. The Java Persistence APIs (JPA) provided by WebSphere Application Server are based on the Apache OJPA project and enable simpler entities to be created using the annotated POJO model.

V7.0 integrates the support previously included in the EJB 3.0 Feature Pack and extends it with new support for Java Transaction API (JTA) 1.1, which enables persistence managers to coordinate updates with transactions.

Web services

New Web services standards help you to develop and deploy SOA-centric Web applications. V7.0 integrates support previously provided in the Web Services Feature Pack and extends it with new support for:

- ▶ Technologies introduced and defined in the WS-I profiles: WS Reliable Secure Profile and Basic Profile 1.2 and 2.0
- ▶ New Java EE 5 WS components focused on application portability
 - JAX-WS 2.1 for simplified Web service development and now a standardized programming model for programming using WS-Addressing
 - JAXB 2.1 for simpler XML to Java mappings, including improved support for XML type substitution

- StAX 1.0 for high-performance streaming parsing
- SAAJ 1.3 for enhancements in support of SOAP 1.2
- JSR 109 1.2 for supporting a declarative model to enhance or override the annotation-based programming model for Web services
- ▶ New Web services standards focused on interoperability
 - OASIS WS-ReliableMessaging for allowing heterogeneous nodes to send messages interoperably and reliably
 - OASIS WS-SecureConversation for an interoperable, efficient, and secure pipeline between services to exchange application messages
 - OASIS WS-Trust for a standard interoperable mechanism to obtain security tokens for Web services
 - OASIS WS-Transaction for standardized transactional Web services for those cases where transactional capabilities are required
 - OASIS WS-SecurityPolicy for exposing interoperably a standard way to define and declare security capabilities and constraints to a service
 - OASIS Kerberos Token Profile for the ability to use Kerberos as part of Web services requests
 - W3C WS-Policy for a standard language for exposing a service's constraints and capabilities
 - W3C WS-Addressing Metadata for declaring interoperably the need to use WS-Addressing in a standardized fashion
 - W3C SOAP 1.2 for using the latest standards for Web services and interoperability with vendors such as Microsoft
 - W3C MTOM and XOP for interoperable and efficient exchange of binary data between services
 - WS-MetadataExchange for a standard mechanism for exposing a way to retrieve metadata about the service
- ▶ Web services performance improvements
- ▶ Simplified administration of Web services and the qualities of service associated with them through policy administration

Programming model enhancements

V7.0 delivers these new programming model enhancements:

- ▶ Session Initiation Protocol (SIP) high availability options improve performance, help to reduce cost, and improve availability of SIP implementations.
- ▶ JAX-WS 2.1 delivers the successor to JAX-RPC and vastly simplifies the creation of new Web services.
- ▶ JAXB 2.1 support simplifies access to XML documents by Java programs.
- ▶ Portlet specification 2.0, Java Specification Requests (JSRs) 286 delivers the next level of the portlet programming model. Among other enhancements, portlets can send and receive events in order to perform changes to the portlet state, and portlets can share session attributes with servlets and JSPs located within the same Web applications.
- ▶ IBM now certifies Spring running in the WebSphere Application Server. WebSphere does not ship or support Spring itself. EJB 3.0 is the Java EE standard for both the local and remote POJO programming model for business logic implemented in Java.

Web 2.0

It extends an SOA by connecting external Web services, internal SOA services, and JEE objects into highly interactive Web application interfaces. To help reduce IT costs and speed time to market, Web 2.0 support enables a supported AJAX development toolkit.

AJAX

This is an open technique for creating rich user experiences in Web-based applications that does not require additional browser plug-ins. AJAX interfaces result in increased customer satisfaction and line-of-business productivity by streamlining user interaction with Web-based applications. Web 2.0 highlights include:

- ▶ Web 2.0-to-SOA connectivity for enabling connectivity from AJAX clients and mashups to external Web services, internal SOA services, and JEE assets. This connectivity also extends enterprise data to customers and partners through Atom Syndication Format (ATOM) and Really Simple Syndication (RSS) Web feeds.
- ▶ AJAX Messaging for connecting AJAX clients to real-time updated data such as stock quotes or instant messaging.
- ▶ AJAX Development Toolkit for WebSphere Application Server based on Dojo with IBM extensions.

Web 2.0 support is provided by the WebSphere Application Server Feature Pack for Web 2.0, an optional product extension for WebSphere Application Server V7.0.

Systems management

WebSphere Application Server addresses the management challenges associated with mission-critical SOA implementations. From dynamic runtime componentization options, to flexible management of extremely large geographically dispersed deployments, or fine-grained options for individual server implementations, V7.0 delivers enhancements to improve the effectiveness and help reduce the cost of systems management.

- ▶ Flexible management helps dramatically reduce the overhead of administration in large and distributed WebSphere Application Server environments. An administration agent within the base and Express editions:
 - Reduce the server footprint associated with administration.
 - Simplify administration of multiple base application server instances.
 - Work with the new Job Manager function of WebSphere Application Server Network Deployment to centralize the distribution of applications, application updates, or application server configuration updates across a large number of WebSphere administrative domains (cells).
- ▶ Flexible management delivers a loosely coupled management model based upon an asynchronous implementation that dramatically increases the scale of administration for application server implementations.
- ▶ Job Manager augments the Deployment Manager of WebSphere Application Server Network Deployment.
- ▶ New Runtime Provisioning Service commissions only those components that are required for a given set of applications. At application server startup, V7.0 can be configured to start only the runtime components needed to support the application portfolio for a given application server. At application installation time, WebSphere examines the application and creates an application-specific *activation plan*. This new granularity in selecting application runtime components can significantly reduce the application server footprint and resources needed for a given application portfolio.

- ▶ Administration console now supports DataPower®. Administrative integration with DataPower simplifies the environment and can reduce administrative costs for clients who augment their WebSphere Application Server implementation with DataPower appliances. This also addresses the requirement for management of multiple DataPower machines.
- ▶ Fine-grained administration security can now be enforced through the administration console. Access can be restricted based on the administrators' role at the cell, node, cluster, or application level. This enables fine-grained control over administrator scope and control. This capability is especially valuable in large cell implementations where multiple administrators are responsible for subsets of the application portfolio running on the cell.
- ▶ Simple properties file-based configuration allows you to extract WebSphere configuration from one environment into a properties file and apply it to same or other environments while parameterizing platform-specific configuration.
- ▶ Wsadmin administrative scripting is further enhanced in WebSphere Application Server V7.0 with the extension of Command Assistant to an even greater number of administrative screens, and with the availability of an extended sample script library to accelerate automation implementations.

Application development and management

WebSphere Application Server V7.0 delivers the tooling support you need to leverage the latest standards and new virtualized application packaging that can help reduce the cost of managing your application environment.

- ▶ The IBM Rational Application Developer (RAD) for WebSphere Assembly and Deploy function replaces the former Application Server Toolkit (AST) function of the application server. Assembly and Deploy is licensed as part of the WebSphere Application Server. It provides a superset of the older AST functionality and is fully integrated within the RAD tooling environment. The remaining RAD for WebSphere functions are included on a trial basis and are easily purchased through a downloadable license key.
- ▶ WebSphere Business Level Applications expands the notion of *application* beyond J2EE. It enables new formats for deployment packages, a new way of representing applications. WebSphere Business Level Applications comprises any application elements hosted by the application server. This allows the creation of a new logical application that represents a meaningful business function. Full lifecycle management is enabled: install, distribute, activate, monitor, update, and remove. WebSphere Business Level Applications significantly improves the management of multiformat applications.

Security

WebSphere Application Server V7.0 extends its security-rich features with security controls that deliver greater granularity management expanded Kerberos Single Sign-on Technologies, enhanced secure proxy, and new audit and compliance features to meet your business needs.

- ▶ WebSphere Secure Proxy is easier and more secure with WebSphere Application Server V7.0. WebSphere Secure Proxy offers a new DMZ Hardened Proxy profile option. It is designed to improve security by minimizing the number of external ports opened, loading only signed JARs, and running as an unprivileged user when binding to well-known ports. Both static and dynamic routes are supported.
- ▶ WebSphere Security Domains offer greater granularity management of security controls and more flexibility in configuring security under centralized management. WebSphere Security Domains can separate WebSphere administrative security and your business application security. For example, business applications can be configured to use *external* user population LDAP registry while the WebSphere administration can use *internal* user population using a completely different LDAP registry. Granularity can be further

expanded between business applications by allowing a separate security configuration between business applications using the new security configuration application scoping to a cell, a cluster, or application server level. This new level of security granularity gives you additional flexibility in the security mechanism implemented across various application portfolios.

- ▶ New security auditing capabilities offer enhanced compliancy and auditing. The auditing capabilities allow tracking a number of security-related events, including administrative actions such as:
 - Security configuration changes
 - Key and certificate management
 - Access control policy changes

You can audit business applications to record security events such as authentication or authorization attempts. This new security logging and auditing capability ensures accountability for administrative actions. In addition, it offers a tamper-proof event log, as well as the ability to report against the auditing events. A new Auditor role optionally restricts access to the security auditing records. Administrators cannot view or change the auditing records or policy, and auditors cannot change the WebSphere configuration or runtime state.

- ▶ Kerberos Authentication and Single Sign-on expands the WebSphere Application Server V6.1 SPENGO support by further integrating the existing support and expanding the Kerberos offering by delivering interoperability and identity propagation with other applications (such as .NET, DB2, and others) that support Kerberos authentication. It enables single sign on (SSO), end-to-end interoperable solutions, and preserves the original requester identity.
- ▶ Common Criteria Assurance Level 4 certification can give you the confidence that products will deliver security functions such as identification, authentication, and user data protection. Common Criteria is a scheme for independent assessment, analysis, and testing of IT products to a set of security requirements used.

Accessibility by people with disabilities

A U.S. Section 508 Voluntary Product Accessibility Template (VPAT) containing details on accessibility compliance can be requested at:

http://www.ibm.com/able/product_accessibility/index.html

13.3.2 Configuration options

The WebSphere Application Server family offers multiple configuration options and support for multiple business models and deployment platforms. WebSphere Application Server configurations are built on a common server foundation that delivers the same full J2EE or J2SE programming model and core application server functionality across the configurations. This offers unparalleled flexibility to grow applications without the cost of rearchitecting, reprogramming, and migrating to different technology bases.

Configuration options include:

WebSphere Application Server - Express

A simple, cost-effective solution for developing, deploying, and managing dynamic Web sites and applications. WebSphere Application Server - Express delivers the full J2EE programming model, and includes an easy-to-use, visual, application development tool from Rational. Designed to meet the unique needs of small to mid-size companies, WebSphere

Application Server - Express is limited to two CPUs per server and includes a set of integrated applications, wizards, and samples to help you get up and running quickly.

WebSphere Application Server

Optimized to ease administration in a scalable, single-server deployment environment, this core J2EE configuration delivers a secure, high-performance transaction engine for dynamic applications. Perfect for standalone, departmental applications and Web services, WebSphere Application Server offers unlimited CPUs per server and unbeatable platform coverage. This product is available through the Passport Advantage® ordering system.

WebSphere Application Server for Developers

The functional equivalent of the core WebSphere Application Server configuration, but licensed for development use only, WebSphere Application Server for Developers is an easy-to-use development environment for building and testing On Demand Business applications. This product is available through the Passport Advantage ordering system.

WebSphere Application Server Network Deployment

For enterprises that need near-continuous availability, advanced management, and automated performance optimization for their mission-critical applications, WebSphere Application Server Network Deployment delivers an environment that is highly available, dynamically scalable and easily managed. Building on the core WebSphere Application Server configuration, WebSphere Application Server Network Deployment offers:

- ▶ World-class clustering and high availability manager to help eliminate the cost of system downtime.
- ▶ Edge components that offer sophisticated load balancing, caching, and centralized security for enhanced performance at the edge of the network.
- ▶ Advanced Web services capabilities to enable you to connect to systems outside the firewall and across disparate application frameworks.
- ▶ Extended Web services management and advanced remote administration to make management of more complex environments less time-consuming and resource-consuming.

This product is available through the Passport Advantage ordering system.

13.3.3 Ordering information

Preload is not available for Web Enablement for i5/OS.

Web Enablement for i5/OS has three version that are orderable: V6.0, V6.1, and V7.0. These versions will run on IBM i 5.4 or 6.1.

Orders for IBM i 5.4

If you place an initial or upgrade order for IBM i, the following software product is defaulted onto your configuration when ordered for hardware models 9406-595, 9117-MMA, and 9119-FHA:

5722-WE2	IBM Web Enablement for IBM i5/OS, V5.3 WebSphere Express V6.0, Supply Feature 5902.
----------	---

This product will no longer be defaulted onto the IBM i configuration when ordered for hardware models 9406-525, 9406-550, 9407-515, 9407-M15, 9408-M25, or 9409-M50. If you need this software product, you will need to explicitly order it in the initial or upgrade order.

Orders for IBM i 6.1

If you place an initial or upgrade order for IBM i, the following software product is defaulted onto your configuration when ordered for hardware models 9406-595, 9117-MMA, and 9119-FHA:

5722-WE2	IBM Web Enablement for IBM i5/OS, V5.3 WebSphere Express V6.1 DVD, Supply Feature 5904
----------	--

The following product will no longer be defaulted onto the IBM i configuration when ordering an IBM i upgrade from version 5 to version 6, for hardware models 9406-525, 9406-550, 9407-515 9407-M15, 9408-M25, or 9409-M50. If you need this software product, you will need to explicitly order it in the upgrade software order.

5722-WE2	IBM Web Enablement for IBM i5/OS, V5.3 WebSphere Express V6.1 DVD, Supply Feature 5904
----------	--

For further information about WebSphere Application Server for System i refer to the following IBM Web site:

<http://www.ibm.com/systems/i/software/websphere/news/sitenews.html>

For more information on all of the functions and packaging of functions (configuration options) under WebSphere Application Server V7.0 and earlier versions, refer to the following Web site:

<http://www.ibm.com/software/webservers/appserv/was/>

13.4 WebSphere Application Server performance data

With WebSphere Application Server V6.0 or later, 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 the IBM i Information Center for the field definitions for these files:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

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>

13.5 Simplified use of SSL with 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 such as:

```
http:rchas70:2001
```

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 setup for *ADMIN, you would connect using the secure ULR format:

```
https:rchas70:2010
```

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 IBM i 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 13-22 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://rchas70:2001`.

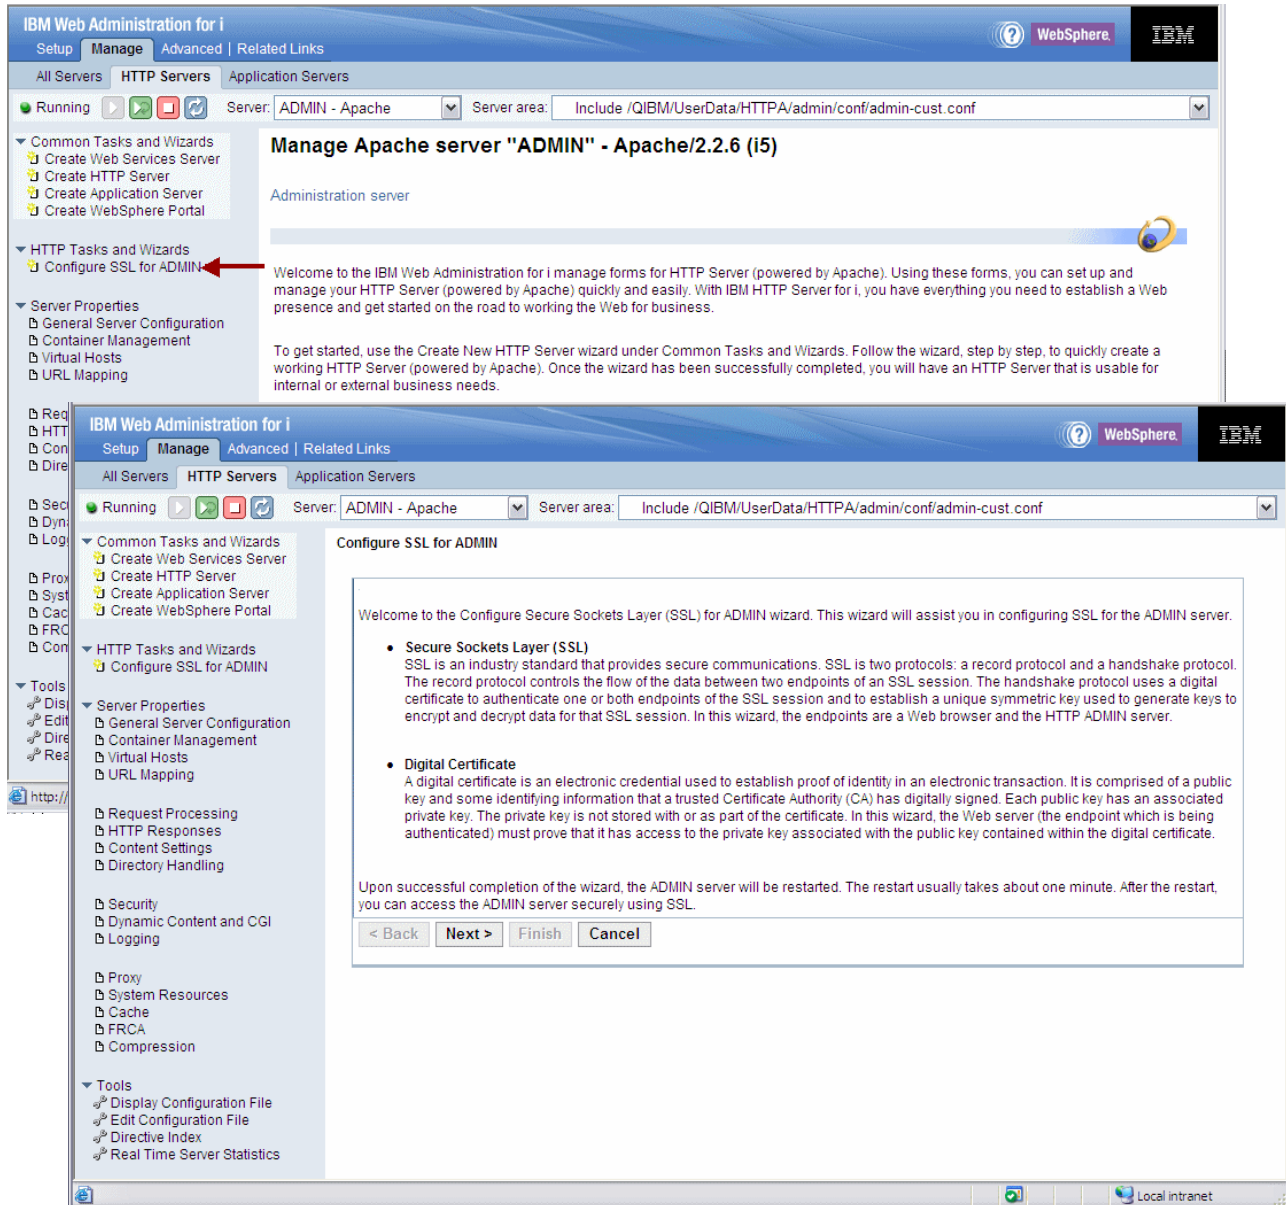


Figure 13-22 Setting up the HTTP Administration server to use SSL (1 of 2)

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 13-23.

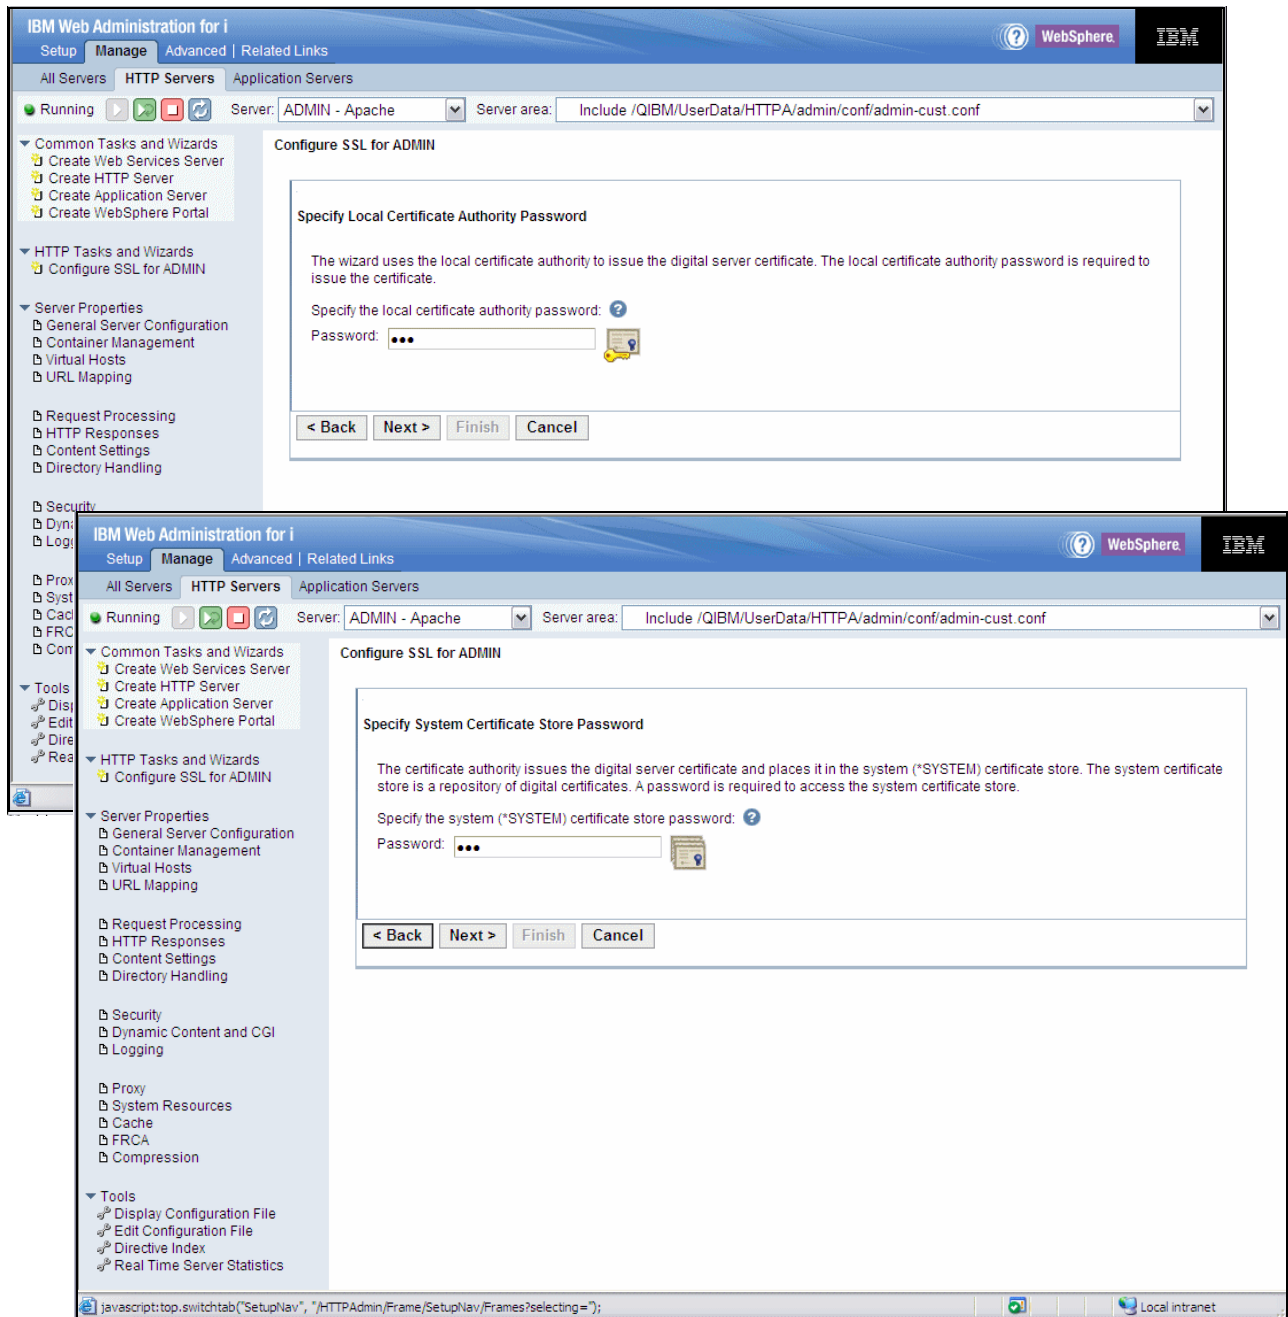


Figure 13-23 Setting up the HTTP Administration server to use SSL (2 of 2)

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 it 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 13-23 on page 315.

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.

If this is the first time you are using the *SYSTEM certificate store, you will be requested to 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 13-24.

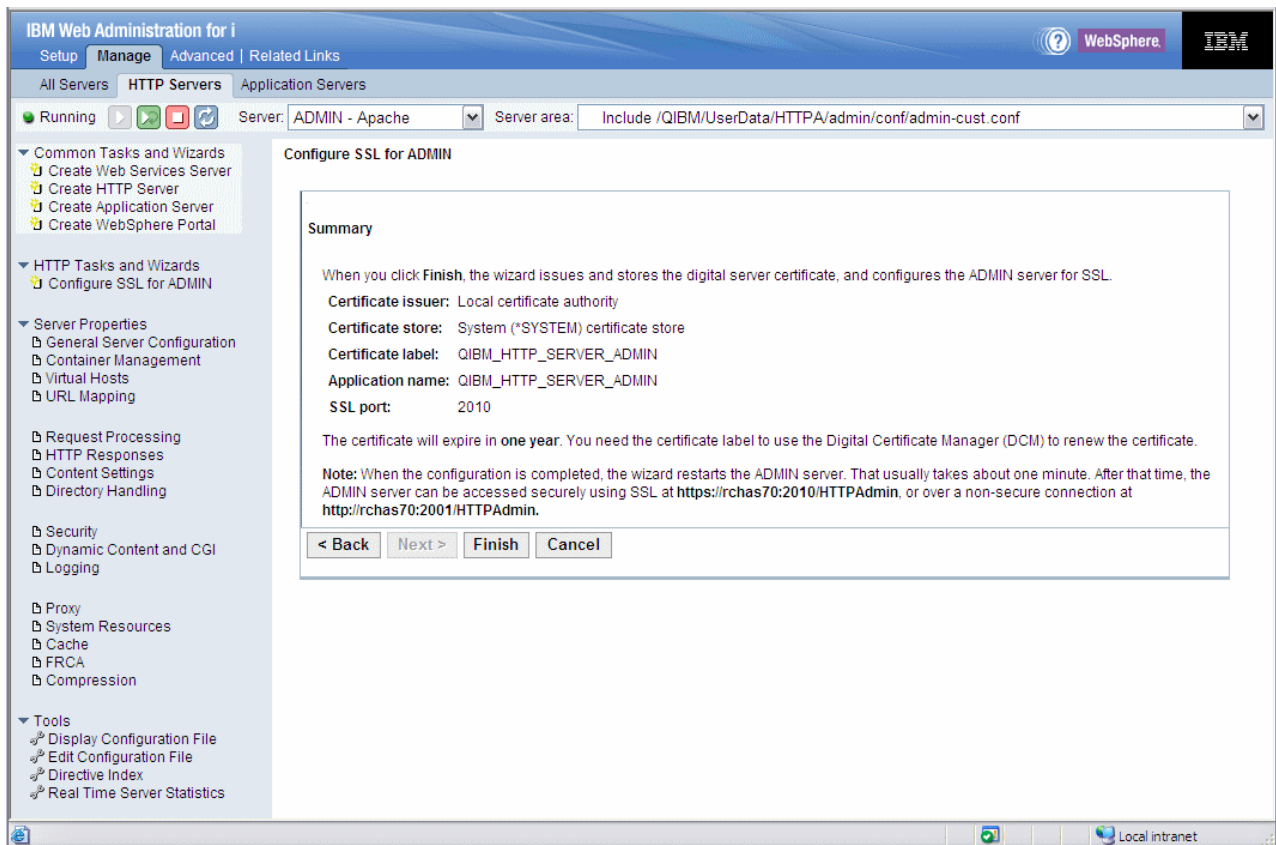


Figure 13-24 Summary: setting up the HTTP Administration server to use SSL

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 insecure requests over port 2001 and secure requests over port 2010.



Application development

This chapter discusses the following topics:

- ▶ “IBM Developer Kit for Java” on page 318
- ▶ “IBM Toolbox for Java” on page 327
- ▶ “Websphere Development Studio” on page 331
- ▶ “Application Development ToolSet” on page 338
- ▶ “Debugger” on page 338
- ▶ “Portable Application Solutions Environment” on page 339
- ▶ “Qshell and utilities” on page 340
- ▶ “XML toolkit” on page 341
- ▶ “Control language” on page 342
- ▶ “Structure Query Language” on page 345
- ▶ “Application programming interfaces” on page 354
- ▶ “Globalization” on page 359
- ▶ “Multithreaded applications” on page 359

14.1 IBM Developer Kit for Java

IBM Developer Kit for Java is optimized for use in the IBM i environment. It uses the compatibility of Java programming and user interfaces, so you can develop your own IBM i applications.

IBM Developer Kit for Java allows you to create and run Java programs on your IBM i. It is a compatible implementation of the Sun Microsystems, Inc. Java Technology.

14.1.1 Release-to-release compatibility

In IBM i 6.1, direct execution (DE) is no longer supported. This change has implications for both the OPTIMIZE and ENBPFRCOL parameters on the Create Java Program (CRTJVAPGM) command.

As of IBM i 6.1, the speed of transformation will not be impacted by the optimization level that you set. In IBM i 6.1, direct processing is no longer supported. This means that the optimization level for Java programs is ignored and the OPTIMIZE(*INTERPRET) parameter value is used when you create a Java program on IBM i 6.1 or a subsequent release.

As in previous releases, a Java program contains a pre-verified version of one or more Java class files. However, in IBM i 6.1, a Java program does not contain machine instructions. At runtime, the Java program is either interpreted from the byte codes or else run with the Just-In-Time (JIT) compiler. Variables can be displayed and modified while debugging.

For target releases prior to IBM i 6.1, the OPTIMIZE parameter specifies the optimization level of the Java program. When creating a Java program for a target release prior to IBM i 6.1, the value of the OPTIMIZE parameter is encapsulated within the Java program, but no machine instructions are generated. The encapsulated value of the OPTIMIZE parameter is used during Java program retranslation on the target release of the operating system.

For Java programs created for target release IBM i 6.1 and later, the ENBPFRCOL parameter value is ignored and ENDPFRCOL(*NONE) is used. To enable performance collection for Java applications in IBM i 6.1 and later releases, you must use the os400.enbpfrcol Java system property.

For a target release prior to IBM i 6.1, you may specify values other than *NONE, but values other than *NONE will not be effective until the Java program is retranslated on the target release of the operating system.

You can still create a Java program with the CRTJVAPGM CL command. The resulting Java program contains class information which is used by the Java runtime interpreter or the Java Just-In-Time (JIT) compiler. Use of the CRTJVAPGM CL command on all .jar and .zip files containing Java classes is optional. However, use of the CRTJVAPGM command is recommended because it helps consolidate and reduce system storage utilization.

Note: The above section applies only to Java applications that are run using the Classic Java Virtual Machine.

14.1.2 Changes to adopted authority support

Support for adopted user profile authority through Java programs is being withdrawn in IBM i 6.1.

In IBM i 6.1, Java applications will no longer be able to adopt user profile authority through Java programs unless IBM Adopt Authority for Java for i5/OS programming request for price quotation (PRPQ) 5799-AAJ is installed. Support for Java adopted authority and 5799-AAJ will be withdrawn in a future release.

IBM i 6.1 should be used as a transition release for Java adopted authority.

Java programs created on releases before to IBM i 6.1 using the Create Java Program (CRTJVAPGM) command and specifying USRPRF(*OWNER), can be saved and restored to IBM i 6.1. However, the adopted authority attributes of the program will not be used unless PRPQ 5799-AAJ is installed. If the PRPQ is installed, the adopt authority attributes in the Java program will be used. Java programs created on IBM i 6.1 with a target release of i5/OS V5R3 or IBM i 5.4 will have the adopt attribute stored in the Java program. When that Java program is saved on an IBM i 6.1 or later system and restored to a i5/OS V5R3 or IBM i 5.4 system, those adopt attributes will be used as they are today on those releases.

Determining if your applications use adopted authority

You can use a tool in releases i5/OS V5R3 and IBM i 5.4 to help you determine whether you have Java applications that will be affected by the changes to adopted authority. The tool works with the JVM to report Java programs that the JVM has logged as using adopted authority. This tool is separate from PRPQ 5799-AAJ, and is available through PTFs.

For information about how to modify your application to accommodate this change, and the available PTFs for each release, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rza/rzahaadopt.htm>

14.1.3 Java Cryptography Extension hardware

The IBM Java Cryptography Extension using Common Cryptographic Architecture for IBM i (IBMJCECCAI5OS) implementation extends the Java Cryptography Extension (JCE) and Java Cryptography Architecture (JCA) to add the capability to use hardware cryptography through the IBM Common Cryptographic Architecture (CCA) interfaces.

The IBMJCECCAI5OS provider plugs into the JCE framework in the same manner as the current providers. For hardware requests, the CCA APIs are invoked by the new native methods. The IBMJCECCAI5OS provider stores CCA RSA key labels in a new Java key store type of JCECCAI5OSKS.

To use hardware cryptography, you must have the following installed on your system:

- ▶ A model 4764 Cryptographic Coprocessor
- ▶ IBM i (5761-SS1) Option 44 - CCA Cryptographic Service Provider
- ▶ Licensed Program Offering (LPO) 5733-CY1 - IBM eServer™ iSeries Cryptographic Device Manager

14.1.4 Java Virtual Machine Tool Interface

The Java Virtual Machine Tool Interface (JVMTI) is an interface for analyzing the Java virtual machine (JVM), and allows communication between the JVM and a native agent.

JVMTI replaces Java Virtual Machine Profiler Interface (JVMPPI) and the Java Virtual Machine Debugger Interface (JVMDI). It contains all the functionality of both JVMDI and JVMPPI, plus new functions. JVMTI was added as part of J2SE 5.0. In JDK 6.

JVMTI support places hooks in the JVM and the Just-In-Time (JIT) compiler, which when activated, provide event information to a profiling agent. The profiling agent is implemented as an integrated language environment (ILE) service program. The profiling agent sends control information to the JVM for enabling and disabling JVMTI events.

For example, the profiler may not be interested in method Entry or Exit hooks and could tell the JVM that it does not want to receive these event notifications. The JVM and JIT have JVMTI event hooks embedded that send event notifications to the profiling agent if the event is enabled. The profiler tells the JVM which events are of interest, and the JVM sends notifications of the events to the profiler when they occur.

A service program, called QJVAJVMTI, which resides in the QSYS library, supports the JVMTI functions.

14.1.5 Java Secure Socket Extension

Java Secure Socket Extension (JSSE) is like a framework that abstracts the underlying mechanisms of both SSL and TLS. By abstracting the complexity and peculiarities of the underlying protocols, JSSE enables programmers to use secure, encrypted communications while at the same time minimizing possible security vulnerabilities. Java Secure Socket Extension (JSSE) uses both the Secure Sockets Layer (SSL) protocol and the Transport Layer Security (TLS) protocol to provide secure, encrypted communications between clients and servers.

Java Secure Socket Extension 6 (JSSE 6) is now available with IBM i 6.1.

When you use the Java 2 Platform, Standard Edition (J2SE) version 6 on your IBM i, JSSE is already configured. The default configuration uses the native IBM i JSSE provider.

For more information on configuring JSSE, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzhaha/rzahajsseconfig16.htm>

14.1.6 New Java CL commands

Several Java-related CL commands have been added for IBM i 6.1:

Work with JVM Jobs (WRKJVMJOB)

The Work with the JVM Jobs (WRKJVMJOB) command can be used to collect performance data, and monitor Java Virtual Machines (JVMs) running in active jobs. The following information or functionality is available when using WRKJVMJOB:

- ▶ The arguments and options with which the JVM was started
- ▶ Environment variables for both ILE and PASE
- ▶ Java lock requests outstanding for the JVM job
- ▶ Garbage collection information
- ▶ Java system properties
- ▶ The list of threads associated with the JVM
- ▶ The partially completed job log for the JVM job
- ▶ The ability to work with spooled input and output files for the JVM job
- ▶ The ability to generate JVM (system, heap, and Java) dumps from a panel option

- ▶ The ability to enable and disable verbose garbage collection from a panel option

The Generate JVM Dump (GENJVMDMP) command

The Generate JVM Dump (GENJVMDMP) command can be used to generate Java, system, and heap dumps. The following types of dumps can be generated:

- *JAVA Generates multiple files that contain diagnostic information for the JVM and the Java applications running in the JVM.
- *SYSTEM Generates a binary format raw memory image of the job that was running when the dump was initiated.
- *HEAP Generates a dump of all the heap space allocations which have not yet been freed.

Note: The issuer of the GENJVMDMP command must be running under a user profile which is the same as the job user identity of the job being worked with, or the issuer of the command must be running under a user profile which has job control (*JOBCTL) special authority.

The Print JVM Job (PRTJVMJOB) command

The Print JVM Job (PRTJVMJOB) command can be used to print Java Virtual Machines (JVMs) running in active jobs.

Note: WRKJVMJOB and GENJVMDMP can be used if you are using the IBM Technology for Java Virtual Machine.

14.1.7 JDBC 4.0

The DatabaseMetaData interface is implemented by the IBM Developer Kit for Java JDBC driver to provide information about its underlying data sources. It is used primarily by application servers and tools to determine how to interact with a given data source. Applications may also use DatabaseMetaData methods to obtain information about a data source, but this is less typical.

JDBC 4.0 is compliant with J2SE 6. Changes were made to the MetaData interface in JDBC 4.0, as follows:

java.sql.DatabaseMetadata changes

In IBM i 6.1, the command language interface (CLI) changes the implementation of the MetaData APIs to also call the SYSIBM stored procedures. Because of this, the JDBC MetaData methods will use the SYSIBM procedures directly on IBM i 6.1, regardless of the JDK level. You will notice the following differences due to this change:

- ▶ The native JDBC driver previously permitted the user `localhost` as the catalog name for most of the methods. In JDBC 4.0, the native JDBC driver will not return any information if `localhost` is specified.
- ▶ The native JDBC driver always returned an empty result set when the nullable parameter of `getBestRowIdentifier` was set to `false`. This parameter now returns the correct result.
- ▶ The values returned by `getColumns` for the `BUFFER_LENGTH`, `SQL_DATA_TYPE`, and `SQL_DATETIME_SUB` columns may be different. These values should not be used in a JDBC application because the JDBC specification defines these columns as unused.

- ▶ The native JDBC driver previously recognized the table and schema parameters of `getCrossReference`, `getExportedKeys`, `getImportedKeys`, and `getPrimaryKeys` as a pattern. Now, the table and schema parameters must match the name as stored in the database.
- ▶ Views used for implementing system-defined views were previously described by `getTables()` as SYSTEM TABLES. To be consistent with the DB2 family, these views are now described as VIEWS.
- ▶ The remarks column for `getProcedures` used to return an empty string if no information was available. In this case, the column now contains null. This behavior is consistent with the DB2 family behavior.
- ▶ Column names returned by `getProcedures` are different. These column names are not defined by the JDBC 4.0 specification. Table 14-1 lists the new column names returned by `getProcedures`.

Table 14-1 New column names returned by `getProcedures` in JDBC 4.0

Column #	Previous name	Name under JDBC 4.0
4	RESERVED1	NUM_INPUT_PARAMS
5	RESERVED2	NUM_OUTPUT_PARAMS
6	RESERVED3	NUM_RESULT_SETS

New QueryCloseImplicit JDBC connection property

IBM i 6.1 introduces performance improvements, including the new native JDBC driver connection property, `QueryCloseImplicit`. By default, this new property will be enabled. Using it provides significant performance improvements for many applications that use native JDBC to query the IBM i database.

Although it is unlikely, with this new connection attribute, applications might see slightly different semantics in certain error scenarios. These error scenarios are reportable to IBM as problems, with or without the new function. With the new function, the error on the `CLOSE` call, which should never happen except in reportable problem situations, is reported on the `ResultSet next()` method call in addition to the `ResultSet close()` method to the JDBC application. Without the new function, the error is only returned on the `ResultSet close()` method.

14.1.8 Support for multiple IBM Developer Kit for Java options

IBM i supports multiple versions of the Java Development Kits (JDKs) and the Java 2 Platform, Standard Edition.

IBM i supports using multiple JDKs simultaneously, but only through multiple Java virtual machines. A single Java virtual machine runs one specified JDK. Only one Java virtual machine can run per job.

For information about how to install more than one JDK at a time, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzhaha/install.htm>

With the Classic JVM, the `java.version` system property determines which JDK to run. Once a Java virtual machine is up and running, changing the `java.version` system property has no effect. If you are using IBM Technology for Java, you select which 5761-JV1 option to run (and

therefore which JDK/bit mode) by setting the JAVA_HOME environment variable. This differs from the Classic JVM, which can use the java.version system property as input. Example 14-1 gives an example for setting the JVM option to run.

Example 14-1 Changing JAVA_HOME variable

To set Java SE 6 32 bit as the default JVM:
 ADDENVVAR ENVVAR(JAVA_HOME) VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit')

Table 14-2 lists the supported options for this release.

Table 14-2 5761-JV1 supported options in IBM i 6.1

5761-JV1 options	JAVA_HOME	java.version
Option 6 - Classic 1.4	/QIBM/ProdData/Java400/jdk14/	1.4
Option 7 - Classic 5.0	/QIBM/ProdData/Java400/jdk15/	1.5
Option 8 - IBM Technology for Java 5.0 32-bit	/QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit	1.5
Option 9 - IBM Technology for Java 5.0 64-bit	/QOpenSys/QIBM/ProdData/JavaVM/jdk50/64bit	1.5
Option 10 - Classic 6	/QIBM/ProdData/Java400/jdk6	1.6
Option 11 - IBM Technology for Java 6 32-bit	/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit	1.6
Option 12 - IBM Technology for Java 6 64-bit	/QOpenSys/QIBM/ProdData/JavaVM/jdk60/64bit	1.6

The default JDK chosen in this multiple JDK environment depends on which 5761-JV1 options are installed. The following order of precedence determines the default JDK:

1. Option 8 - IBM Technology for Java 5.0 32-bit
2. Option 9 - IBM Technology for Java 5.0 64-bit
3. Option 7 - Classic 5.0
4. Option 11- IBM Technology for Java 6 32-bit
5. Option 12 - IBM Technology for Java 6 64-bit
6. Option 10 - Classic 6
7. Option 6 - Classic 1.4

Figure 14-1 on page 324 shows IBM i with multiple JDK environments installed.

Display Installed Licensed Programs			System: RCHAS70
Licensed Program	Product Option	Description	
5722DR1	*BASE	IBM DIRECTOR	
5722IP1	*BASE	IBM Infoprint Server for iSeries	
5722IP1	1	PS to AFP Transform	
5761JC1	*BASE	IBM Toolbox for Java	
5761JS1	*BASE	IBM Advanced Job Scheduler for i5/OS	
5761JV1	*BASE	IBM Developer Kit for Java	
5761JV1	6	Java Developer Kit 1.4	
5761JV1	7	Java Developer Kit 5.0	
5761JV1	8	J2SE 5.0 32 bit	
5761JV1	9	J2SE 5.0 64 bit	
5761JV1	10	Java Developer Kit 6	
5761JV1	11	Java SE 6 32 bit	
5761JV1	12	Java SE 6 64 bit	
5761LSV	*BASE	Extended Integrated Server Support	

Figure 14-1 IBM i 6.1 5761-JV1 installed options

14.1.9 Installing IBM Technology for Java Virtual Machine

IBM Technology for Java Virtual Machine is available in both 32-bit and 64-bit versions. Use these instructions to install IBM Technology for Java Virtual Machine, which is included in licensed program 5761-JV1. Licensed program 5761-JV1 is shipped with the IBM i installation media. To access the IBM Technology for Java option, perform the following steps:

1. Enter the Go Licensed Program (GO LICPGM) command and select Option 10 (Display).
2. If you do not see 5761JV1 listed, perform the following steps:
 - a. Enter the GO LICPGM command on the command line.
 - b. Select Option 11 (Install licensed program).
 - c. Choose Option 1 (Install) for licensed program 5761-JV1 *BASE.
 - d. Repeat the previous step for each additional option you want to install.
3. Add the appropriate environment variable. At a command line, enter one of the following commands:
 - a. ADDENVVAR ENVVAR(JAVA_HOME)
VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit')
 - b. ADDENVVAR ENVVAR(JAVA_HOME)
VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk50/64bit')
 - c. ADDENVVAR ENVVAR(JAVA_HOME)
VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit')
 - d. ADDENVVAR ENVVAR(JAVA_HOME)
VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk60/64bit')

If you are unsure what JVM you are currently using, you can check using the following methods. If you see IBM J9 VM in the result, you are using IBM Technology for Java.

- ▶ Look in the job log for the job containing the JVM. There will be a message that states what JVM you are using. A WebSphere Application Server V6.1 application server profile example is shown in Example 14-2 on page 326.


```

Message ID . . . . . : JVAB302      Severity . . . . . : 10
Message type . . . . . : Diagnostic
Date sent . . . . . : 11/14/08      Time sent . . . . . : 14:42:5

Message . . . . . : Java Virtual Machine is IBM Technology for Java.
                    PID(1451)
Cause . . . . . : JAVA_HOME environment variable is
                  /QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit

```

Figure 14-2 Identifying the Java version from a job log

- ▶ As part of the Java command you are using to run your application, add **-showversion**. You will see one additional line that shows the JVM you are using.
- ▶ From **qsh** or **qp2term**, run **java -version**.

Note: Before migrating to or from the Classic JVM, which is a 64-bit virtual machine, to the 32-bit version of IBM Technology for Java, consider that there may be limitations when using the 32-bit environment. For example, the amount of addressable memory is much smaller. In 32-bit mode, the Java object heap cannot grow much larger than 3 gigabytes. You will also be limited to running approximately 1000 threads.

IBM 64-bit Java Virtual Machine

The IBM Java Virtual Machine (JVM) is extended with the availability of the IBM technology for Java 64-bit JVM with IBM i 6.1. The 64-bit JVM delivers improved performance for Java-based applications and helps solution providers with application portability, configuration and service-tool compatibility. The benefits of the new 64-bit JVM are realized when deploying new Java applications on IBM i or when migrating existing Java applications that have been running in the classic JVM.

IBM 64-bit JVM uses the same underlying implementation as 32-bit, but with no practical limit to heap size. In IBM 64-bit JVM object references are 64 bits instead of 32 bits.

IBM 64-bit Java virtual machine provides the following benefits:

- ▶ Helps solution providers with application portability.
- ▶ Works with popular industry tools.
- ▶ Improves performance over the classic JVM.

For JVM performance information and information on how to choose the correct JVM for your use, refer to Chapter 6, “Performance” on page 125.

14.1.10 List of Java system properties

Java system properties determine the environment in which the Java programs run. They are similar to system values or environment variables in IBM i. They are available through a properties file.

Starting a JVM sets the system properties for that instance of the JVM.

You can choose to use the default values for Java system properties or you can specify values for them by using the following methods:

- ▶ Adding parameters to the command line (or the Java Native Interface (JNI) invocation API) when you start the Java program.
- ▶ Using the QIBM_JAVA_PROPERTIES_FILE job-level environment variable to point to a specific properties file, as shown in Example 14-2.

Example 14-2 Setting QIBM_JAVA_PROPERTIES_FILE value

```
ADDENVVAR ENVVAR(QIBM_JAVA_PROPERTIES_FILE)
          VALUE(/QIBM/userdata/java400/mySystem.properties)
```

- ▶ Creating a SystemDefault.properties file that you create in your user.home directory.
- ▶ Using the /QIBM/UserData/java400/SystemDefault.properties file.

IBM i and the JVM determine the values for Java system properties by using the following order of precedence:

1. Command line or JNI invocation API
2. QIBM_JAVA_PROPERTIES_FILE environment variable
3. user.home SystemDefault.properties file
4. /QIBM/UserData/Java400/SystemDefault.properties
5. Default system property values

- SystemDefault.properties file

The SystemDefault.properties file is a standard Java properties file that enables you to specify default properties of your Java environment.

- List of Java system properties

For the complete list of Java system properties, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzaha/sysprop2.htm>

14.1.11 Java Runtime.exec

The IBM Technology for Java Virtual Machine has changed the default for output from the Java Runtime.exec() function. When this function runs as an IBM i PASE executable program, the output of the child process defaults to the character coding that is specified by the file.encoding property. This property is normally an ASCII-character encoding. In the previous release, the output of the child process used an EBCDIC-character encoding that matched the job CCSID. This changed behavior enables your Java application to read output from the child process without accounting for the CCSID differences. This change does not affect cases where you use Runtime.exec() to call an ILE program, whose default output might still be EBCDIC.

If you want the Runtime.exec() of a PASE executable program to output EBCDIC character encoding, set the environment variable QIBM_JAVA_EXEC_CONV=T.

14.1.12 Miscellaneous changes

The following capabilities are available to Java programs running under IBM i 6.1:

Internet Protocol version 6

Internet Protocol version 6 (IPv6) is now fully supported by both the IBM Developer Kit for Java Classic and IBM Technology for Java JVMs. For more information on IPv6 concepts, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzai2/rzai2ipv6concepts.htm>

Time zone configuration

Time zone configuration has changed in IBM i 6.1. The `user.timezone` system property setting provided by the `QTIMZON` system value can be overridden by either specifying the `user.timezone` value explicitly on the command line or in the `SystemDefault.properties` file or by using `LOCALE` to set the `user.timezone` value. The `LOCALE` method was used before `QTIMZON` system value support was available on IBM i, but is no longer recommended.

Native Abstract Windowing Toolkit information

When using JDK 1.4, Native Abstract Windowing Toolkit (NAWT) does not support locale- and language-specific fonts and character sets. When using NAWT, ensure that you comply with the following requirements:

- ▶ Use only characters that are defined in the ISO8859-1 character set.
- ▶ Use the `font.properties` file. The `font.properties` file resides in the `/QIBM/ProdData/Java400/jdknn/lib` directory, where `nnn` is the version number of the J2SE that you are using. Do not use any of the `font.properties.xxx` files, where `xxx` is a language or another qualifier.

When using JDK version 1.5 and later, the fonts and character sets used are determined at Java startup time by the value of the `file.encoding` property. If `file.encoding` is not explicitly set, it is given an appropriate default value based on the job CCSID.

System debugging for IBM Technology for Java

New debugging information has been added for the IBM Technology for Java JVM. For more information, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rza/aha/rzahadebugit4j.htm>

14.2 IBM Toolbox for Java

IBM Toolbox for Java is a set of Java classes that allow you to use Java programs to access data on your system. You can use these classes to write client/server applications, applets, and servlets that work with data on your system. You can also run Java applications that use the IBM Toolbox for Java classes on the IBM i Java virtual machine (JVM).

IBM Toolbox for Java uses the IBM i Host Servers (for example, the `QZDASOINIT` jobs for ODBC and JDBC functions and the `QZRCSRVS` jobs for remote command execution) as access points to the system. Because IBM Toolbox for Java uses communication functions built into Java, you do not need to use System i Access for Windows to use IBM Toolbox for

Java. Each server runs in a separate job on the server, and each server job sends and receives data streams on a socket connection.

IBM Toolbox for Java is available in the following forms:

- ▶ The licensed program for IBM Toolbox for Java, 5761-JC1 installs on IBM i 5.4 and later versions of IBM i. From a client, IBM Toolbox for Java connects back to i5/OS V5R3 and subsequent versions of IBM i.
- ▶ IBM i also includes the non-graphical classes of IBM Toolbox for Java optimized for use when running IBM Toolbox for Java classes on the IBM i Java virtual machine (JVM). So if, for example, you do not have a need for the graphical functionality of the licensed program, you can still easily use IBM Toolbox for Java.
- ▶ IBM Toolbox for Java is also available in an open source version. You can download the code and get more information from the JTOpen Web site at:

<http://jt400.sourceforge.net>

14.2.1 Using IBM i enhancements in the IBM Toolbox for Java JDBC

The IBM Toolbox for Java JDBC can utilize the following enhancements in IBM i:

- ▶ Support for JDBC 4.0
- ▶ Query storage limit
- ▶ Decimal float (DECFLOAT) data type
- ▶ Passing the client type and application name to the server
- ▶ Maximum length of cursor names extended
- ▶ Generated key support
- ▶ Improved default value support
- ▶ Increased maximum in GROUP BY clause
- ▶ Batch update support

For details about each of these IBM i enhancements, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahh/rzahhjdbcenhance08.htm>

14.2.2 New classes

The following classes have been added. All classes listed here are in package `com.ibm.as400.access`, unless otherwise noted.

The AS400JDBCManagedConnectionPoolDataSource class

The JDBC connection pool is a managed connection pool that allows you to reuse connections.

The `AS400JDBCManagedConnectionPoolDataSource` class simplifies connection pool maintenance by eliminating the need for user applications to implement their own management code. The connection pool is used in combination with the JDBC `DataSource` interface mechanism.

The `AS400JDBCManagedConnectionPoolDataSource` class contains an implementation of the `javax.sql.DataSource` interface. This class manages the functionality of the connection

pooling mechanism and allows you to access it like a normal DataSource. Setting up and using the self-managed JDBC connection pool manager involves the usage of Java Naming and Directory Interface (JNDI) to bind the connection pool managed datasource. In order for self-managed JDBC connection pooling to function, JNDI is necessary.

The following properties allow you to configure the attributes of the connection pool:

- ▶ InitialPoolSize
- ▶ MinPoolSize
- ▶ MaxPoolSize
- ▶ MaxLifetime
- ▶ MaxIdleTime
- ▶ PropertyCycle
- ▶ ReuseConnections

The AS400JDBCManagedDataSource class

This class can be used to produce Connection objects that will automatically participate in connection pooling, and are managed by the Toolbox's built-in connection pooling manager.

The FileAttributes class

This class represents the set of file attributes that can be retrieved and set with the Get Attributes (Qp0IGetAttr) and Set Attributes (Qp0ISetAttr) APIs.

The object must exist and the caller must have authority to it. Only attributes supported by the specific file system, object type, and system operating release can be retrieved or set.

The HistoryLog class

This class represents a history log on the system. It is used to get a list of messages in a history log. The close() method should be called when you are finished retrieving messages from the list.

The ObjectReferences class

This class represents the set of information about integrated file system references on an object that can be retrieved with the Retrieve Object References (QP0LROR) API.

A reference is an individual type of access or lock obtained on the object when using integrated file system interfaces. An object may have multiple references concurrently held, provided that the reference types do not conflict with one another.

This class will not return information about byte range locks that may currently be held on an object.

The user must have execute (*X) data authority to each directory preceding the object whose references are to be obtained. The user must have read (*R) data authority to the object whose references are to be obtained.

UDFS class

This class represents a user-defined file system on the system. It can be used for creating a file system that can be made visible to the rest of the integrated file system name space with the mount() method. It also returns information about a mounted file system or UDFS.

14.2.3 Enhanced classes

Significant enhancements were made to the following classes. All classes listed here are in package `com.ibm.as400.access`.

CommandCall for Unicode enabled commands

Represents an IBM i command object. This class allows the user to call any non-interactive CL command. Results of the command are returned in a message list.

The IFSFile class

This class represents an object in the integrated file system of IBM i. As in `java.io.File`, `IFSFile` is designed to work with the object as a whole. For example, use `IFSFile` to delete or rename a file, to access the file's attributes (is the object a file or a directory, when was the file last changed, is the file hidden, and so forth), or to list the contents of a directory. Use `IFSFileInputStream` or `IFSRandomAccessFile` to read data from the file, and `IFSFileOutputStream` or `IFSRandomAccessFile` to write data to the file.

The ISeriesNetServer class

This class represents the `NetServer` service on a system. It allows a user to query and modify the state and configuration of the `NetServer`.

If the `NetServer` job on the i5/OS system is not started, the list methods may return incomplete results. To determine whether the `NetServer` job is started, use the `isStarted` method. To start the `NetServer`, use the `start` method.

The Permission class

The `Permission` class is provided to retrieve a user's authority to an object.

To improve performance, the `Permission` object caches authority changes until the `commit()` method is called. When `commit()` is called, all changes up to that point are sent to the system.

The ObjectList class

The `ObjectList` class represents a list of system objects in a specific library, multiple libraries, or system-wide.

The ObjectDescription class

This class represents a system `QSYS` object and its attributes. `ObjectDescription` objects can be constructed individually, or generated from an `ObjectList`. An object's attributes can be retrieved by calling `getValue()` and passing one of the integer attribute constants defined in this class.

The User class

This class represents a user profile object on the system. It internally calls the Retrieve User Information (`QSYRUSRI`) API for the methods that retrieve user profile information.

14.2.4 Compatibility

The IBM Toolbox for Java no longer ships the `x4j400.jar` (IBM XML parser). It is recommended that applications use one of the following JAXP-compliant XML parsers:

- ▶ The XML parser that is built into JDK 1.4 and higher
- ▶ The Apache Xerces XML parser available from xml.apache.org

- ▶ One of the XML parsers that ship on IBM i under /QIBM/ProdData/OS400/xml/lib

IBM Toolbox for Java no longer supports running in the default JVM in Netscape Navigator or Microsoft Internet Explorer. Running applets that use IBM Toolbox for Java classes in a browser requires that you install a plug-in such as the most recent Sun Java 2 Runtime Environment (JRE) plug-in.

IBM Toolbox for Java no longer includes data400.jar. The classes that were in data400.jar are now in jt400.jar. Remove data400.jar from your CLASSPATH statements.

You cannot use this release of IBM Toolbox for Java to deserialize some objects that you serialized using releases before V5R1.

If you are using Secure Sockets Layer (SSL) to encrypt data flowing between the client and the server, you must use Java Secure Socket Extension (JSSE).

To use all the IBM Toolbox for Java classes, use the Java 2 Platform, Standard Edition (J2SE). Using the vaccess classes or the Graphical Toolbox requires using the Swing package, which comes with J2SE. Using PDML requires that you run version 1.4 or later of the Java Runtime Environment.

For more information about the requirements for running IBM Toolbox for Java, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahh/rzahnm0402.htm>

14.3 Websphere Development Studio

WebSphere Development Studio (WDS) for IBM i offers significant enhancements to the IBM i host Integrated Language Environment (ILE) compilers.

This release includes a major repackaging of its components, as shown in Figure 14-3 on page 332.

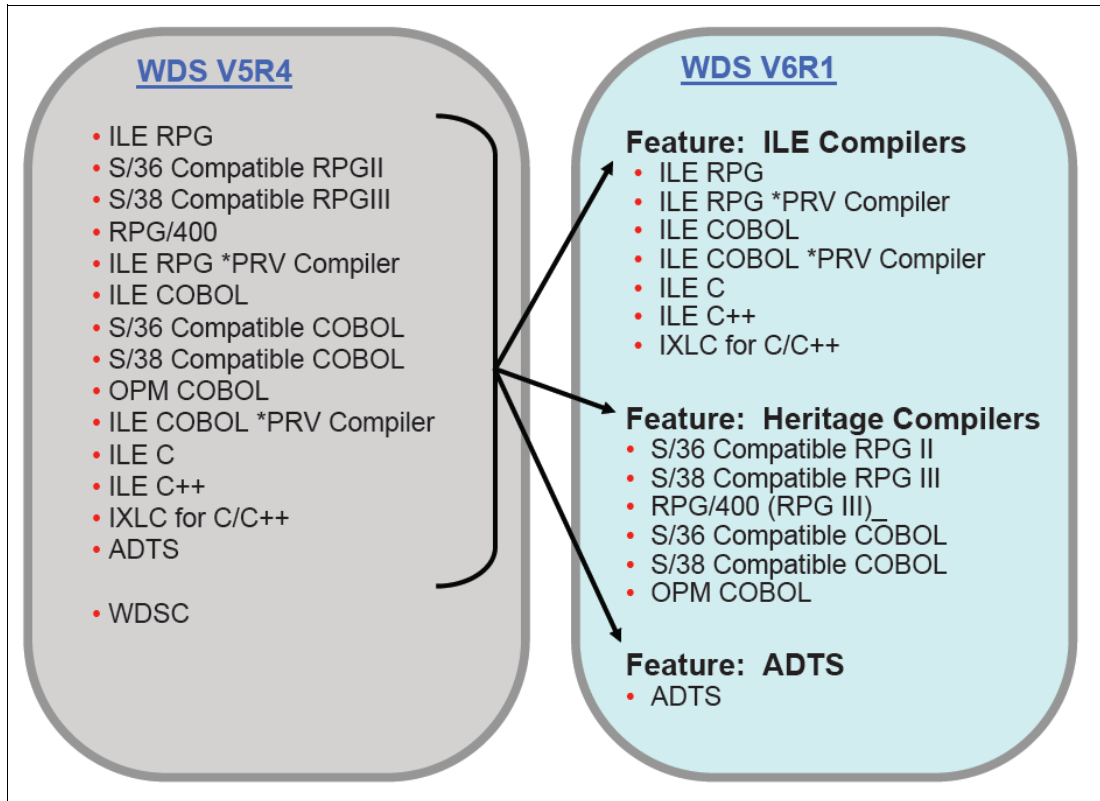


Figure 14-3 New packaging for WebSphere Development Studio in IBM i 6.1

The product no longer bundles all application development tools and compilers together. Instead, it contains three options that are ordered separately. The components of each option are provided in Table 14-3.

Table 14-3 WDS features in IBM i 6.1

Option	Components
ILE Compilers	<ul style="list-style-type: none"> ▶ Option 31 ILE RPG ▶ Option 35 ILE RPG *PRV compiler ▶ Option 41 ILE COBOL ▶ Option 45 ILE COBOL *PRV ▶ Option 51 ILE C ▶ Option 52 ILE C++ ▶ Option 56 IXLC for C/C++
Heritage Compilers	<ul style="list-style-type: none"> ▶ Option 32 S/36 compatible RPG II ▶ Option 33 S/38 compatible RPG III ▶ Option 34 RPG/400® ▶ Option 42 S/36 compatible COBOL ▶ Option 43 S/38 compatible COBOL ▶ Option 44 OPM COBOL
Application Development Toolset	<ul style="list-style-type: none"> ▶ Option 21 Application Development Toolset

Also included in this new release are enhancements to the ILE RPG, COBOL, and C and C++ compilers.

As of IBM i 6.1, the WebSphere Development Studio Client is no longer included as part of the WebSphere Development Studio for IBM i software.

14.3.1 ILE RPG compiler

The ILE RPG compiler is part of the ILE Compilers option of IBM WebSphere Development Studio for System i. It is used for compilation of RPG IV, the latest generation of RPG.

The following list describes the enhancements to ILE RPG in IBM i 6.1.

ILE RPG can run multi-threaded

When `THREAD(*CONCURRENT)` is specified on the Control specification of a module, it has the following attributes:

- ▶ It can run concurrently in multiple threads.
- ▶ Multiple threads can run in the module at the same time.
- ▶ By default, static variables will be defined so that each thread will have its own copy of the static variable.
- ▶ Individual variables can be defined to be shared by all threads using `STATIC(*ALLTHREAD)`.
- ▶ Individual procedures can be serialized so that only one thread can run them at one time, by specifying `SERIALIZE` on the Procedure-Begin specification.

Define a main procedure that does not use the RPG cycle

Using the `MAIN` keyword on the Control specification, a subprocedure can be identified as the program entry procedure. This allows an RPG application to be developed where none of the modules uses the RPG cycle, and allows an RPG program to be called recursively.

The prototype for the main subprocedure must have the `EXTPGM` keyword; the main subprocedure can only be called by a program call.

Note: Other than being the program-entry procedure, the main subprocedure is like any other subprocedure. It does not use the RPG cycle.

Program Call Markup Language can be stored in the module

Program Call Markup Language (PCML) can now be stored in the module as well as in a stream file. By using combinations of the `PGMINFO` command parameter and/or the new `PGMINFO` keyword for the Control specification, the RPG programmer can choose where the PCML information should go. If the PCML information is placed in the module, it can later be retrieved using the `QBNRP11` API.

Note: This enhancement is also available through PTFs for IBM i 5.4, but only through the Control specification keyword.

For a complete list of new and changed language elements since IBM i 5.4, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzasd/sc09250825.htm&tocNode=int_187807

Files defined in subprocedures

Files can be defined locally in subprocedures. I/O to local files can only be done with data structures; I and O specifications are not allowed in subprocedures, and the compiler does not generate I and O specifications for externally described files.

By default, the storage associated with local files is automatic; the file is closed when the subprocedure returns. The `STATIC` keyword can be used to indicate that the storage associated with the file is static, so that all invocations of the subprocedure will use the same file, and if the file is open when the subprocedure returns, it will remain open for the next call to the subprocedure.

Qualified record formats

When a file is defined with the `QUALIFIED` keyword, the record formats must be qualified by the file name, `MYFILE.MYFMT`. Qualified files do not have I and O specifications generated by the compiler; I/O can only be done through data structures.

When files are qualified, the names of the record formats can be the same as the formats of other files. For example, you can have `FILE1.FMT1` and `FILE2.FMT1`.

Files defined like other files

Using the `LIKEFILE` keyword, a file can be defined to use the same settings as another File specification, which is important when passing a file as a parameter. If the file is externally-described, the `QUALIFIED` keyword is implied. I/O to the new file can only be done through data structures.

Files passed as parameters

A prototyped parameter can be defined as a File parameter using the `LIKEFILE` keyword. Any file related through the same `LIKEFILE` definition may be passed as a parameter to the procedure. Within the called procedure or program, all supported operations can be done on the file; I/O can only be done through data structures.

EXTDESC keyword and EXTFILE(*EXTDESC)

The `EXTDESC` keyword identifies the file to be used by the compiler at compile time to obtain the external description of the file; the filename is specified as a literal in one of the forms `'LIBNAME/FILENAME'` or `'FILENAME'`. This removes the need to provide a compile-time override for the file.

The `EXTFILE` keyword is enhanced to allow the special value `*EXTDESC`, indicating that the file specified by `EXTDESC` is also to be used at runtime.

EXTNAME library support

The `EXTNAME` keyword is enhanced to allow a literal to specify the library for the external file. `EXTNAME('LIBNAME/FILENAME')` or `EXTNAME('FILENAME')` are supported. This removes the need to provide a compile-time override for the file.

EXFMT allows a result data structure

The `EXFMT` operation is enhanced to allow a data structure to be specified in the result field. The data structure must be defined with usage type `*ALL`, either as an externally-described data structure for the record format (`EXTNAME(file:fmt:*ALL)`), or using `LIKEREC` of the record format (`LIKEREC(fmt:*ALL)`).

Larger limits for data structures, character, UCS-2 and graphic variables

Data structure limits have changed as follows:

- ▶ Data structures can have a size up to 16,773,104.
- ▶ Character definitions can have a length up to 16,773,104. The limit is 4 less for variable-length character definitions.
- ▶ Character definitions can have a length up to 16,773,104. The limit is 4 less for variable-length character definitions.
- ▶ UCS-2 definitions can have a length up to 8,386,552 UCS-2 characters. The limit is 2 less for variable-length UCS-2 definitions.
- ▶ Graphic definitions can have a length up to 8,386,552 DBCS characters. The limit is 2 less for variable-length graphic definitions.
- ▶ The VARYING keyword allows a parameter of either 2 or 4 indicating the number of bytes used to hold the length prefix.

%ADDR(varying: *DATA)

The %ADDR built-in function is enhanced to allow *DATA as the second parameter to obtain the address of the data part of a variable-length field.

Limits for DIM and OCCURS and for character, UCS-2 and DBCS literals

The following limits have increased:

- ▶ An array or multiple-occurrence data structure can have up to 16,773,104 elements, provided that the total size is not greater than 16,773,104.
- ▶ Character literals can now have a length up to 16380 characters.
- ▶ UCS-2 literals can now have a length up to 8190 UCS-2 characters.
- ▶ Graphic literals can now have a length up to 16379 DBCS characters.

TEMPLATE keyword for files and definitions

The TEMPLATE keyword can be coded for file and variable definitions to indicate that the name will only be used with the LIKEFILE, LIKE, or LIKEDS keyword to define other files or variables. Template definitions are useful when defining types for prototyped calls, since the compiler only uses them at compile time to help define other files and variables, and does not generate any code related to them.

Template data structures can have the INZ keyword coded for the data structure and its subfields, which will ease the use of INZ(*LIKEDS).

Relaxation of some UCS-2 rules

The compiler will perform some implicit conversion between character, UCS-2 and graphic values, making it unnecessary to code %CHAR, %UCS2 or %GRAPH in many cases. This enhancement is also available through PTFs for i5/OS V5R3 and IBM i 5.4. Implicit conversion is now supported for:

- ▶ Assignment using EVAL and EVALR
- ▶ Comparison operations in expressions
- ▶ Comparison using fixed form operations IFxx, DOUxx, DOWxx, WHxx, CASxx, CABxx, COMP

Note: Implicit conversion was already supported for the conversion operations MOVE and MOVEL.

UCS-2 variables can now be initialized with character or graphic literals without using the %UCS2 built-in function.

Eliminate unused variables from the compiled object

New values *UNREF and *NOUNREF are added to the OPTION keyword for the CRTBNDRPG and CRTRPGMOD commands, and for the OPTION keyword on the Control specification. The default is *UNREF. *NOUNREF indicates that unreferenced variables should not be generated into the RPG module. This can reduce program size, and if imported variables are not referenced, it can reduce the time taken to bind a module to a program or service program.

14.3.2 ILE COBOL

The ILE COBOL compiler gives you powerful functions to enhance your programming tasks. The following list describes the enhancements made to ILE COBOL in IBM i 6.1:

National UCS-2 CCSID support

The NTLCCSID parameter has been added to the CRTCBMOD and CRTBNDCBL commands, and to the PROCESS statement, to allow you to specify the UCS-2 CCSID to be used for National data items. With this parameter, you can specify a CCSID other than the default 13488, such as CCSID 1200, to be used for National items.

PCML in module support

The PGMINFO parameter on the CRTCBMOD and CRTBNDCBL commands has been enhanced to allow you to specify the location where you want to put the generated PCML. When the user specifies *PCML as the first parameter for the PGMINFO keyword, a second parameter specifying a location of *STMF, *MODULE, or *ALL, can also be specified:

- *STMF Causes the PCML to be put into the stream file specified on the INFOSTMF parameter.
- *MODULE Causes the PCML to be put into the generated module.
- *ALL Causes the PCML to be put in all of these locations.

PROCESS statement option PGMINFO

This option allows the user to request that PCML be added to the module, and can be specified as PGMINFO(PCML MODULE). If the user had requested the PCML be added to a streamfile from the create command, the PCML will be added to both the module and the streamfile.

Complex OCCURS DEPENDING ON debugger support

Support has been added so that the system debugger and the client debugger can now debug complex OCCURS DEPENDING ON (ODO) arrays.

Large program support

The compiler has been enhanced so that larger programs and programs containing a very large number of data items can now be compiled.

Note: Large program support is also available in i5/OS V5R3 and IBM i 5.4 with PTFs:

- ▶ i5/OS V5R3: SI27834
- ▶ IBM i 5.4: SI26591

14.3.3 ILE C and ILE C++

The ILE C compiler is a high-performance, high-function compiler for C development on IBM i. The ILE C++ compiler provides additional features not found in the C language. In IBM i 6.1 both compilers were enhanced with the following new features and modifications:

Decimal floating-point support

Decimal floating support, available only in ILE C, contains the following features:

- ▶ Allows floating-point computations to be performed using decimal arithmetic (base 10).
- ▶ Avoids potential rounding errors converting binary floating-point data to/from human readable formats.
- ▶ Conforms to the proposed decimal formats and arithmetic described in the IEEE 754 draft.
- ▶ Includes support emulated on POWER5 processors; supported directly in hardware on POWER6 processors.
- ▶ Adds support to ILE C compiler and runtime, based on Draft Technical Report 24732 submitted to the ISO/IEC JTC1/SC22/WG14—Programming Language C committee.
- ▶ New data types:
 - `_Decimal32`, 4 bytes, 7 digits precision, -95/+96 exponent
 - `_Decimal64`, 8 bytes, 16 digits precision, -383/+384 exponent
 - `_Decimal128`, 16 bytes, 34 digits precision, -6143/+6144 exponent
- ▶ Includes new decimal floating-point literals: `1.0df`, `1.0dd`, and `1.0dl`.
- ▶ Converts to/from integers, binary floating-point, packed decimal types.
- ▶ Includes `DECFLTRND` option on C compiler commands (`CRTCMOD`, `CRTBNDC`) to control compile-time decimal floating-point rounding mode.
- ▶ Includes new length modifiers for the `scanf/printf` family of functions:
 - `printf("dfp32 value = %HE\n", dfp32)`
 - `printf("dfp64 value = %DE\n", dfp64)`
 - `printf("dfp128 value = %DDE\n", dfp128)`
- ▶ Includes new `strtod32`, `strtod64`, `strtod128` APIs defined.
- ▶ Includes new decimal floating-point macros in `float.h` and `math.h`: `DEC32_MIN`, `DEC32_MAX`, `DEC_INFINITY`, `DEC_NAN`, and others.

Thread local storage support

Thread local storage support, for both ILE C and ILE C++, has the following features:

- ▶ Includes language level support of thread local variables.
- ▶ Includes similar support to that provided by the POSIX `pthread_getspecific` and `pthread_setspecific` APIs, except easier to use and faster, since the support is integrated into the operating system and the language.

- ▶ Makes threaded programming in IBM i easier by providing a simple method for restricting access to a variable to a single thread.
- ▶ Adds support to the ILE C and ILE C++ compilers.
- ▶ Includes any variable declared with static storage duration can be changed to a thread local variable. This includes global variables and static local variables. Automatic variables cannot be declared as thread local variables.
- ▶ Includes the ability to declare a thread local variable by adding the `__thread` storage specifier to the declaration of the variable. Example 14-3 contains two thread local variables:

Example 14-3 Declare thread local variable

```
__thread int tl_global = 0;
int threadfunc(void) {
    static __thread int tl_static_auto = 0;
    ...
}
```

ILE C and C++ compiler INLINE option changes

For the Create Bound C Program (CRTBNDC) and Create C Module (CRTCMOD) commands, the default mode for the `INLINE` option has changed from `*NOAUTO` to `*AUTO`. If the `INLINE` option is set to `*ON` and no inline mode is specified, the inline mode is `*AUTO` instead of `*NOAUTO`. The use of `*AUTO` permits additional functions to be inlined that were not previously inlined. Because the `INLINE` option is set to `*OFF` by default, only objects created with the `INLINE` option set to `*ON` and with no inline mode specified are affected by this change.

For the Create Bound C++ Program (CRTBNDCPP) and Create C++ Module (CRTCPPMOD) commands, the default mode for the `INLINE` option has changed from `*NOAUTO` to `*AUTO`. If the `INLINE` option is set to `*ON` and no inline mode is specified, the inline mode is `*AUTO` instead of `*NOAUTO`. The use of `*AUTO` permits additional functions to be inlined that were not previously inlined. Because the `INLINE` option is set to `*OFF` by default, only objects created with the `INLINE` option set to `*ON` and with no inline mode specified are affected by this change.

14.4 Application Development ToolSet

There are no enhancements in IBM i 6.1 for the Application Development ToolSet (ADTS). It is recommended that you consider use of the new Eclipse-based workstation tool Rational Developer for i for development of traditional IBM i applications. To learn about Rational Developer for i and some of the benefits it has over ADTS, refer to 15.2, “Rational Developer for i” on page 364.

14.5 Debugger

IBM i Debugger provides a graphical user debugging environment on the IBM i platform. It is used for debugging and testing programs that run on IBM i, including programs written in ILE and OPM languages, Java, and programs that run in the PASE environment. It is packaged in a JAR file with IBM Toolbox for Java.

The IBM i Debugger changes in IBM i 6.1 are:

► Updated installation instructions:

The Software requirements for installing IBM i Debugger have been updated to reflect the supported versions of Java. You must have one of the following installed on your client:

- Java 2 Platform, either the Standard Edition (J2SE) or the Enterprise Edition (J2EE), version 1.4 or later
- Java 2 Runtime Environment (JRE), Standard Edition, version 1.4 or later

For more information about IBM i Debugger, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzamq/rzamqmain.htm&tocNode=int_96110

14.6 Portable Application Solutions Environment

The Portable Application Solutions Environment (PASE) is used to port IBM AIX applications to the IBM i platform with minimal effort. This support is installed in IBM i 6.1 as 5761-SS1, option 33.

PASE provides an integrated runtime environment that allows you to run selected applications without the complexity of managing operating systems, such as AIX or Linux. It also provides industry-standard and de facto-standard shells and utilities that provide you with a powerful scripting environment. The following new functions and changes are available:

14.6.1 Derivation of PASE

IBM i PASE for IBM i 6.1 is derived from AIX 5L™ 5.3, Technology Level 6.

14.6.2 Compilers

The following compiler products are now supported to run on PASE:

- IBM XL C/C++ Enterprise Edition for AIX, V9.0
- IBM XL C Enterprise Edition for AIX, V9.0
- IBM XL Fortran Enterprise Edition for AIX, V11.1

You can install the XL C/C++ Enterprise Edition for AIX, XL C Enterprise Edition for AIX, or XL Fortran Enterprise Edition for AIX product from an installation image. For more information about installing the AIX compilers from an installation image, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzalfr/rzalfinstcomp.htm>

14.6.3 Snapcore

A new utility, snapcore, is used for gathering information for a core file.

14.6.4 New or changed PASE runtime functions

New or changed PASE runtime functions in IBM i 6.1 include:

The ioctl function

The ioctl function is changed to include information for IPV6 interfaces in results from the SIOCGIFCONF command. IPV6 interfaces that also have an associated IPV4 interface on the same line description return the IPV4 interface address in dotted-name form in field ifr_name. IPV6 interfaces on lines that do not have an associated IPV4 interface return the line description name in field ifr_name.

The ioctl function is also enhanced to support the CSIOCGIFCONF, OSIOCGIFCONF, and SIOCGSIZIFCONF commands with results equivalent to the results of these ioctl commands on AIX.

Real-time semaphore interfaces

PASE is enhanced to support runtime functions, such as sem_close, sem_destroy, sem_getvalue, sem_init, sem_open, sem_post, sem_timedwait, sem_trywait, sem_unlink, and sem_wait (using the same system support as ILE functions that are named alike).

Execution priority

The mapping between the nice values and IBM i Run Priority (RUNPTY) value is changed. A single unit of nice value always maps to two units of RUNPTY (instead of a variable mapping to one to four units of RUNPTY). By default, within an interactive job, the PASE nice command runs a job at RUNPTY(40) instead of RUNPTY(60). Users who want the nice command to run work at batch RUNPTY(50) must specify a nice value of 15.

Execution priority values returned by PASE runtime functions (for example, field pi_pri in structures returned by getprocs) are now RUNPTY values. This causes the PASE ps command to show RUNPTY values under the heading PRI.

PASE applications can now control thread execution priority using the pthread_setschedparam API or an attributes object passed to the pthread_create API. The i5/OS operating system restricts thread priority, so applications can only set priority values that make a thread equal to or less prioritized than the RUNPTY for the job.

14.7 Qshell and utilities

IBM i supports a number of different program shells. These include the Qshell, Bourne, Korn, and C Shell. This support is installed, in IBM i 6.1, as 5761-SS1, option 30.

Qshell is a command environment based on POSIX and X/Open standards. It consists of the following two parts:

- ▶ The shell interpreter (or qsh) is a program that reads commands from an input source, interprets each command, and then runs the command using the services of the operating system.
- ▶ The utilities (or commands) are external programs that provide additional functions and can be quite simple or very complex.

14.7.1 Qshell

The following changes were made to Qshell support (5761-SS1, option 30):

ipcs - Report interprocess communication status

- ▶ Added support for large (64-bit) shared memory segments to the SEGSZ column. This is required to display the size of 64-bit shared memory segments correctly.
- ▶ Added support for shared memory page size to the -E option.

Qshell PATH variable changes

Before IBM i 6.1, if the PATH variable was not set, a default Qshell PATH of /usr/bin: was used. Starting in IBM i 6.1, the default Qshell PATH is:

```
/usr/bin:./QOpenSys/usr/bin.
```

Note: If the command you specify does not begin with a slash (/), qsh searches the directories listed in the order specified until it finds the command to run. The value of the variable is a colon separated list of directories. The current working directory is specified by a period (.) or a null directory before the first colon, between two colons, or after the last colon.

14.8 XML toolkit

The XML Toolkit for IBM System i, licensed program offering (LPO) 5733-XT2, V1R2M0, is designed to provide key development enablement components, XML parsers, to assist in the use of XML for general application or business-to-business (B2B) solutions. The XSL transformer allows you to reformat or reorganize XML documents to other formats (for example, XML or HTML).

14.8.1 New version of XML Toolkit

XML Toolkit for IBM System i (5733-XT2) is the replacement for XML Toolkit for iSeries (5733-XT1). XML Toolkit for iSeries (5733-XT1) is supported through IBM i 5.4. Only XML Toolkit for System i (5733-XT2) is supported in IBM i 6.1.

5733-XT1 has an integrated Web services server. 5722-XT2 does not because IBM i 6.1 has integrated Web services server support. For information about this integrated Web services server support, see 12.7, “New Web services server” on page 282.

5733-XT2 has several options that can be installed all at once or selectively. The options are:

- ▶ 5733XT2 *BASE 5050 XML Toolkit for IBM System i5
- ▶ 5733XT2 *BASE 2924 XML Toolkit for IBM System i5
- ▶ 5733XT2 1 5050 XML version 5.6 and XSL version 1.10
- ▶ 5733XT2 2 5050 XML 5.6 and XSL 1.10 Development
- ▶ 5733XT2 3 5050 XML Scripting for IBM System i5
- ▶ 5733XT2 4 5050 XML version 5.6.3 and XSL version 1.10.2
- ▶ 5733XT2 5 5050 XML 5.6.3 and XSL 1.10.2 Development
- ▶ 5733XT 6 5050 XML version 5.6.3 and XSL version 1.10.2

- ▶ 5733XT2 7 5050 XML 5.6.3 and XSL 1.10.2 Development

Options 4 and 5 of 5733-XT2 may be installed on systems running i5/OS V5R3 and later. Options 6 and 7 may only be installed on systems running IBM i 6.1 and later. Other than the differences in target releases, the functionality of these two sets of options is the same.

5733-XT2 contains the latest XML and XSLT code from Apache and IBM for XML4C 5.6.3 and XSLT4C 1.10.2. There are also two new service programs for clients who want to improve the XML and XSLT code performance.

14.9 Control language

Control language (CL) allows system programmers and system administrators to write programs using IBM i and related products' commands.

14.9.1 New commands

The following new commands are added:

The Close Database File (CLOSE) command

The Close Database File (CLOSE) command can be used to close a database file that is implicitly opened by the Receive File (RCVF) command.

The CLOSE command can be used to process the records in the database file more than once. When the file is closed, the next RCVF command implicitly reopens the file and reads a record. It can also be used for a database file that is already closed or was never opened; no error message is sent. When the CL program or procedure ends, all database and display files that are implicitly opened in the CL program or procedure were implicitly closed. Example 14-4 uses the CLOSE command with a database file.

Example 14-4 Close a database file specifying an Open identifier

```
DCLF FILE(*LIBL/MYFILE2) OPNID(FILE2)
:
RCVF OPNID(FILE2)
:
CLOSE OPNID(FILE2)
RCVF OPNID(FILE2)
```

The Include CL Source (INCLUDE) command

The Include CL Source (INCLUDE) command can be used to embed another source file when you compile a CL source program using either the CL compiler or the ILE CL compiler.

The CL source to be embedded can be located in another member of the same source file that is identified on the Source file (SRCFILE) parameter of the CL compiler commands or a different source file. The CL compiler commands include Create CL Program (CRTCLPGM), Create CL Module (CRTCLMOD), and Create Bound CL Program (CRTBNDCL). The CL source to be embedded cannot contain a Program (PGM), an End Program (ENDPGM), or an INCLUDE command.

If you use the Create CL Program (CRTCLPGM) command to run the CL compiler and the CL source contains one or more INCLUDE commands, you can run the Retrieve CL Source (RTVCLSRC) command at a later time to retrieve either the original CL source (which

contains just the INCLUDE commands) or the expanded CL source (which contains the embedded CL source commands).

Example 14-5 shows the use of the INCLUDE command. The command in this example will cause the CL program source in member DCLSET1 to be included at compile time.

Example 14-5 Include CL source using default INCLUDE file

```
INCLUDE SRCMBR(DCLSET1) SRCFILE(*INCFILE)
```

Note: Member DCLSET1 must be located in the file specified for the INCFILE parameter on the CL command that invoked the CL compiler.

14.9.2 Command enhancements

The following commands are changed:

Declare Processing Options (DCLPRCOPT)

With the Declare Processing Options (DCLPRCOPT) command, you can set many of the same create option parameters that are on the CL compiler commands. The CL compiler commands include Create CL Program (CRTCLPGM), Create CL Module (CRTCLMOD), and Create Bound CL Program (CRTBNDCL).

For more information about the DCLPRCOPT command, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rbam6/rbam6dclprcopt.htm&tocNode=int_102349

Command Definition (CMD)

Several create option parameters on the Create Command (CRTCMD) command can also be specified on the CMD statement. Values that are specified on the CMD statement take precedence over any values specified on the CRTCMD command.

Example 14-6 sets the prompt text for the command from message UCD0002 in message file MYCMDPMT, which is located using the library list.

All prompt text messages defined for the command are retrieved dynamically when the command is prompted. Error messages sent from the Dependency (DEP) statements in the command definition are found in message file MYCMDMSG, which is located using the library list. The text for the command object is the same as the command prompt text. Only the first two parameters of the command can be specified in positional form without the associated parameter keyword.

Example 14-6 Defining prompt text and Command Creation options

```
CMD PROMPT(UCD0002) PMTFILE(MYCMDPMT *DYNAMIC)+  
MSGF(MYCMDMSG) TEXT(*CMDPMT) MAXPOS(2) +  
PRDLIB(MYAPPLIB) HLPID(*CMD) +  
HLPPNLGRP(MYAPPHLP)
```

Library MYAPPLIB is automatically added to the library list before the command is run and is removed from the library list when the command is completed. The help identifier for the command help modules in panel group MYAPPHLP starts with the name of the command object that is created by the CRTCMD command.

Dynamic prompt message

Starting from IBM i 6.1, CL commands for the operating system and most IBM products use dynamic retrieval of prompt text messages to allow a single copy of the command to handle all installed national language versions. With the Unicode support in control language (CL), the command processing program (CPP) can always get its data in either extended binary-coded decimal interchange code (EBCDIC) or UTF-16, regardless of how the data is passed to CL. For more information, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rbam6/dynmsg.htm?tocNode=int_102428

Unicode support

Unicode support was added to the CL commands. This support is only for data type *CHAR or *PNAME. With the Unicode support in control language (CL), the command processing program (CPP) can always get its data in either extended binary-coded decimal interchange code (EBCDIC) or UTF-16, regardless of how the data is passed to CL. For more information about unicode support in CL, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rbam6/rbam6unicode.htm?tocNode=int_97855

New messages sent to the QSYSMSG message queue

The system can send several new messages to the QSYSMSG message queue. Messages that are sent to the QSYSMSG message queue can be categorized into these groups: messages sent to the QSYSMSG message queue instead of QSYSOPR; messages sent to QSYSMSG, QSYSOPR, or both; and messages always sent to both QSYSMSG and QSYSOPR. For the complete list of messages, refer to the IBM i Information Center at:

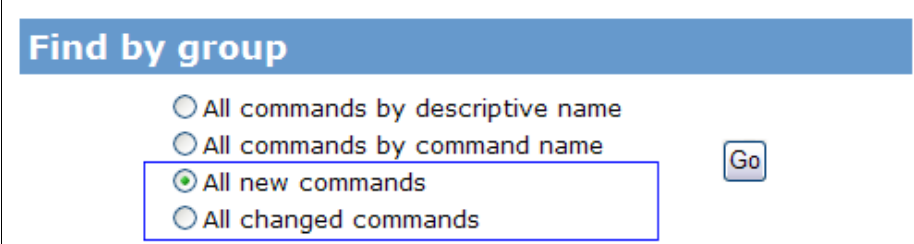
http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rbam6/msmq.htm?tocNode=int_102476

Summary of new and changed commands

You can display a list of all new commands or all changed commands using the CL command finder at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/clfinder/finder.htm>

Figure 14-4 on page 344 shows how to display the new and changed commands using the CL command finder.



The screenshot shows a web interface titled "Find by group". It contains four radio button options: "All commands by descriptive name", "All commands by command name", "All new commands", and "All changed commands". The "All new commands" option is selected, indicated by a green dot. To the right of the options is a "Go" button. A blue rectangular box highlights the "All new commands" and "All changed commands" options.

Figure 14-4 CL command finder

14.10 Structure Query Language

Structured Query Language (SQL) is a standardized language for defining and manipulating data in a relational database. In accordance with the relational model of data, the database is perceived as a set of tables, relationships are represented by values in tables, and data is retrieved by specifying a result table that can be derived from one or more base tables.

14.10.1 Call Level Interface

IBM i 6.1 includes new and significantly changed information for the SQL Call Level Interface (CLI):

- ▶ DB2 CLI's internal workspace limits have been increased in IBM i 6.1. This allows applications to run with a greater number of active handles than previous releases.
- ▶ Support for extended indicators was added through a new connection attribute. See the attribute `SQL_ATTR_EXTENDED_INDICATORS` under the API `SQLSetConnectAttr()` for more information.
- ▶ Support for new connection attributes that specify client information special register values, has been provided. See the new attributes under API `SQLSetConnectAttr()`.
- ▶ CLI applications can now explicitly free locators by using the new connection attribute, `SQL_ATTR_FREE_LOCATORS`, provided by the API `SQLSetConnectAttr()`.
- ▶ The decimal floating point data type is now supported in DB2 CLI. Support for changing a decimal floating point's rounding mode is also part of IBM i 6.1.
- ▶ DB2 CLI now allows cursor names of up to 128 bytes, to be specified on the `SQLSetCursorName()` API.
- ▶ DB2 CLI added full Unicode support in IBM i 6.1. This is provided through the Wide APIs. More information about this support can be found in the IBM Information Center article *Unicode in DB2 CLI* at:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/cli/rzadcliunicode.htm>
- ▶ New statement attributes provide additional information for setting and retrieving blocking information. These attributes can be found under the `SQLSetStmtAttr()` API.

SQL Call Level Interface API changes

The most notable API change was made to the constant values that represent `BINARY` and `VARBINARY` data types. Table 14-4 shows the changes made to the constant values.

Table 14-4 Changes to the value of constants

Constant name	Old value	New value
<code>SQL_BINARY</code>	96	-2
<code>SQL_VARBINARY</code>	97	-3

Important: This change requires that applications that use CLI, which reference these constants, to recompile their programs that contain the reference before running the program on the IBM i 6.1 release. In addition, any application that has a hard-coded value of 96 to represent `SQL_BINARY` needs to change to use the new value. Any application that has a hard-coded value of 97 to represent `SQL_VARBINARY` needs to change to use the new value.

Applications that are not recompiled before running on IBM i 6.1 will see SQL0804 if an incorrect data type is used in the application.

Other API changes for IBM i 6.1 include:

SQLBindCol	Bind a column to an application variable.
SQLBindParam	Bind a buffer to a parameter marker.
SQLBindParameter	Bind a parameter marker to a buffer.
SQLColumns	Get column information for a table.
SQLForeignKeys	Get the list of foreign key columns.
SQLGetCol	Retrieve one column of a row of the result set.
SQLGetTypeInfo	Get data type information.
SQLPrimaryKeys	Get primary key columns of a table.
SQLProcedureColumns	Get input/output parameter information for a procedure.
SQLProcedures	Get list of procedure names.
SQLSetConnectAttr	Set a connection attribute.
SQLSetCursorName	Set cursor name.
SQLSpecialColumns	Get special (row identifier) columns.
SQLStatistics	Get index and statistics information for a base table.
SQLTables	Get table information.

For more information about the API changes, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/cli/rzadpwhatsnew.htm&tocNode=int_96860

Compatibility changes

Changes were made to the internal implementation of the CLI metadata functions. These changes were made for compatibility with other DB2 CLI implementations and affect some of the result set information returned for these functions. The list of affected APIs, their affected result columns, and the behavior change can be found at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/cli/rzadpwhatsnew.htm&tocNode=int_96860

To revert to using the old metadata implementation, set the SQLSetConnectAttr() API's connection attribute SQL_ATTR_OLD_MTADTA_BEHAVIOR to SQL_TRUE.

14.10.2 SQL programming

In IBM i 6.1 the following new functions and changes were added to SQL programming on IBM.

Close SQL cursor SET OPTION changes

In IBM i 6.1, the Close SQL Cursor (CLOSQLCSR) SET OPTION default has changed from *ENDACTGRP to *ENDMOD for SQL routine bodies. When CLOSQLCSR is not specified, SQL procedures, scalar functions, and triggers are created with CLOSQLCSR=*ENDMOD. In previous releases, CLOSQLCSR=*ENDACTGRP was used. For SQL UDTFs, the CLOSQLCSR option is ignored and the UDTF is always created with *ENDACTGRP.

Note: Because of this change, the default behavior is different in the following case: Statements prepared in an SQL routine are closed when the routine exits. A statement cannot be prepared on one call of an SQL code body built with CLOSQLCSR=*ENDMOD and run on another call.

Row change timestamp column

When you create a table, you can now define a column to be a row change timestamp column. Every time a row is added or changed in the table, the row change timestamp column value is set to the timestamp of the insert or update operation.

Row change expressions

The new expressions ROW CHANGE TIMESTAMP and ROW CHANGE TOKEN can be used to determine when a row was last changed.

Full outer join

DB2 for IBM i now supports full outer joins. Previously, you could only simulate a full outer join by using a left outer join and a right exception join.

SELECT from INSERT

When you insert one or more rows into a table, you can now select the result rows of the insert operation by specifying the INSERT statement in the FROM clause of a SELECT statement.

Decimal floating point support

DB2 for IBM i SQL now supports the decimal floating-point data type:

- ▶ Allowable conversions of data types
- ▶ Data retrieval errors
- ▶ Parameter passing conventions for stored procedures and user-defined functions

Source stream file support

Either a source file member or a source stream file can now contain the source for a DB2 for IBM i application program. The SQL statement processor now allows statements to be run from a source stream file as well.

Other functional changes to the SQL programming information

These new functional changes in IBM i 6.1 include:

- ▶ The default value, the hidden attribute, and the row change timestamp attribute can also be included in the LIKE or AS clause of the CREATE TABLE statement.
- ▶ Most expressions allowed by SQL can now be used in the definition of the key columns for indexes.
- ▶ New special registers have been added to the Special registers in SQL statements topic.
- ▶ The Concurrency topic has been updated to include considerations for using the SKIP LOCKED DATA clause.
- ▶ The Indicator variables and stored procedures topic has been updated to document support for extended indicators.

Further information

For more information about the enhancements for SQL programming in IBM i 6.1, refer to the IBM i Information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/sq1p/rbafywhatsnew.htm?tocNode=int_96819

14.10.3 Embedded SQL programming

New and significant changes were added to Embedded SQL programming in IBM i 6.1:

- ▶ The ILE RPG precompiler now recognizes scoped variables. They are read using host variables in ILE RPG applications that use SQL.
- ▶ The ILE COBOL precompiler now supports the UCS-2 variable type:
 - Graphic host variables in COBOL applications that use SQL
 - Host structure in COBOL applications that use SQL
 - Host structure array in COBOL applications that use SQL
- ▶ The use of stream files is now supported on the Create SQL ILE COBOL Object (CRTSQLCBLI), Create SQL ILE C Object (CRTSQLCI), Create SQL ILE C++ Object (CRTSQLCPPI), and Create SQL ILE RPG Object (CRTSQLRPGI) commands.
- ▶ Extended indicators are supported so that an application that uses indicator variables can pass default or unassigned values. Read Indicator variables are used to assign special values.
- ▶ The decimal floating-point data type is now supported by the C precompiler. See the following topics for more information:
 - Numeric host variables in C and C++ applications that use SQL
 - Host structure declarations in C and C++ applications that use SQL
 - Host structure arrays in C and C++ applications that use SQL

For more information about the above enhancements, refer to IBM i information Center at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajp/rzajpwhatsnew.htm&tocNode=int_96806

14.10.4 SQL Messages and codes

In IBM i 6.1 new messages and codes were added. Several additions were made to the Listing of SQLSTATE values and the Listing of SQL messages.

SQL routines that contain references to SQLCODE and SQLSTATE in an IF-THEN-ELSE statement will not work the same. SQLCODE and SQLSTATE will be set to 0 and the diagnostic area cleared before the first statement is executed in the THEN or ELSE clause. See Example 14-7.

Example 14-7 SQL routine with reference to SQLCODE

```
IF SQLCODE < 0 THEN
    SET SAVE_CODE = SQLCODE;
END IF;
```

If SQLCODE is negative when the IF is executed, SAVE_CODE is set to 0. In previous releases, SAVE_CODE was set to the negative value contained in SQLCODE.

Note: These changes apply only to functions and procedures that are built using the LANGUAGE SQL keyword on the CREATE PROCEDURE, CREATE FUNCTION, CREATE TRIGGER, ALTER PROCEDURE, or ALTER FUNCTION SQL statements. They do not apply to higher-level language programs.

These changes can be found in IBM i Information Center at the following links:

- ▶ Listing of SQLSTATE values

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzala/rzalaccl.htm>

- ▶ Listing of SQL messages:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzala/rzalaml.htm>

14.10.5 Grouping and On-Line Analytical Processing

Many BI applications and On-Line Analytical Processing (OLAP) tools involve hierarchical, multidimensional aggregate views of transaction data.

Grouping is used in conjunction with grouping sets and super groups. The GROUPING aggregate function returns a value that indicates whether a row returned in a GROUP BY answer set is a row generated by a grouping set that excludes the column represented by expression.

IBM i 6.1 grouping and OLAP capabilities allow data to be grouped in multiple ways with a single SQL request:

- ▶ GROUPING SETS
- ▶ ROLLUP
- ▶ CUBE

GROUPING SETS

GROUPING SET on GROUP BY clause enables DB2 to return aggregates for multiple sets of grouping columns.

GROUPING SETS((Country, Region), (Country, Store)) will result in the data being summarized at the following levels

- ▶ (Country, Region)
- ▶ (Country, Store)

CUBE and ROLLUP can be used in combination with Grouping Sets.

Important: These types of combinations can result in an exponential growth in the number of grouping sets returned by a query; combine carefully.

ROLLUP

ROLLUP on GROUP BY clause results in DB2 returning aggregates for each level of the hierarchy implicitly represented in the grouping columns.

ROLLUP(Country, Region) will result in the data being summarized at the following levels:

- ▶ (Country, Region)
- ▶ (Country)
- ▶ () << represents Grand Total

CUBE

CUBE on GROUP BY clause results in DB2 returning aggregates for all possible distinct combinations represented by the grouping columns. It returns results at multiple intersection points.

CUBE(Country, Region) will result in the data being summarized at the following levels:

- ▶ (Country, Region)
- ▶ (Country)
- ▶ (Region)
- ▶ () << represents Grand Total

14.10.6 SQL reference

The following link on IBM i Information Center highlights the major new features in the SQL reference topic:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/db2/rbafzwhatsnew.htm&tocNode=int_96909

This section includes details of some of the major new features and enhancements.

VALUES Clause in FROM

Using the VALUES clause in a SELECT creates a result table by specifying the actual values, using expressions, for each column in a row in the result table. See Example 14-8.

Example 14-8 Values clause in FROM

```
SELECT
:
FROM
IBM i 6.1db2.conference,
(values ('i5/OS', 'Technical'))AS X (R1, R2)
```

Support for Full Outer Join

Full Outer Join combines the results of both left and right outer joins. The result contains all rows from both tables, showing NULL values for missing matches. This support is equivalent to UNION ALL of LEFT OUTER JOIN and RIGHT EXCEPTION JOIN. See Example 14-9.

Example 14-9 Support for full outer join

```
SELECT      PART,
            SUPPLIER,
            PARTS.PROD#,
            PRODUCT
FROM
            PARTS FULL OUTER JOINPRODUCTS ON PARTS.PROD# = PRODUCTS.PROD#
```

INSERT in FROM

Allows a programmer to dynamically generate and populate temporary tables as part of the query definition. This support provides an easy way to return generated values.

You can retrieve the values for rows that are being inserted by specifying the INSERT statement in the FROM clause of a SELECT statement.

SKIP LOCKED DATA

The skip-locked-data-clause specifies to skip rows on which incompatible locks are held by other transactions. These rows can belong to any table referenced in the statement. By default, all DB2 requests wait for conflicting locks to be released. It gets the next row in the transaction table for update and skip any rows that are currently locked, as shown in Example 14-10.

Example 14-10 Skip locked data

```
DECLARE X CURSOR FOR
      SELECT * FROM
      FROM transaction
      WITH CS
      SKIP LOCKED DATA
      OPEN X;
FETCH X INTO ...
```

Note: This clause is only used under isolation levels CS and RS. It is ignored when an isolation level of NC, UR, or RR is in effect. It applies to row level locks. If an isolation clause is not specified, the default isolation level is used. This clause can also be used on the DELETE and UPDATE statements, specifying to skip rows on which incompatible locks are held by other transactions.

EXTENDED INDICATOR Variables

To indicate that the associated host variable contains the null value, the indicator variable is set to a negative value of -1, -2, -3, -4, or -6. A CALL statement with indicator variables is processed as follows:

- ▶ If the indicator variable is negative, a default value is passed for the associated host variable on the CALL statement and the indicator variable is passed unchanged.
- ▶ If the indicator variable is not negative, the host variable and the indicator variable are passed unchanged.

Table 14-5 lists the values and meaning of each indicator.

Table 14-5 EXTENDED INDICATOR variables

Indicator Values	Meaning
0	Value is provided
-1, -2, -3, -4, and -6	Null value
-5	Default value
-7	Column is treated as if it had not been specified in the statement

Prior to this support you had to update all columns or use a different UPDATE statement for every combination of columns you wanted to update. It applies to UPDATE and INSERT VALUES, also allowed on PREPARE.

Scalar functions

A scalar function takes input arguments and returns a single value result. A scalar function can be used wherever an expression can be used.

Datetime format

The following enhancements have been made to the Datetime format:

- ▶ The `ADD_MONTHS` and `LAST_DAY` have been changed in IBM i 6.1 to return a timestamp when the input argument is a timestamp; otherwise a date is returned. In IBM i 5.4, a date was always returned.
- ▶ The `NEXT_DAY` function has been changed in IBM i 6.1 to return a date when the input argument is date; otherwise a timestamp is returned. In IBM i 5.4, a timestamp was always returned.
- ▶ `TIMESTAMP_FORMAT` returns a timestamp from a character string representation of a timestamp according to the specified format of the string
- ▶ Enhancements to the `VARCHAR_FORMAT`, which returns a character string representation of a timestamp, with the string in a specified format.
- ▶ `MONTHS_BETWEEN` returns an estimate of the number of months between two expressions.
- ▶ `ROUND_TIMESTAMP` returns a timestamp that is the expression, rounded to the unit, specified by the format-string.
- ▶ `TRUNC_TIMESTAMP` returns a timestamp that is the expression truncated to the unit specified by the format-string.

String handling

The following enhancements have been made to String handling:

- ▶ `ASCII` returns the ASCII code value of the leftmost character of the argument as an integer.
- ▶ `CHR` returns the character that has the ASCII code value specified by the argument. If the expression is 0, the result is the blank character (X'20').
- ▶ `ENCRYPT_AES` returns a value that is the result of encrypting a data-string using the AES encryption algorithm.

14.10.7 SQL troubleshooting

In IBM i 6.1 the following new question has been added to the IBM i database frequently asked questions topic:

Can an SQL trigger access the trigger buffer?

The body of an SQL trigger can use the `REFERENCING NEW ROW` and `REFERENCING OLD ROW` clauses to reference the fields in the before and after record images that are stored in the trigger buffer. However, an SQL trigger program cannot reference the trigger buffer directly, nor can it reference the other information in the buffer. For example, the trigger time and trigger event are not available to an SQL trigger.

For the list of IBM i database frequently asked questions refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzate/rzatefaq.htm>

14.10.8 SQL Query Engine enhancements

DB2 for IBM i provides two query engines to process queries: Classic Query Engine (CQE) and SQL Query Engine (SQE).

Within SQE, several more components were created and other existing components were updated. Additionally, new data access methods are possible with SQE that are not supported under CQE

IBM i 6.1 enhancements for SQL Query Engine include:

Restriction elimination

The following restrictions have been eliminated:

National Language Sort Sequences

The NLSS support in IBM i is frequently used in overseas markets where the application requires the character data be sorted in a manner that matches the local language instead of the default *HEX ordering, which closely mirrors the sorting of the English alphabet.

Translation support

The low-level translation capability is required by commonly used SQL built-in functions such as Upper and Lower.

Remaining restrictions

The following restrictions remain:

ICU 2.6.1 Sort sequences

In IBM i 6.1, the SQE has added support for sort sequence. However, its ICU sort sequence support supports only ICU version 3.4 and not version 2.6.1.

Non-SQL interfaces (OPNQRYF, Query/400, QQQRY API)

These interfaces continue to be processed by the Classic Query Engine (CQE).

Logical file reference on a FROM clause

Logical file references on a FROM clause will continue to be processed by the Classic Query Engine (CQE).

Select/omit logical files defined on tables

QAQQINI option changes in IBM i 6.1: the default value for the QAQQINI option for IGNORE_DERIVED_INDEX has changed from *NO to *YES. The default behavior now runs supported queries through the SQE even if a select/omit logical file index is created over any of the tables in the query. In IBM i 6.1, the SQE optimizer supports many types of derived indexes.

The use of the QAQQINI option for IGNORE_DERIVED_INDEX applies to only select/omit logical file indexes.

Technological advances

In IBM i 6.1, the SQL query engine has been enhanced by the following technological advances:

Self-learning query optimization

A self-learning query optimizer is able to analyze poor-performing query plans and dynamically adjust its internal algorithms based on feedback to select a better query plan on future executions.

Self-adjusting query execution

Optimizer will build an index over a hash table to minimize page fault issues associated with a poorly sized hash table. Stats manager will update stale stats in small tables in real time.

14.11 Application programming interfaces

IBM i application programming interfaces (APIs) allow your application program written in a high-level language to use specific data or functions of the IBM i operating system.

14.11.1 API changes

In IBM i 6.1 Application Programming Interfaces has new or significantly changed information. A list of new APIs or changed APIs for IBM i 6.1 can be found using the API finder, as shown in Figure 14-5. API finder is included in IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/apifinder/finder.htm>

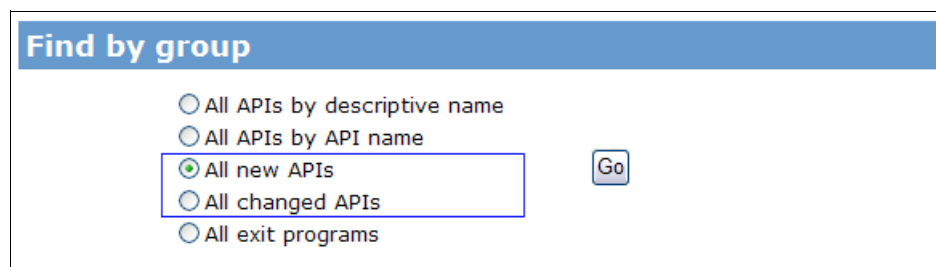


Figure 14-5 New and changed API finder

In addition to the list of new and changed APIs, the following categories of APIs have been enhanced in IBM i 6.1:

14.11.2 Cluster

A new subcategory named Cluster Administrative Domain has been added. The cluster administrative domain APIs enable the management of resources that need to be maintained consistently across the nodes in a clustered environment. A cluster node can only belong to a single cluster administrative domain. Within this domain, a user can specify objects or attributes of the operational environment, which will be maintained synchronously across the nodes in the domain.

The Add Monitored Resource Entry (QfpadAddMonitoredResourceEntry), Remove Monitored Resource Entry (QfpadRmvMonitoredResourceEntry), and Retrieve Monitored Resource

Information (QfpadRtvMonitoredResourceInfo) APIs have been moved to this subcategory from the Server Support category.

For more information about Cluster Administrative Domain, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/apis/clust6a1T0C.htm>

14.11.3 Problem Management

A new subcategory named QTRC Trace has been added. The Problem Determination APIs have been moved to the Problem Management category from the UNIX-Type category.

The QTRC Trace APIs provide support for writing application trace data. Using these APIs, you can do the following tasks:

- ▶ Retrieve the current active trace level for a component.
- ▶ Write component trace data as a string of text, hexadecimal dump, or call stack.

The QTRC Trace APIs include:

- ▶ Get Active Trace Level (QtrcGetActiveLevel())
- ▶ Write Call Stack Trace Data (QtrcWriteStack())
- ▶ Write Hexadecimal Dump Formatted Trace Data (QtrcWriteHexDumpF())
- ▶ Write Hexadecimal Dump Trace Data (QtrcWriteHexDump())
- ▶ Write Text Trace Data (QtrcWriteText())
- ▶ Write Text Trace Data Using Pointer-based Print Formatted (QtrcWriteTextPPrintF())
- ▶ Write Text Trace Data Using Print Formatted (QtrcWriteTextPrintF())
- ▶ Write Text Trace Data Using Variable Argument List Print Formatted (QtrcWriteTextVPrintF())

For more information about QTRC Trace APIs, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/apis/qtrctrace.htm>

14.11.4 Security

The NetWare Authentication Entry APIs have been removed from the Security – Network Security category.

14.11.5 Data conversion APIs

In the Information Center, the Data Conversion APIs have been moved from the National Language Support category to the UNIX-Type category. These APIs have also been renamed Character Conversion APIs.

Character Conversion APIs are:

- ▶ Code Conversion (iconv())
- ▶ Code Conversion Allocation (QtqIconvOpen())
- ▶ Code Conversion Allocation (iconv_open())

- ▶ Code Conversion Deallocation (iconv_close())
- ▶ Convert Data (QDCXLATE)

14.11.6 Miscellaneous

The miscellaneous APIs are the APIs that do not logically fall in a specific part of the IBM i reference information. In the Information Center, the Miscellaneous category no longer includes a link to the User Application APIs. A new subcategory named Service-related APIs has been added to the Miscellaneous category.

Note: The service-related APIs must be used only when recommended by an IBM service representative for collecting problem-related information.

For more information about service-related APIs, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/apis/misc1d.htm>

14.11.7 API documentation updates in the Information Center

Miscellaneous technical updates have been made to the following API documentation:

- ▶ Cluster - Cluster Control:
 - Change Cluster Node Entry (QcstChangeClusterNodeEntry) API
 - End Cluster Node (QcstEndClusterNode) API
 - Remove Cluster Node Entry (QcstRemoveClusterNodeEntry) API
- ▶ Cluster - Cluster Resource Group:
 - Change Cluster Resource Group (QcstChangeClusterResourceGroup) API
 - Initiate Switchover (QcstInitiateSwitchOver) API
 - Remove Node From Recovery Domain (QcstRemoveNodeFromRcvyDomain) API
 - Start Cluster Resource Group (QcstStartClusterResourceGroup) API
- ▶ Configuration:
 - Open List of ASPs (QYASPOL) API
 - Start Disk Management Operation (QYASSDMO) API
- ▶ Database and File:
 - Cancel EDRS Request (QxdaCancelEDRS) API
- ▶ UNIX-Type - Integrated File System:
 - stat()--Get File Information
- ▶ Other API documentation updates to the following APIs:
 - Convert Data (QDCXLATE) API
 - Create and Send MIME E-mail (QtmsCreateSendEmail) API
 - Log Software Error (QPDLOGGER) API
 - Open Virtual Terminal Path (QTVOPNVT) API
 - Report Software Error (QpdReportSoftwareError) API
 - Retrieve Job Information (QUSRJOBI) API

- Set User Policy (QPDETPOL) API
- Start Disk Management Operation (QYASSDMO) API

14.11.8 Machine Interface instructions

This section gives an overview of some of the changes to the Machine Interface (MI) architecture for IBM i 6.1.

- ▶ Decimal floating-point has been added to the CPYNV, RETCA, SETCA, SETDPAT, and SSCA instructions.
- ▶ The order of unlocks can be specified on the LOCK and UNLOCK instructions.
- ▶ Teraspace locks will be handled differently in some cases:
 - Teraspace location locks can only be obtained as thread-scoped locks. Process-scoped locks on teraspace locations will no longer be allowed. See the LOCKSL and LOCKTSL instructions for additional information.
 - Materialize instructions will interpret teraspace addresses from the context of the thread in which they are running. See the MATAOL, MATOBJLK, and MATSELLK instructions for additional information.
 - If the current thread does not belong to the process whose locks are being materialized, a null pointer value will be returned by the MATPRLK instruction for any teraspace locations.
 - Mapping a teraspace location to a single-level storage address is no longer allowed. Starting in IBM i 6.1, the LOCKSL and UNLOCKSL instructions are the preferred instructions for locking and unlocking space locations, including teraspace locations. LOCKTSL and UNLCKTSL will not be updated to support any future enhancements.
- ▶ The DESMTX instruction will allow mutexes to be destroyed by means of a mutex copy.
- ▶ The system sequence number can be materialized using the MATJSAT instruction.
- ▶ The maximum number of entries in a user profile has increased and can be materialized using the MATUP instruction.
- ▶ The timestamp of the last automatic reclaim can be materialized using the MATQAT instruction.
- ▶ New fields have been added to the MATSOBJ instruction to indicate whether the object is in the format required by the system.
- ▶ The process CCSID, country identifier, and language identifier can be materialized using the MATPRATR instruction.
- ▶ Scaled processor time can be materialized using the MATMIF and MATPRATR instructions.
- ▶ Coordinated Universal Time (UTC) is supported on the system. The time and date at which the system should automatically power on has been changed from local system time to Coordinated Universal Time in MATMATR instruction option hex 013C. See the MATTODAT, MATMDATA, and MATMATR instructions for additional information.
- ▶ The system can defer activation of service programs until one of their functions is called. The activation information can be materialized using the MATBPGM and MATPTR instructions.
- ▶ New options have been added to the MATMDATA instruction to materialize the following information:
 - The machine default page size

- The teraspace segment size
- The time-of-day clock as local time with a value that is not unique
- The time-of-day clock as Coordinated Universal Time (UTC)
- ▶ An independent index can be created with a maximum entry length of 32,000 bytes using the CRTINX instruction and the maximum entry length can be materialized using the MATINXAT instruction.
- ▶ Operations to an independent index can be prevented from user state. See the CRTINX, DESINX, FNDINXEN, and RMVINXEN instructions for additional information.
- ▶ The system supports a longer serial number for some hardware. An *extended serial number* field has been added to options 12 and 20 of the MATRMD instruction. It is recommended that current users of MATRMD options 12 and 20 use the new extended serial number field.
- ▶ Additional journal attributes can be materialized using the MATJOAT instruction.
- ▶ Remote journal can use dataport services. Dataport services information can be materialized using the MATJPAT instruction.
- ▶ Additional entries can be stored in an authorization list and an authorization list can be materialized into an independent index. See the MATAL, MATAUU and MATAUOBJ instructions for additional information.
- ▶ New fields for *bi-di formatting* and *unicode case level version* have been added to the XLATEMB instruction.
- ▶ Additional information can be materialized using the following MATMATR instruction options:
 - Hex 01E0 Partitioning information
 - Hex 01F4 On-demand processor information
 - Hex 0210 Shared processor pools information (new option)
 - Hex 0220 Processor attributes (new option)
- ▶ Additional information can be materialized using the following MATRMD instruction options:
 - Hex 12 Auxiliary storage information
 - Hex 20 Auxiliary storage information including varied-off Independent ASPs
 - Hex 26 Processor utilization data
 - Hex 27 Shared processor pools utilization information
 - Hex 28 Multiprocessor utilizations
 - Hex 2F Resource affinity selection override (new option)
 - Hex 30 Resource information by name

14.11.9 Cryptographic services

The cryptographic services APIs perform cryptographic functions in the IBM i operating system or on the 2058 Cryptographic Accelerator, as specified by a user.

The cryptographic services APIs can help you ensure the following:

- ▶ Privacy of data
- ▶ Integrity of data

- ▶ Authentication of communicating parties
- ▶ Non repudiation of messages

IBM dropped support of IBM Cryptographic Support for AS/400® (5722-CR1) in IBM i 6.1. For information on how to migrate CR1 applications to the Cryptographic Services APIs, refer to the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/apis/qc3MigrateCR1.htm>

14.12 Globalization

IBM i support for the International Components for Unicode (ICU) libraries has been updated to be based on ICU version 3.4. This includes the support for the ICU-based sort support.

The ICU 3.4 tables follow the naming of I34xx_yy. You can see these choices by looking at the table objects on the operating system. Use the following command to see the ICU 3.4 tables:

```
WRKTBL TBL(QSYS/I34*)
```

Note: If your applications use ICU sort support, you need to update them to use the 3.4 version of the support for better performance.

The system bidirectional layout services have been updated to an enhanced version. This update might cause different conversion results if bidirectional data is being processed.

14.13 Multithreaded applications

On IBM i, a job represents a process. Each process has at least one thread (task) in which the program runs its code. The first thread in a process is referred to as the initial thread. Some processes are capable of supporting additional threads, which are called secondary threads.

You can now use thread local storage (TLS), a new method to create data that is private to a thread. Thread local storage is similar to thread-specific storage, except that thread local storage is implemented by the operating system, whereas thread-specific storage is implemented with APIs. As a result, thread local storage is faster and easier to use than thread-specific storage. ILE C, ILE C++, and ILE RPG compilers support thread local storage.

Threads cannot share certain resources. Data that threads cannot share among themselves are called thread-private data. These threads can be defined using either thread-specific data or thread local storage.

Data that is defined using thread local storage is similar to thread-specific data. Threads do not share thread local storage, but all functions within a thread can access the storage for that thread. Thread local storage is different from thread-specific data in that thread local storage is managed by the language environment, whereas thread-specific data is managed by the application.

In IBM i 6.1a new storage variable is added, thread local storage.

A thread local storage variable is visible to all source files in your application, a single source module, or a single function, depending on how the variable is declared. Thread local storage

is not visible to other threads in the process. When a thread is started, copies of thread local storage variables are created for the thread.

For a code example implementing thread local storage, refer to IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahw/rzahwex1.htm>



Rational products for IBM i

IBM Rational Software Development products for IBM i offer tools to help customers accelerate software innovation and delivery by allowing extension and reuse of existing applications to create Web applications and Web services without costly rewrites. These tools facilitate collaboration between dispersed software teams, allowing you to leverage existing IT skills to create multi-platform SOA and Web 2.0.

The latest Rational products for IBM i are:

- ▶ Rational Developer for i for SOA Construction v7.5
 - Rational Developer for i v7.5
 - Rational Business Developer v7.5 and Enterprise Generation Language (EGL)
- ▶ Rational Host Access Transformation Services for 5250 v7.5
- ▶ Rational Host Integration Solution v7.5
- ▶ Rational Team Concert™ for i

Version 7.5 is based on the functionality of Eclipse 3.4 and also supports the Rational Jazz™ team collaboration platform with Rational Team Concert for i.

This chapter discusses the new functions and enhancements delivered in version 7.5 of the products and provides examples of some of these features.

15.1 Rational Developer for i for SOA Construction

Rational Developer for IBM i for SOA Construction (RDi SOA) is a new software bundle that enables IBM i users to develop Web applications. It comprises:

- ▶ Rational Developer for i
- ▶ Rational Business Developer
- ▶ Rational Host Access Transformation Services (HATS) Toolkit

Note: HATS Toolkit includes IBM WebFacing Tool for IBM i.

Rational Developer for i for SOA Construction enables you to build innovative rich Web applications and user interfaces (UI), so you can link to the back-end RPG, COBOL, and Enterprise Generation Language (EGL) business applications, as shown in Figure 15-1.

RDi SOA V7.5 provides the following features:

- ▶ Simplifies and accelerates development of Web and SOA applications.
- ▶ Exploits advances in middleware and technology while writing business applications.
- ▶ Supports edit/compile/debug for IBM i RPG and COBOL application development.
- ▶ Can be used to modernize and extend existing RPG and COBOL business applications to the Web and mobile devices, or into Web services.
- ▶ Includes tools for development of EGL applications.
- ▶ HATS toolkit can be used to modernize 5250 user interfaces.

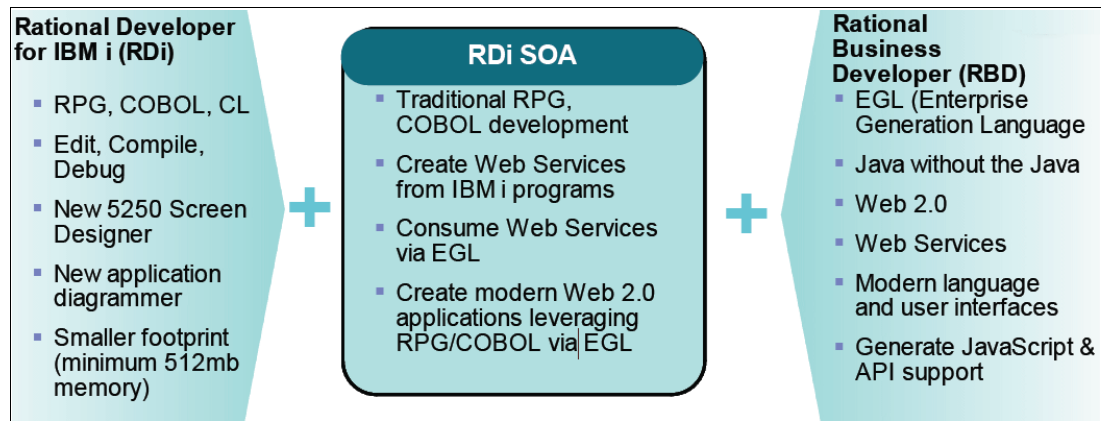


Figure 15-1 Rational Developer for i for SOA construction

15.1.1 Comparison between RDi and RDi SOA

Table 15-1 provides a comparison between the features included in RDi and RDi SOA.

Table 15-1 Comparison between RDi and RDi SOA

Item	RDi	RDi for SOA
Description	Offers edit, compile and debug capabilities to develop applications for IBM i.	A bundle offering to accelerate Web and SOA development for IBM i platform.

Item	RDi	RDi for SOA
Operating System	Windows	Windows
System requirements	<p>Intel Pentium® III 800 MHz processor minimum (higher is recommended) 512 MB RAM; 1 GB RAM recommended. Minimum: 800 MB of disk space is required for product package installation.</p> <ul style="list-style-type: none"> - Additional disk space is required for the resources that you develop. Disk space requirements can be reduced or increased depending on the features that you install. - Additional disk space is required if you download the product package to install this product. - For Windows: Additional disk space is required if you use FAT32 instead of NTFS. - For Windows: An additional 500 MB of disk space is required in the directory pointed to by your environment variable TEMP; 1024 x 768 resolution minimum using 256 colors (or higher for best results); Microsoft mouse or compatible pointing device; CD-ROM drive. 	<p>Processor: Minimum 800 MHz Pentium III (or higher for best results). Memory: Minimum 1 GB RAM, 2 GB works well. Disk space: Minimum 3 GB of disk space is required for product package installation. Additional disk space is required for the resources that you develop. Note. Disk space requirements can be reduced or increased depending on the features that you install. Additional disk space is required if you download the product package to install this product. 1024 x 768 resolution minimum using 256 colors (or higher for best results). Microsoft mouse or compatible pointing device. CD-ROM drive. For Linux: An additional 500 MB of disk space is required in the /tmp directory. Display: 1024 x 768 display minimum using 256 colors (or higher for best results). Processor: Minimum 800 MHz Pentium III (or higher for best results). Memory: Minimum 1 GB RAM, 2 GB works well.</p>
Web services wizards		✓
Integrated Debugger		✓
Rational Business Developer		✓
Edit, compile and debug for RPG, COBOL, C and C++	✓	✓
Application Diagram	✓	✓
Screen Designer	✓	✓
HATS Toolkit		✓
Eclipse-based IDE	✓	✓
Pricing	<p>Refer to:</p> <p>https://www-112.ibm.com/software/howtobuy/buyingtools/paexpress/Express?P0=E1&part_number=D03HCLL,&catalogLocale=en_US&Locale=en_US&country=USA&S_TACT=none&S_CMP=none</p>	<p>Refer to:</p> <p>https://www-112.ibm.com/software/howtobuy/buyingtools/paexpress/Express?P0=E1&part_number=D03L5LL,&catalogLocale=en_US&Locale=en_US&country=USA&S_TACT=none&S_CMP=none</p>

15.2 Rational Developer for i

The IBM Rational Developer for IBM i (RDi) is a new Lightweight eclipse-based workstation offering for IBM i application development. It provides a comprehensive integrated development environment for edit/compile/debug of traditional RPG/COBOL/C/C++ IBM i applications.

Rational Developer for i offers many capabilities previously found only in the WebSphere Development Studio Client (WDSC) Advanced Edition offering. Rational Developer for i includes application diagram, screen designer technology preview phase 2, single sign-on functionality for Web applications, using enterprise identity mapping (EIM), an IBM i program called Java Platform 2, Enterprise Edition Connector (J2C), and integration with IBM Rational ClearCase® software for IBM i.

Some of the benefits of moving from 5250-based tools to RDi include:

- ▶ Integrated set of tools for developing IBM i applications
 - Tools based on Eclipse open source IDE
 - Integrated with change management systems
 - Open source tools
 - Vendor tools, such as the X-Analysis tool from Databorough, an IBM Rational Business Partner
- ▶ Improved productivity for native development
 - Closely integrated edit, compile, run/debug experience
 - Tools to help better understand large applications
- ▶ Lower learning curve for other technologies
 - Learn the Workbench tools with RPG/COBOL development
 - Then branch into Web, XML and Web Services development using the same tools platform

Rational Developer for i V7.5 is based on Eclipse 3.4.1 and includes the following enhancements.

15.2.1 Remote System Explorer

The following are new enhancements in Remote System Explorer (RSE):

IBM i 6.1 host support

Level of Program Verifiers, Syntax Checkers, Content assist, Editor wizards and Outline view are updated to the IBM i 6.1 host level.

Hardware requirement reduction

- ▶ Intel Pentium III 800 MHz processor minimum (higher is recommended)
- ▶ 512 MB RAM; 1 GB RAM recommended
 - The memory requirement of 512 MB is sufficient for most users.
 - The higher memory is recommended for Windows Vista users.

Note: 800 MB of disk space is required for product package installation. Additional disk space is required for the resources that you develop.

- ▶ Disk space requirements can be reduced or increased depending on the features that you install.
- ▶ Additional disk space is required if you download the product package to install this product.
- ▶ An additional 500 MB of disk space is required in the directory pointed to by your environment variable TEMP.

Rational ClearCase integration

Rational Developer for IBM i provides Rational ClearCase integration with IBM i artifacts. This support is available through IBM i Projects which can store local source members and save files (*SAVF).

15.2.2 Application Diagram

The Application Diagram was previously included in WebSphere Development Studio Client Advanced Edition for System i V7.0. The Application Diagram capabilities have now been moved to Rational Developer for i. The new enhancements in Application Diagram are:

- ▶ **Functional zoom**

The Application Diagram now graphically shows subroutines and subprocedures in their source member. The diagram initially opens up showing only source members, programs and service programs, along with any direct (copybooks and bindings) and indirect (procedure and program calls) relationships between them. Source members can now be expanded to see the lower-level details. Expanding a source member shows the call graph for that source member within the graphical representation of the source member in the diagram.

- ▶ **Program calls**

The previous release of the application diagram only showed bound procedure and subroutine calls. In version 7.5 the Application Diagram is updated to also show program calls for ILE RPG, ILE COBOL, and CL source.

- ▶ **CL support**

CL source members are now scanned and their information is added to the generated diagram.

15.2.3 Screen Designer

The Screen Designer was previously included in WebSphere Development Studio Client Advanced Edition for System i V7.0. The Screen Designer capabilities have now been moved to Rational Developer for i.

In RDi V7.1 the Screen Designer technology preview phase 2 contains improved usability enhancements. It now reflects more closely its intended appearance and function:

- ▶ Technology Preview Keyword specific property pages
- ▶ The ability to specify keyword properties when adding a keyword
- ▶ A toolbar button on the design page for configuring preferences
- ▶ Partial preview page function, including indicator sets and field values

- ▶ Convenient buttons to control screen size from both the design page and the preview page
- ▶ Screen information stored as XML comments in the display file

15.2.4 Screen Designer technology preview phase 2

In V7.5, Screen Designer technology preview phase 2 contains improved usability enhancements. It now reflects more closely its intended appearance and function. The enhancements include:

- ▶ Performance enhancements
- ▶ Defect fixes

15.2.5 Debug

The enhancements in debug include:

- ▶ Select and debug each RPG thread in the application, view and change variables, and pause a thread.
- ▶ View and modify RPG varying length variables.
- ▶ Show the data item information in COBOL OCCURS DEPENDING ON (ODO) properly and more easily.
- ▶ View and modify decimal floating points in ILE C.

15.2.6 IBM i Web services and Java tools

The Program Call wizard now enables the generation of Program Call Markup Language (PCML) and direct program call Java beans, which use the ProgramCall class to invoke the IBM i program. PCML program call Java beans use the ProgramCallDocument class to invoke the IBM i program. Generated program call beans are more extensible. Communications between generated program call beans and IBM i now support SSL.

15.2.7 Integration with IBM Rational Application Developer

The integration with IBM Rational Application Developer for WebSphere Software V7.5 allows shell sharing with other IBM Rational V7.5 products.

Note: When Rational and WebSphere products are installed on a single system, they share a single common framework. This approach is known as shell sharing.

The first product that you install also installs the common framework. When you install additional products, the existing framework is used and product-specific plug-ins only are installed.

It also enables a series of service oriented architecture (SOA) wizards to ease the creation of SOA wrappers of an IBM i program or service program.

15.2.8 Change management

There are increased capabilities in the IBM i Project support for centralizing IBM i resources in a common change management repository:

- ▶ You can now create a single IBM i Project to contain all of the source of your application, if your workstation's disk space and communication speed with IBM i allow. In the previous release, only a small subset of files could be contained.
- ▶ IBM i Project can now store an increased number of source files and save files, if your workstation's disk space and communication speed with IBM i allow.

15.2.9 Other WebSphere Development Studio Client functions

Other functions from WebSphere Development Studio Client Advanced Edition for System i that are now included in IBM Rational Developer for i are:

- ▶ IBM i program call J2C connector
- ▶ Single sign-on using EIM for Web applications
- ▶ Screen designer technology preview phase 1
- ▶ ClearCase and IBM Rational Developer for i integration for i

15.2.10 Improved performance

Performance has been improved in the following areas:

- ▶ System actions, such as expanding Library objects in the Remote System view, are now faster than in the previous release.
- ▶ Long-running actions are now performed in the background to allow you to continue editing or navigating the IDE.
- ▶ Startup time is reduced when not shell sharing with other larger IBM Rational products.
- ▶ IBM i Remote System Explorer tools are now based on the open source Remote System Explorer included in the Eclipse Target Management project.

15.2.11 Reduced resource requirements

The new version requires less resources by:

- ▶ Reducing the hard disk and main memory requirements for the developer personal computer.
- ▶ Improving performance at startup and some system actions.

15.3 Using Rational Developer for i

This section provides an overview of the welcome page displayed when RDi is started. The welcome page is the first page you see when you first launch eclipse, as shown in Figure 15-2. Its purpose is to introduce you to the product. Welcome content typically includes an overview of the product and its features, tutorials to guide you through some basic tasks, samples to get you started, and more.

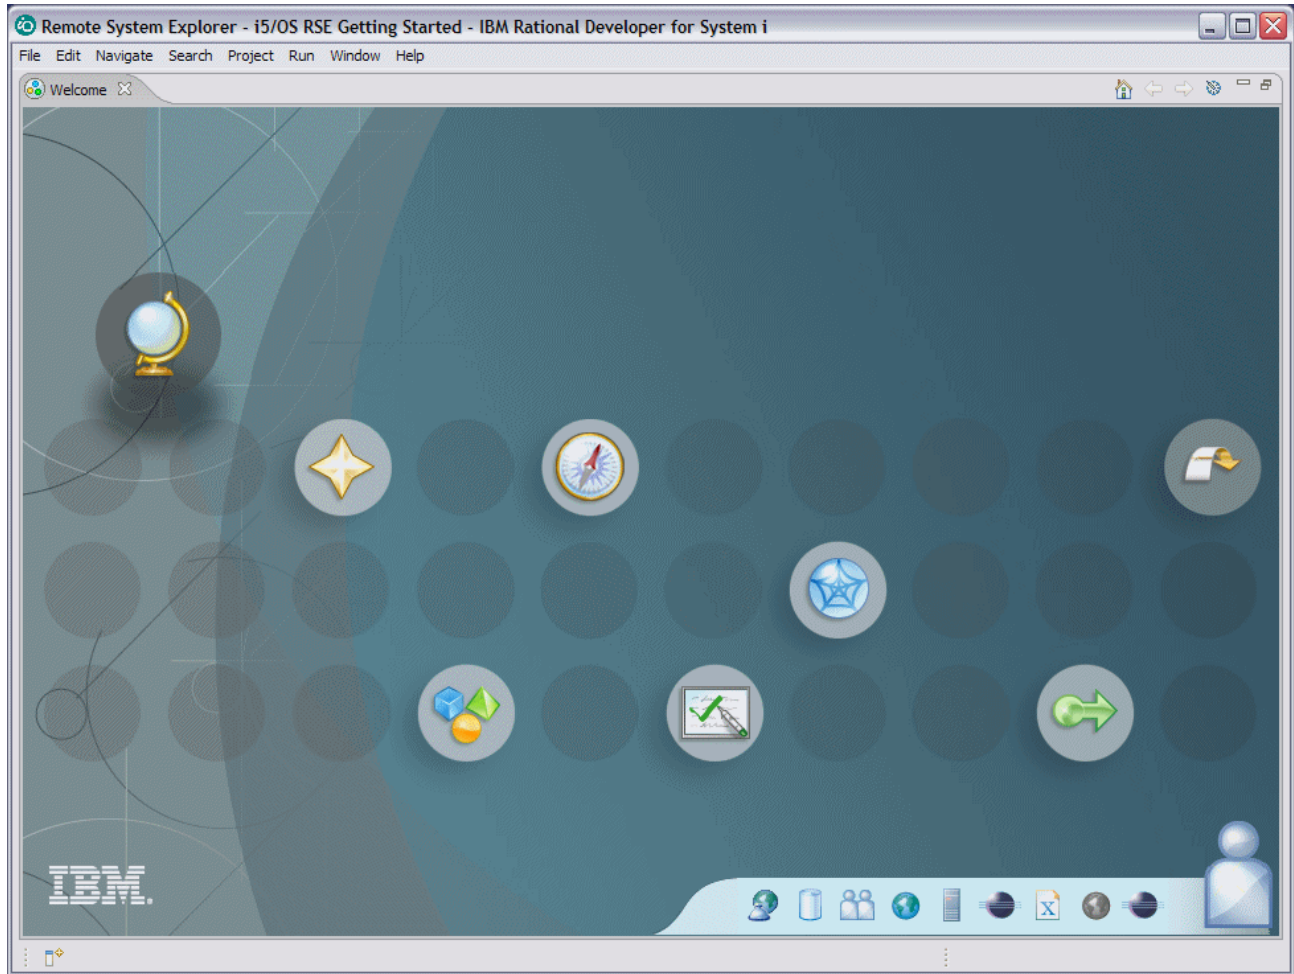


Figure 15-2 RDi welcome panel

From the welcome page you can select the following actions:

- ▶ Overview
- ▶ What's new
- ▶ First Steps
- ▶ Web Resources
- ▶ Tutorials
- ▶ Migrate
- ▶ Go to Workbench
- ▶ Enable Roles

For detailed information about each feature, refer to RDi product help. RDi help can be reached from the Workbench menu by choosing the path **Help** → **Help Contents** → **Rational Developer for System i**.

You can also refer to the IBM Rational Developer for IBM i Web page at:

<http://www.ibm.com/software/awdtools/developer/rdi/>

15.3.1 Remote System Explorer

The Remote System Explorer (RSE) is more than just a graphical user interface. It is a workbench perspective that provides access to all development resources of your IBM i

server. You can think of it as an enhanced and more flexible workstation version of the Programming Development Manager (PDM). It allows you to effectively manage and organize all the IBM i resources and applications needed by application developers in an easy and user-friendly way.

When you first open the workbench, the RSE is open by default. RSE allows you to work with your libraries, objects, members, jobs, and IFS files, as well as edit source members, compile programs, and debug applications. The key components of an RSE perspective are shown in Figure 15-3.

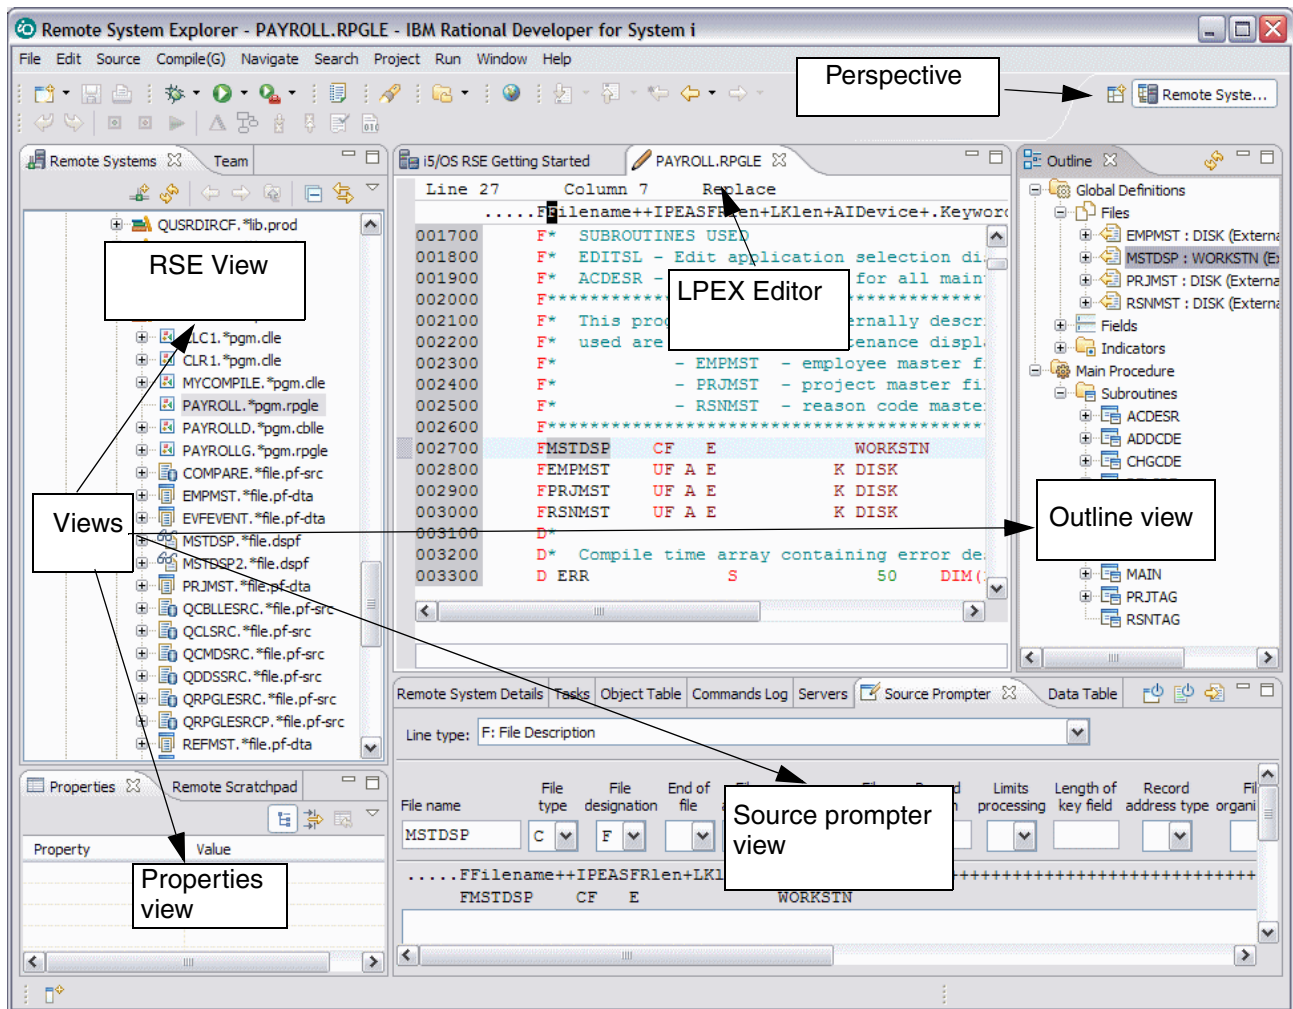


Figure 15-3 Remote System Explorer perspective

From the Remote System Explorer in RDi you can:

Connect to multiple IBM i servers

You can create connections to IBM i servers from RSE by expanding the **New Connection** entry in the tree, and then expanding **i5/OS** as shown in Figure 15-4.

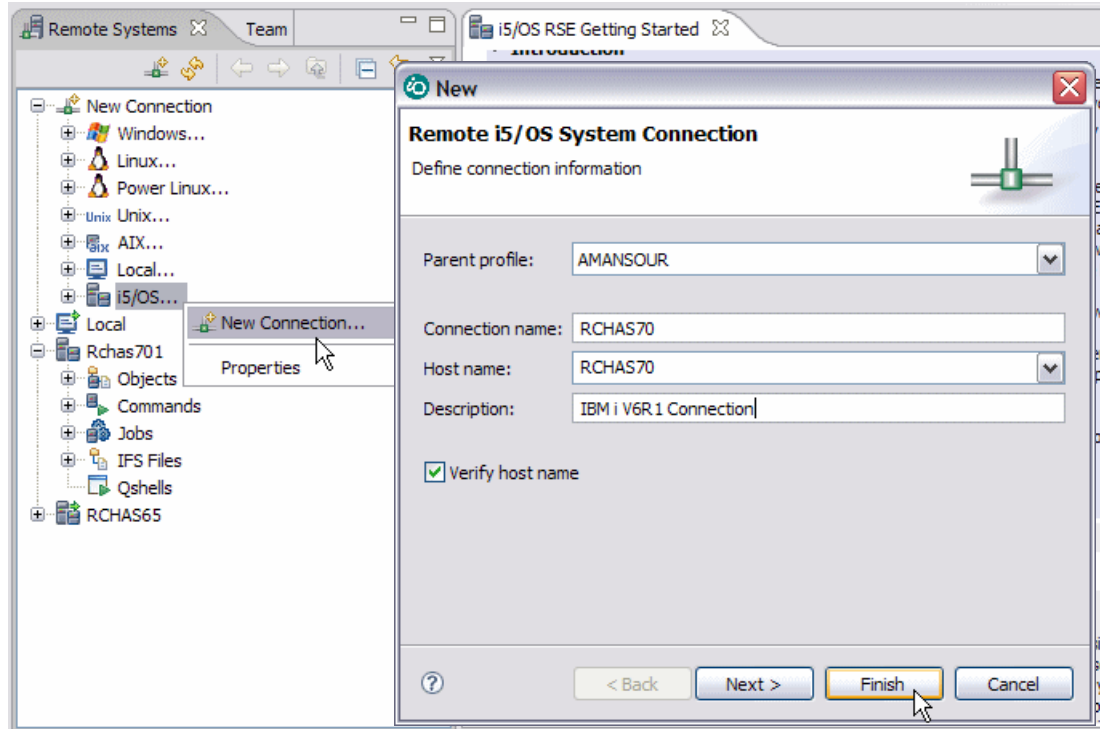


Figure 15-4 Create connection to IBM i from RSE

Work with subsystems

After you configure a connection to the IBM i, you can easily connect and expand your new connection to reveal your subsystems. Subsystems are predefined filters that group the various types of remote resources that can be explored in the remote system. There are four subsystems, as shown in Figure 15-5 on page 371:

- Objects** A PDM-like group that allows access to libraries, objects, and members.
- Commands** Contains predefined commands and enables you to create command sets, each of which contains one or more often-used commands. When run, all commands in a command set are sent to the remote system and executed, and the results are logged in the Commands log.
- Jobs** Enables you to see various jobs, subset by job attributes, and to perform operations on those jobs.
- IFS Files** Enables you to explore folders and files in the Integrated File System of the remote i5 system.

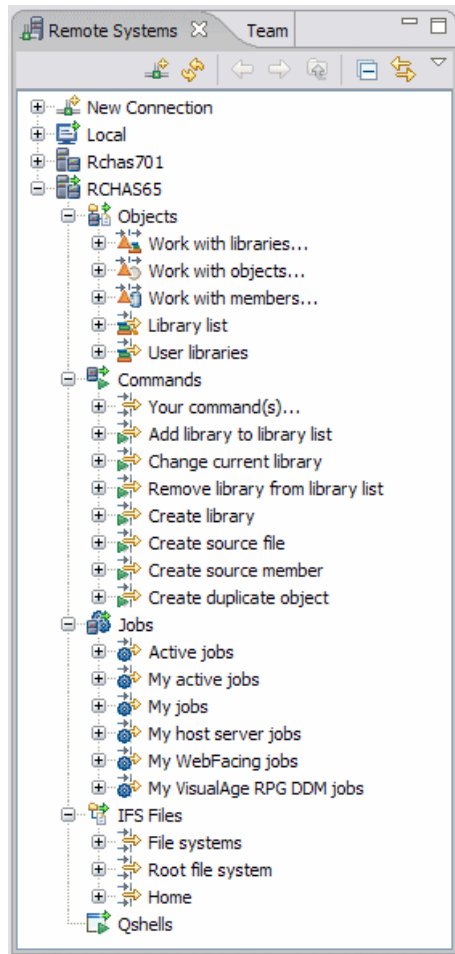


Figure 15-5 IBM i subsystems from RSE

Table views

Table views are an alternative for displaying information in the RSE. Table views take the currently selected object in the Objects, Jobs, or IFS Files subsystems as input and display the contents in a table. With this view, you can see item properties at the same time since they are displayed as rows in a table, as shown in Figure 15-6 on page 372. You can perform the following tasks with the Table view:

- ▶ Navigation
- ▶ Command prompting and running
- ▶ Direct editing of table entries
- ▶ Object and source sub-setting and positioning
- ▶ View locking and unlocking
- ▶ Programming Development Manager actions
- ▶ Displaying of additional information
- ▶ List printing

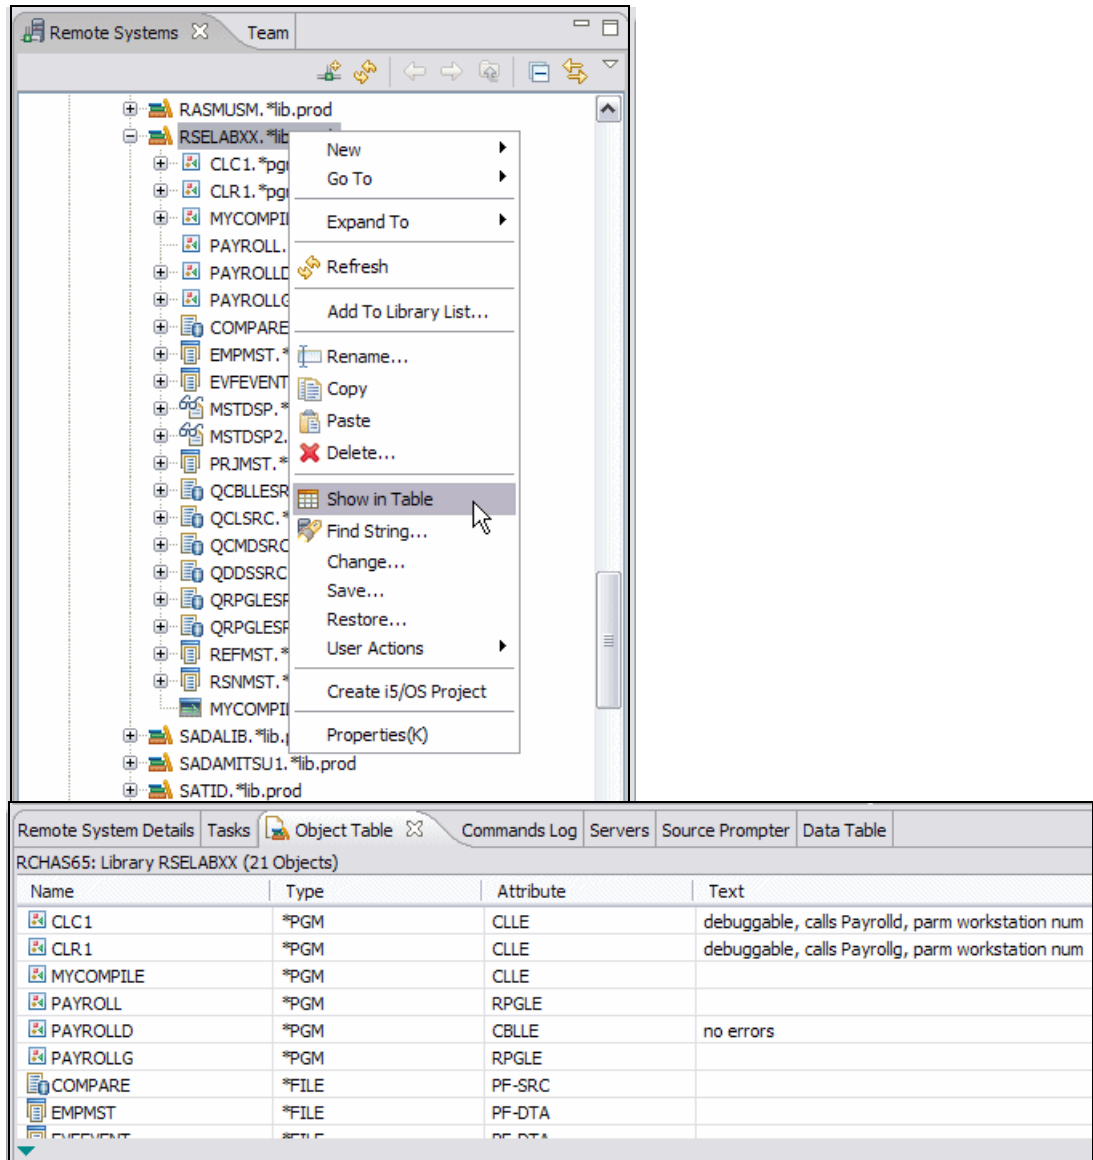


Figure 15-6 RSE Table view

Other Remote System Explorer functions

Some other functions that you can perform from RSE are:

- ▶ Compile source members
- ▶ Run programs
- ▶ Debug programs
- ▶ Add, change, and remove libraries from the library list

15.3.2 i5/OS Projects

i5/OS Projects allow for disconnected development. A live connection is required only when code updates or builds are needed, or when you need to view remote resources for a project. If you are working on a remote laptop, or are at home with a slower connection, you might choose to work in disconnected mode. With this option, you explicitly copy files locally while

you have a live connection, close the connection, work on the files at your leisure, and upload them back to the server once you have finished. While working in disconnected mode, you can still check source code for syntax and semantic errors and connect only to submit a compile when you are ready to create the program object.

The i5/OS Projects perspective has four main views as shown in Figure 15-7:

- ▶ i5/OS Project Navigator
- ▶ Remote Systems
- ▶ Properties
- ▶ Tasks

To open the i5/OS Projects perspective, follow these steps:

1. From the workbench menu, select **Window** → **Open Perspective**.
2. Select **i5/OS Projects**. If this option does not appear on this submenu, select **Other...** The Select Perspective window opens. Select **i5/OS Projects** from this window and click **OK**.

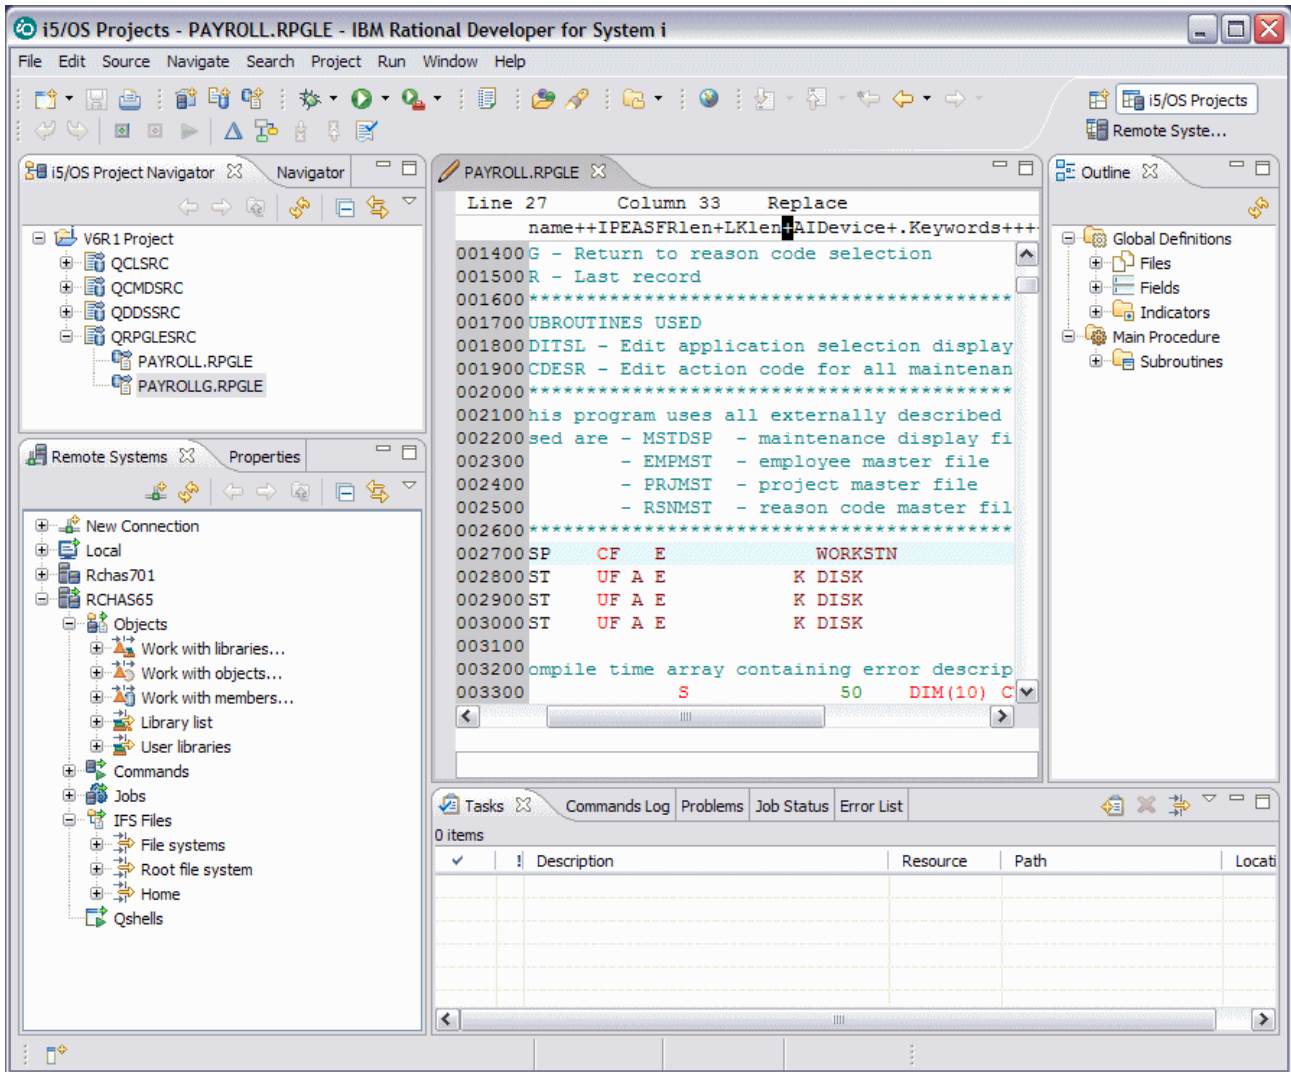


Figure 15-7 i5/OS Projects

15.3.3 Application Diagram

The Application Diagram provides a graphical view of the different resources in an IBM i native application and their relationships to each other. There are two different diagrams that you can look at in the Application Diagram view:

- ▶ Source Call Diagram
- ▶ Program Structure Diagram

The Source Call Diagram takes ILE RPG, ILE COBOL, and CL source as input and displays a call graph showing subroutine and procedure calls. An example of this is shown in Figure 15-8 on page 375, where we have selected the PAYROLL RPG program and visualize the application diagram.

The Program Structure Diagram takes program and service program objects as input and displays the binding relationships between them as well as the modules bound into each program and service program.

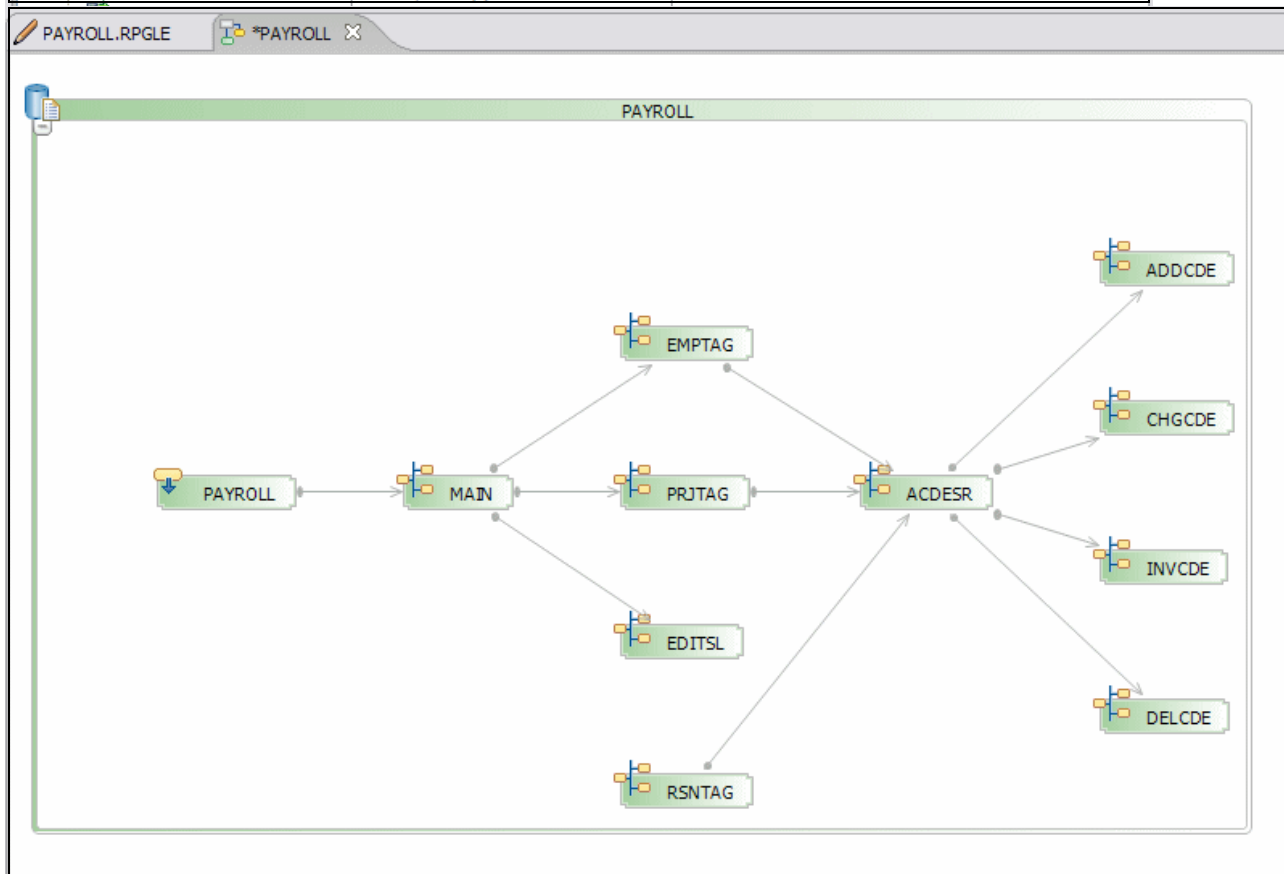
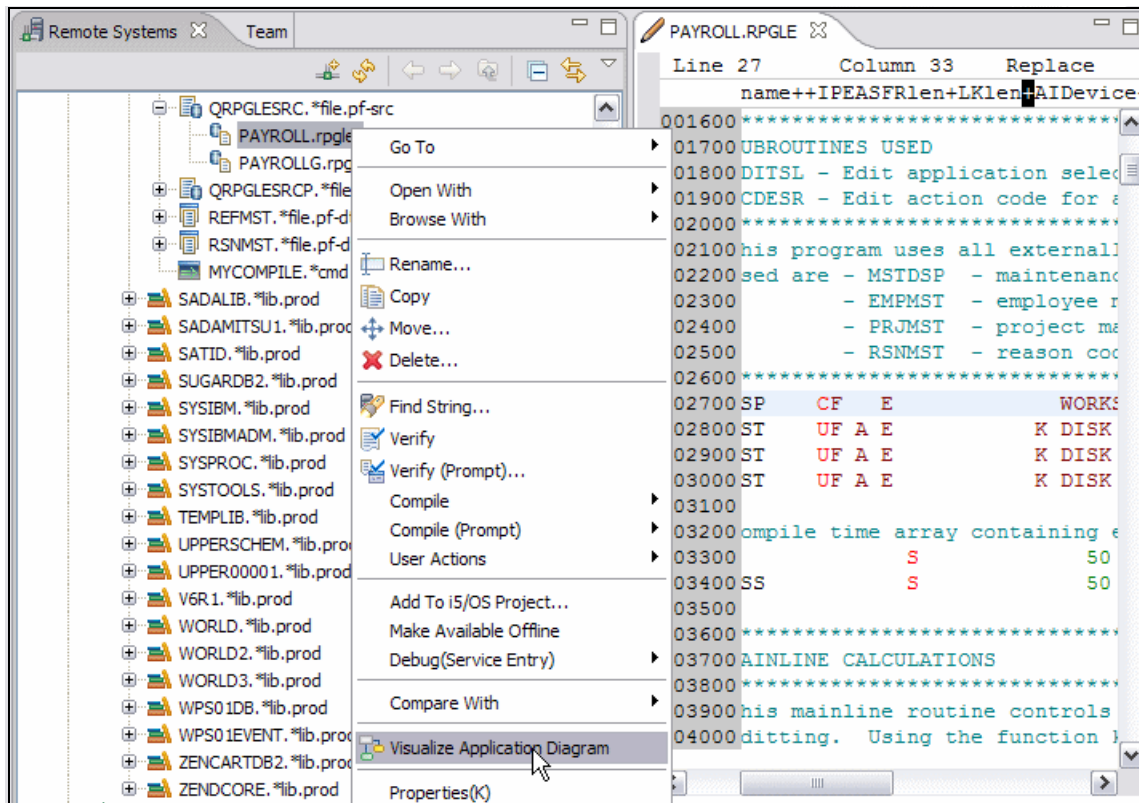


Figure 15-8 Application Diagram

Note: You can use both the Outline and Properties views with the Application Diagram viewer. The Outline view initially shows a thumbnail view of the entire diagram. This can be useful when the entire diagram does not fit within the visible space of the viewer.

15.3.4 Screen Designer

The Screen Designer is a technical preview that provides the capability to graphically design and modify the content of DDS display files.

The Screen Designer design page provides an integrated palette. Items on the palette may be dragged from the palette and dropped onto the design area to create a new item.

You may launch the Screen Designer either from the Remote Systems view or the i5/OS Projects Navigator view, by right-clicking any DDS source member of a source type Display File (DSPF) or DDS menu (MNUDDS) and selecting **Open With** → **Screen Designer (Technology Preview)** from the context menu, as shown in Figure 15-9. This opens the Screen Design view shown in Figure 15-10 on page 377.

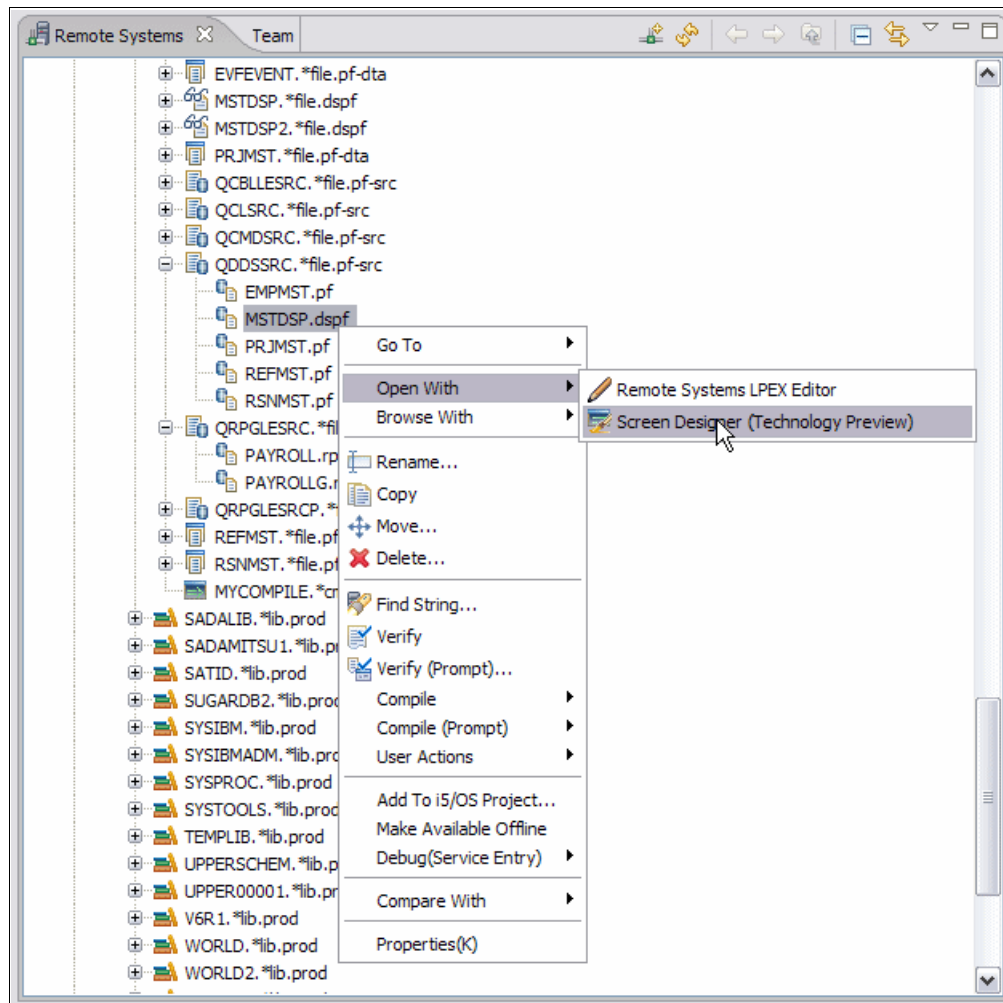


Figure 15-9 Launching Screen Designer

When the Screen Designer is opened, one of two tabbed pages may be used to edit a display file. The Design page may be used to graphically modify panels, records and fields. The

Source page is available for directly editing the display file source code. Screen Designer updates all related views when changes are made either to the Source page, the Design page, or the Properties view.

The Screen Designer design page provides an integrated palette. Items on the palette may be dragged from the palette and dropped onto the design area to create a new item. The palette may be made to automatically hide or be shown, or may be set to remain hidden or shown. Alternatively, the Workbench's Palette view may be used. Whenever the Palette view is closed, the Screen Designer's integrated palette is shown. Whenever the Palette view is opened, the Screen Designer's integrated palette is hidden.

Above the design area, a Controls area is available for defining and managing panels. Panels are only a design time concept and a panel is defined as a group of records. An All records panel contains all records and is always present and cannot be deleted.

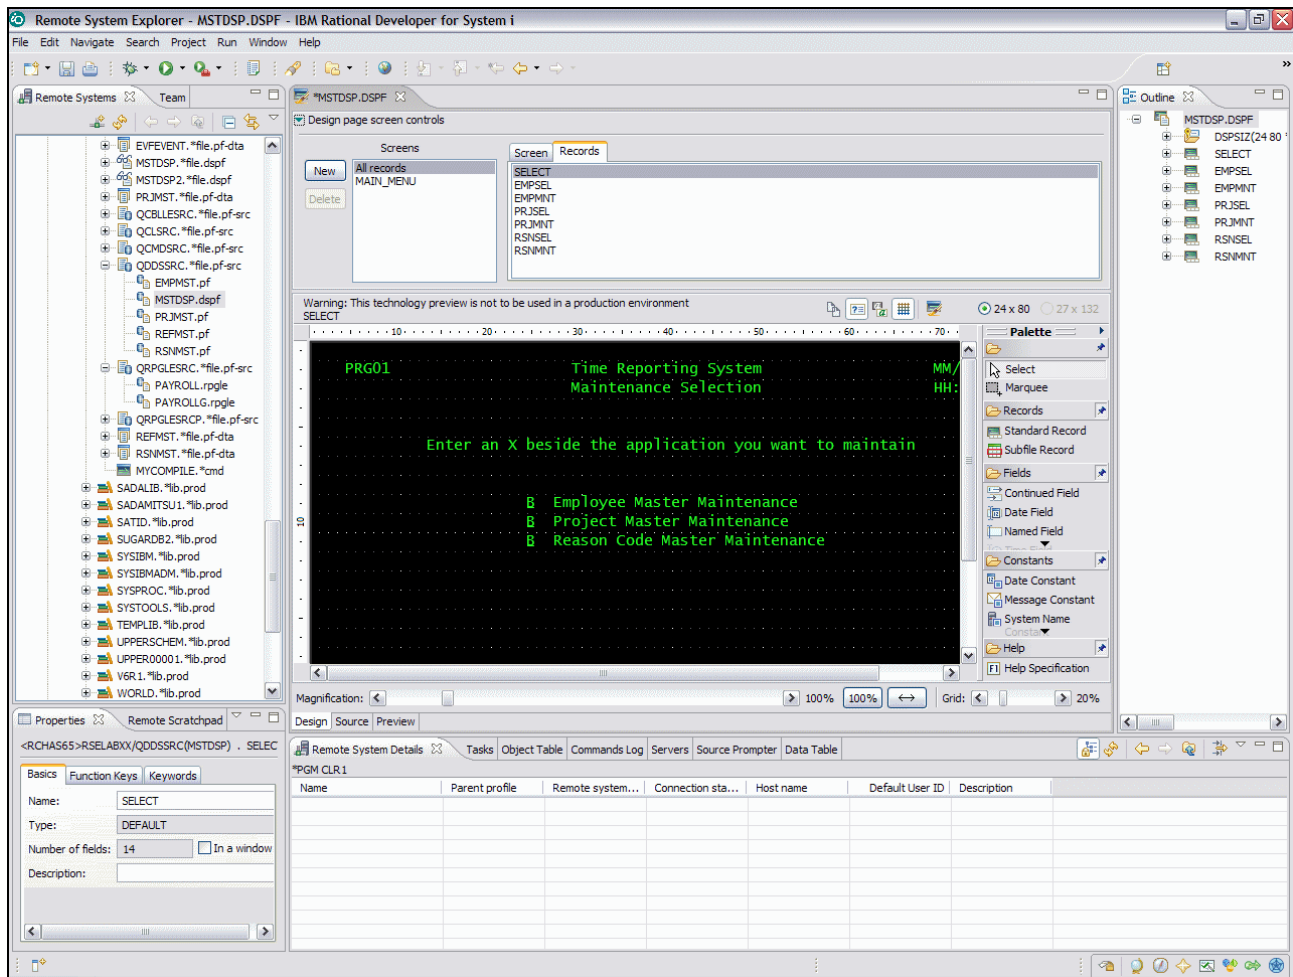


Figure 15-10 The Screen Designer in the Remote System Explorer perspective

15.3.5 Integrated IBM i Debugger

With the Integrated IBM i Debugger you can debug an application that is running on an IBM i system. It provides an interactive graphical interface that makes it easy to debug and test your IBM i programs. It is fully integrated into the workbench.

You can also set breakpoints before running the debugger, by inserting breakpoints directly into the source while editing.

The Integrated System i debugger client user interface also enables you to control program execution. For example, you can run programs, set line, watch and service entry breakpoints, step through program instructions, examine variables, and examine the call stack.

You can also debug multiple applications, which may be written in different languages, from a single debug window. Each session you debug is listed separately in the Debug view.

15.3.6 Web Services and Java Tools for IBM i

IBM i Web development tools give you the ability to create new e-business applications that use a Web-based front end to communicate with the business logic in an ILE or non-ILE language program residing on a IBM i server.

IBM i Web services are self-contained, modular applications that can be described, published, located, and invoked over the Internet, employing “just-in-time” integration of services.

The Web Service wizard creates a bottom-up IBM i Program Web Service, with Web service operations that call one or more server programs or service program procedures on the IBM i.

The Web Service wizard works in the context of a Web project and allows for creation, deployment, testing, generation of a proxy, and publication to a Universal Description, Discovery, and Integration (UDDI) registry of Web services.

Important: Web services are only available with Rational Developer for i for SOA Construction.

An example of how to access Web service creation is shown in Figure 15-11 on page 379.

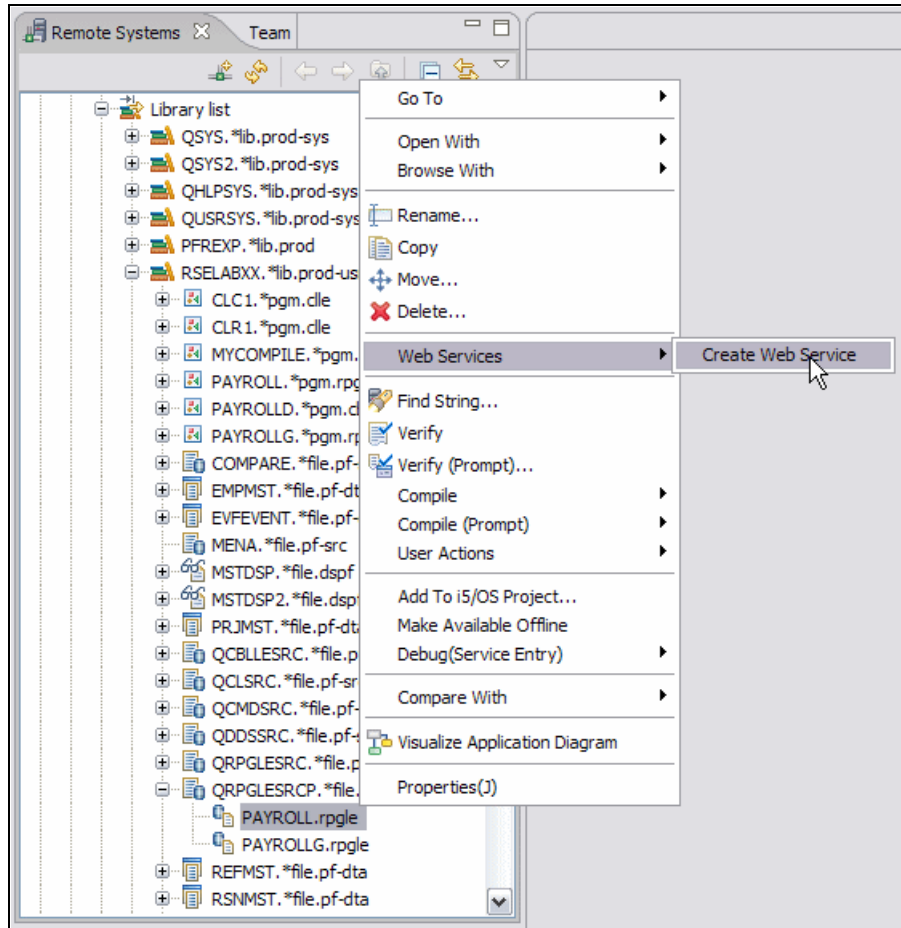


Figure 15-11 Create Web service

15.4 Rational Business Developer

IBM Rational Business Developer (RBD) V7.1 delivers an Eclipse integrated developer environment (IDE) for EGL. EGL is a modern language designed to shield IBM i application programmers from the technical complexities of Web and service-oriented architecture (SOA) middleware and standards. EGL provides a simpler and more abstract development paradigm to help accelerate development of robust and sophisticated Web applications and service-oriented solutions.

EGL allows you to embrace and leverage the strengths of Java and Java 2 Platform Enterprise Edition (J2EE) without forcing your entire team of developers to learn and master the low-level intricacies of these platforms. IBM Rational Business Developer includes state-of-the-art tools that help simplify creation and debugging of EGL applications and services. It has an EGL compiler that transforms the EGL source into modules that can be deployed to either a Java runtime or to ILE.

By hiding the details of the runtime platforms and associated middleware, EGL frees you to focus on the business problem rather than on the underlying implementation technologies. If you have little or no experience with Java and Web technologies, you can use EGL to create enterprise-class services and applications. If you have little or no experience with IBM i programming, you can create highly optimized IBM i services and applications.

EGL is designed to address a full spectrum of business application requirements by allowing development of:

- ▶ **Business Services**

The language built-in notion of service enables the creation and consumption of services in a simplified and straightforward way.

- ▶ **Web applications**

Tight integration of EGL with Java Server Faces (JSF) framework and tooling enables the creation of Web applications without the need to know either Java or the details of the JSF framework.

- ▶ **Reports**

EGL integration with BIRT, an open source reporting engine, allows the creation of professional reports.

- ▶ **Batch systems**

The language includes the notion of batch program, which can be generated to run without user interaction; for example, to produce reports or to batch database or file loads and updates.

- ▶ **Traditional 5250 transactional applications**

EGL includes language constructs. IBM Rational Business Developer includes tools to support the development of 5250 character-based dialogs and relative business logic.

IBM Rational Business Developer empowers you to increase your productivity for IBM i application development through:

- ▶ **Abstraction**

EGL provides concise and powerful notations that help to eliminate tight coupling. It reduces the amount of coding required to interface systems and middleware. This abstraction simplifies and speeds up your development work.

- ▶ **Declarative programming**

EGL includes a certain level of declarative specifications to help you reduce repetitive and error-prone coding. For example, validation rules associated to a data item trigger validation to be executed whenever the item is used, in a Web page or a 5250 panel.

- ▶ **Language**

EGL is a comprehensive but easy-to-learn language. It is modern, modular, and readable. It has a rich library of built-in functions to boost your productivity for commonly required operations, such as date and time math, string manipulation, and others. In addition, the language is extensible and offers full interoperability with other languages, including EGL interfaces to native Java, and seamless invocation of RPG or COBOL programs or any ILE procedure.

- ▶ **Tools**

IBM Rational Business Developer contains a rich set of Eclipse-based capabilities, including EGL source-level debugging, powerful smart editing, visual construction, graphical navigation, and automated generation of Create/Read/Update/Delete (CRUD) applications from Unified Modeling Language (UML) models or from data schemas.

- ▶ **SOA**

EGL is designed to facilitate services development and deployment. A simplified SOA development paradigm is built into the language itself and complemented with the tools. You can create services without the need to know Web service protocols and standards

such as Web Services Description Language (WSDL), Simple Object Access Protocol (SOAP), or Extensible Markup Language (XML).

15.4.1 Rational Business Developer enhancements

IBM Rational Business Developer includes the following enhancements to further extend the ability to build modern business-critical IBM i solutions:

- ▶ EGL support for accessing Data Queues and Data Areas
- ▶ Ability to create Portlets in EGL
- ▶ Message Queuing (MQ) support through EGL language abstractions
- ▶ Invocation of ILE procedures
- ▶ Ability to create Web applications for deployment to the IBM i Integrated Application Server
- ▶ JSF enhancements (for example, AJAX, type ahead, and security)
- ▶ IDE productivity enhancements: Formatting, dynamic error markers, and debugger hot swapping
- ▶ Improved Web services development
- ▶ Ability to invoke EGL programs from Java

For more information and material about Rational Business Developer, refer to the IBM Web site at:

<http://www.ibm.com/software/awdtools/developer/business/index.html>

15.4.2 Enterprise Generation Language enhancements

EGL offers the following enhancements:

- ▶ Support for legacy technologies
Support for Text User Interface, Web transactions, Java wrappers, and migration from VisualAge® Generator is included in this version.
- ▶ Editor enhancements
The EGL editor can format your code automatically. It also highlights syntax errors as you type.
- ▶ Data access technologies
In this version of EGL, you can access DL/I, IMS™, and MQ data sources.
- ▶ Runtime data access connections
The EGL Runtime Data Source wizard helps you connect to a database, sets your build descriptor options based on that connection, and sets the options necessary to use that connection at runtime.
- ▶ Web application security
EGL provides support for J2EE container-managed security through system functions that verify the remote user's name, authentication type, and role. EGL also provides a code sample that demonstrates how to access security details from an LDAP-compliant server.

15.5 Using Enterprise Generation Language

In this section we introduce a simple dynamic Web site using EGL. This site has two pages: one to display a list of records in a database and another to allow you to change the data from one of those records.

In this tutorial, you learn how to complete these tasks:

- ▶ Create and configure an EGL project
- ▶ Create EGL source code that accesses a data source
- ▶ Create two simple Web pages that access data in a relational database
- ▶ Pass a parameter from one Web page to another
- ▶ Test an application on a Web application server

To be able to create EGL-related files and projects in Rational Business Developer you will have to enable the EGL Development capability. To enable this capability:

1. In RBD Workspace select **Window** → **Preferences**.
2. In the Preferences Dialog expand **General** → **Capabilities**.
3. Select **EGL Developer** as shown in Figure 15-12 and click **OK**.

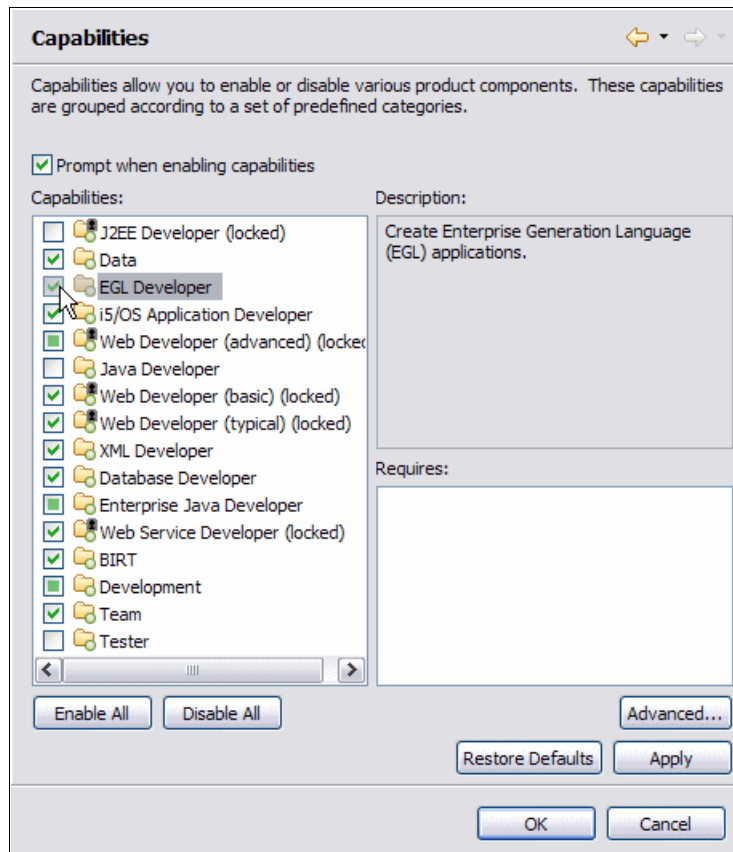


Figure 15-12 Enable EGL Developer Capabilities

The project will have a Web page interface, so we need to create an EGL Web project, not an EGL project. An EGL Web project combines the features of a Dynamic Web project and an

EGL project. If you are using WebSphere Application Server, you will also create an EAR project.

Create an EGL project:

1. From RBD Workspace click **File** → **New** → **Project**.
2. Select **EGL Project** as shown in Figure 15-13 on page 383, and click **Next**.

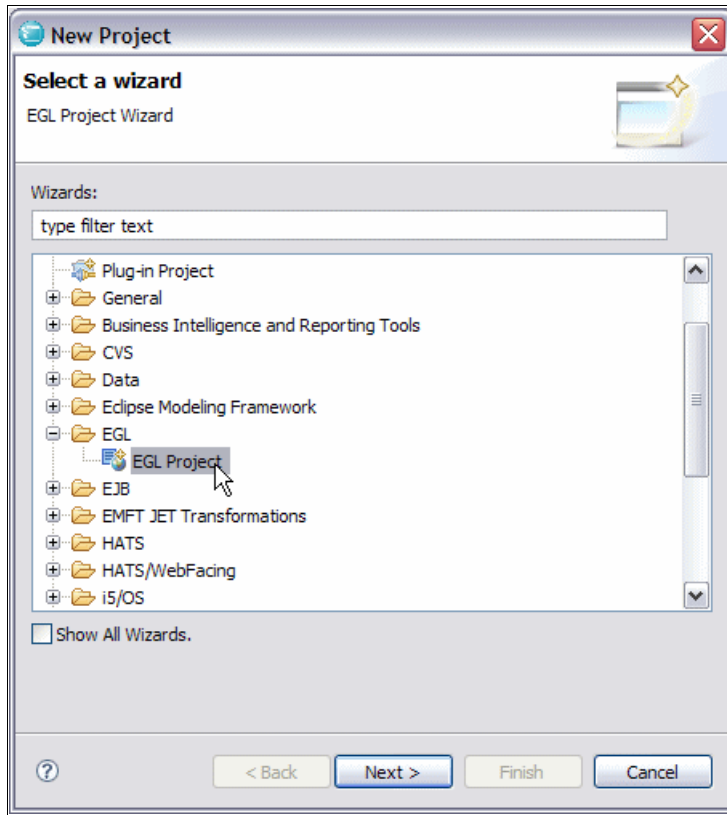


Figure 15-13 New EGL Project Wizard

You have to select a target runtime for your project. This application server can be WebSphere Application Server or Apache Tomcat.

Note: If you are using WebSphere Application Server, you will also create an EAR project in addition to the EGL Web project.

The EAR project contains information about deploying an application in the J2EE framework, including how to run it on a server and how to connect it to data sources.

In this example we create an EGL application that will be used as a Web service, and use an existing Web service in an EGL application. This example shows you how to:

- ▶ Create an EGL project
- ▶ Create and Test EGL Web services
- ▶ Consume a third-party Web service

EGL implements two categories of services:

- ▶ EGL service
For applications written entirely in EGL
- ▶ EGL Web service
Created using EGL, but can be accessed from both EGL and non-EGL clients.

Start the Rational Developer for i for SOA workbench, and select a workspace that will contain all of the files that you use for your development project, and perform the following steps:

1. In RDi open the Web perspective from **Window** → **Open Perspective** → **Other** → **Web** as shown in Figure 15-14.

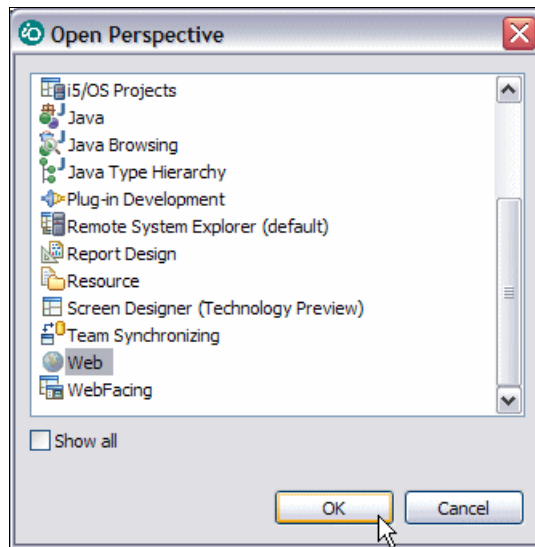


Figure 15-14 Open Web Perspective

2. Select **File** → **New** → **Project**, and from the dialog select **EGL Project** as shown in Figure 15-15, and click **Next**.

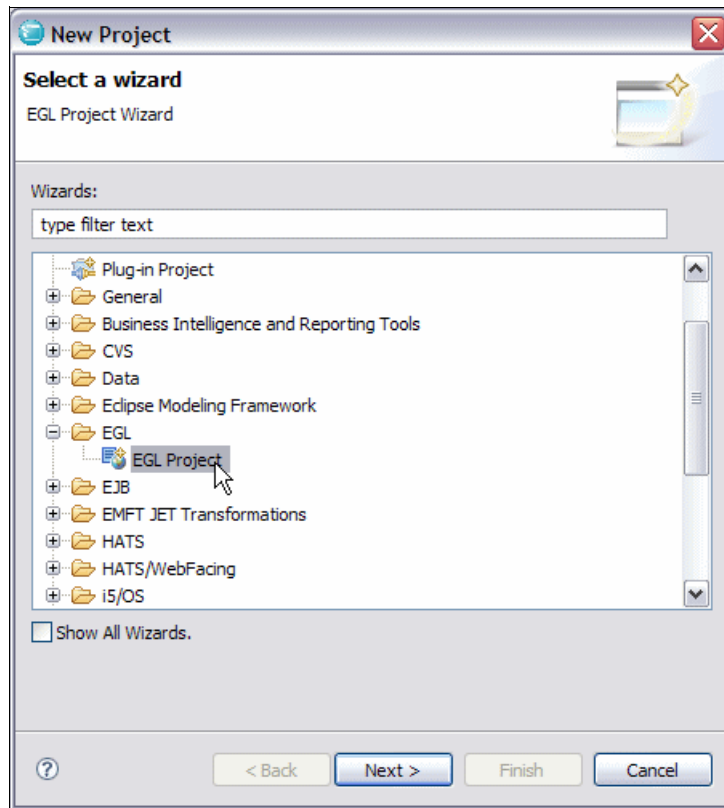


Figure 15-15 New EGL Project Wizard

3. In the New EGL Project dialog, select **Web Project**, give a name to the project and click **Next**.

4. In the EGL Web project dialog select **WebSphere Application Server v6.1** → **Create new project build descriptor**, as shown in Figure 15-16, and click **Finish**.

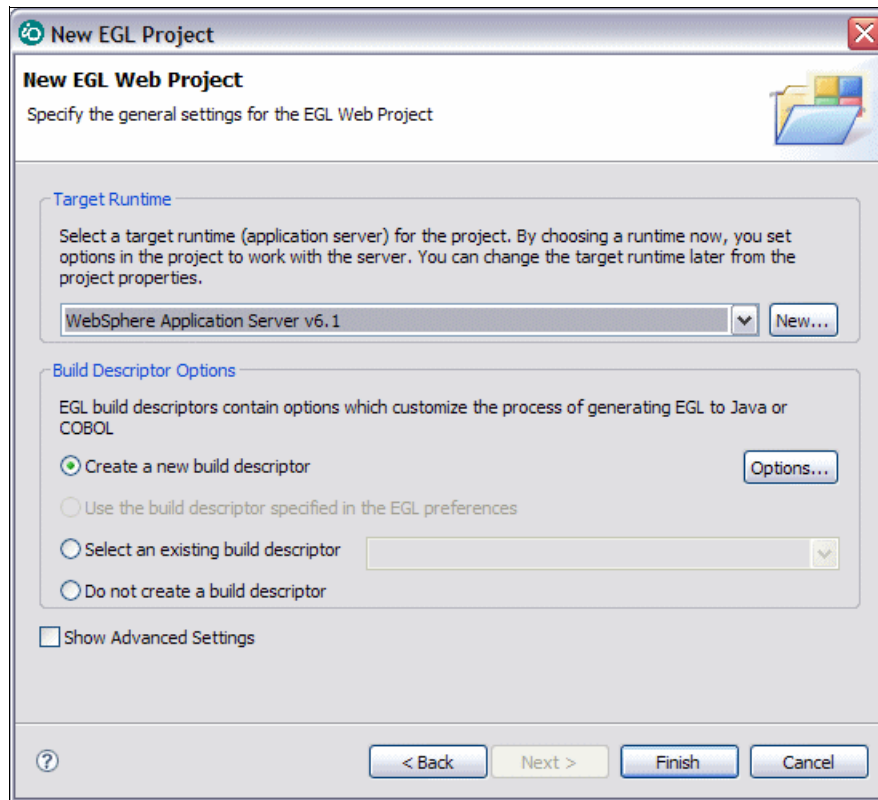


Figure 15-16 EGL Web Project Settings

5. In the project explorer, expand the new project, right-click the **EGLSource** folder and select **New** → **Service**, as shown in Figure 15-17.

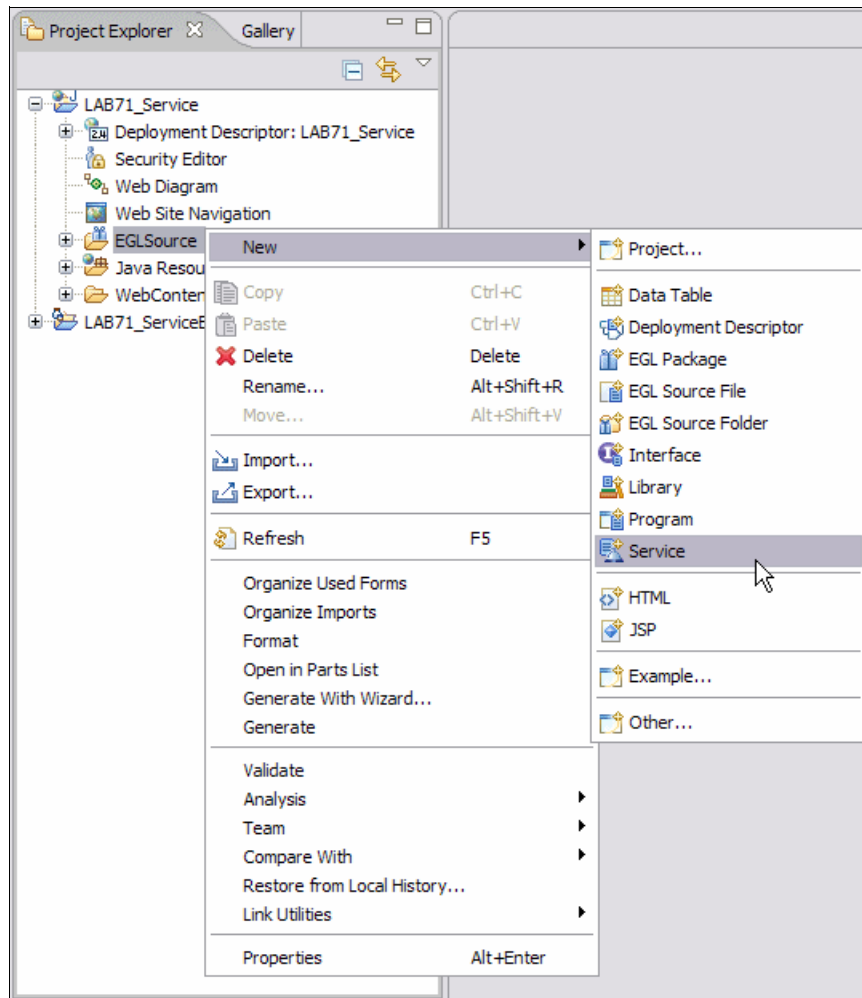


Figure 15-17 Create New Service

6. In the New EGL Service Part dialog, enter the Package name, the EGL source file name, and click **Finish**.

7. Create some logic in this source file to provide the functionality for this service. In this example a very simple service is used. The service checks whether or not a ZIP code entry is a valid California ZIP code, as shown in Figure 15-18.

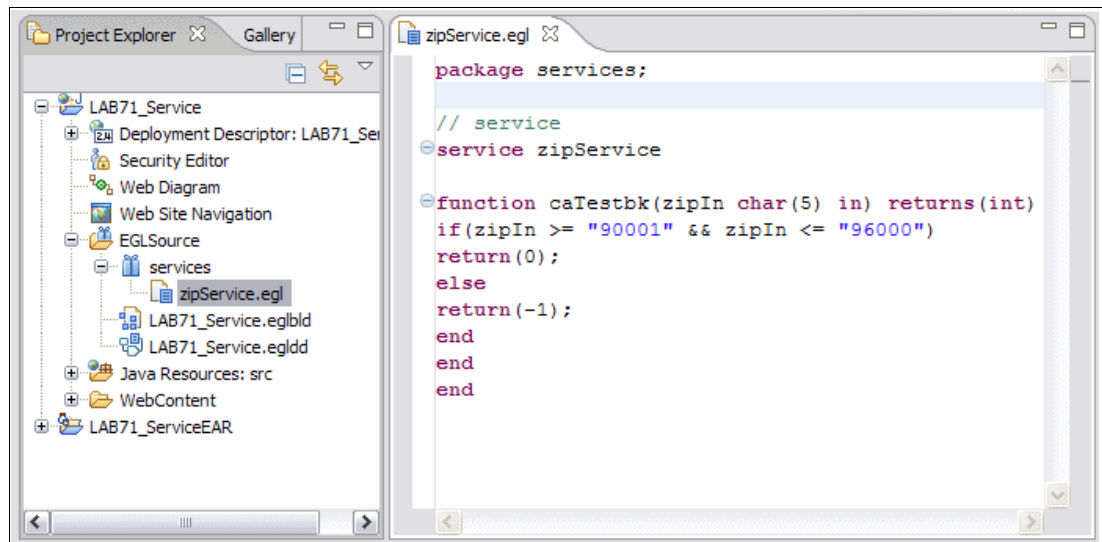


Figure 15-18 Source to validate ZIP code

8. To describe this service to the EGL environment, you need to add it to an EGL Deployment descriptor file, which has been created already during project creation. To do this:
 - a. Expand the EGL Source directory.
 - b. Locate the ProjectName.egl dd file, and double-click the file to open it in the editor as shown in Figure 15-19.

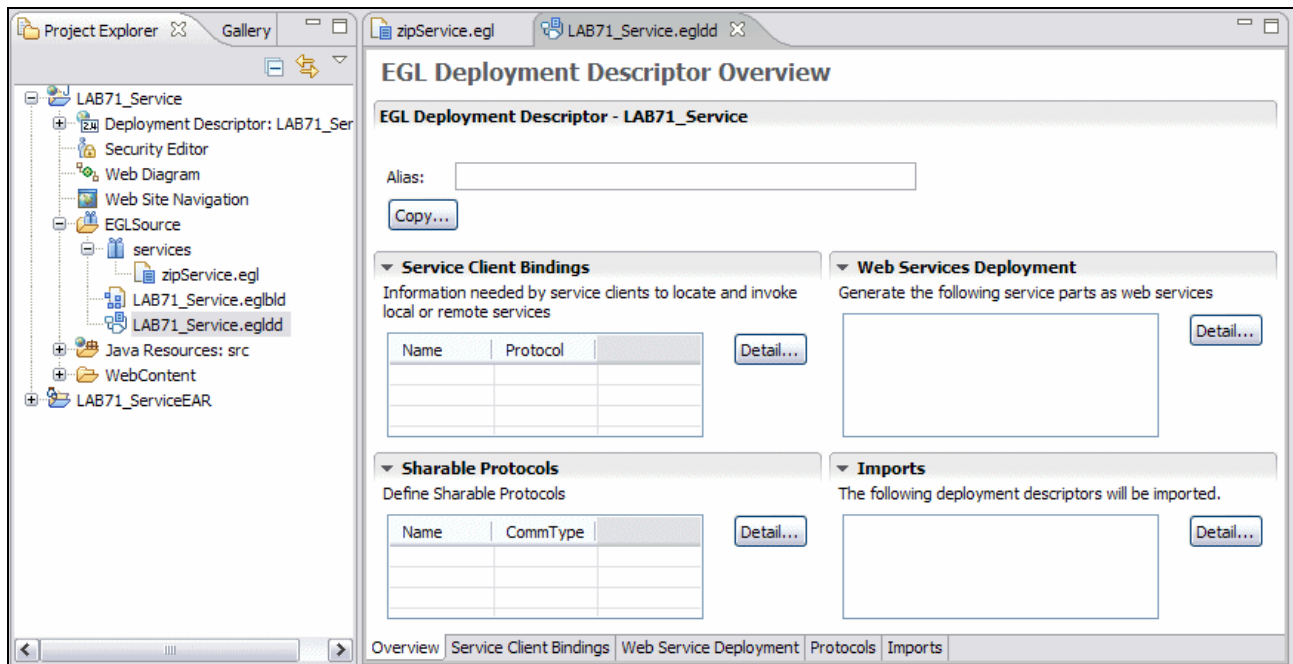


Figure 15-19 EGL Deployment Descriptor

9. Click the **Web Service Deployment** tab, and click **Add** in the window that displays. This will show the Add Web Services dialog, as shown in Figure 15-20.
 - a. Select your service in the **EGL service parts found** list.
 - b. Click **Add**. The service is moved to the EGL service parts to be generated as Web services list.
 - c. Click **Finish**. Save and close the DeploymentDescriptor file.

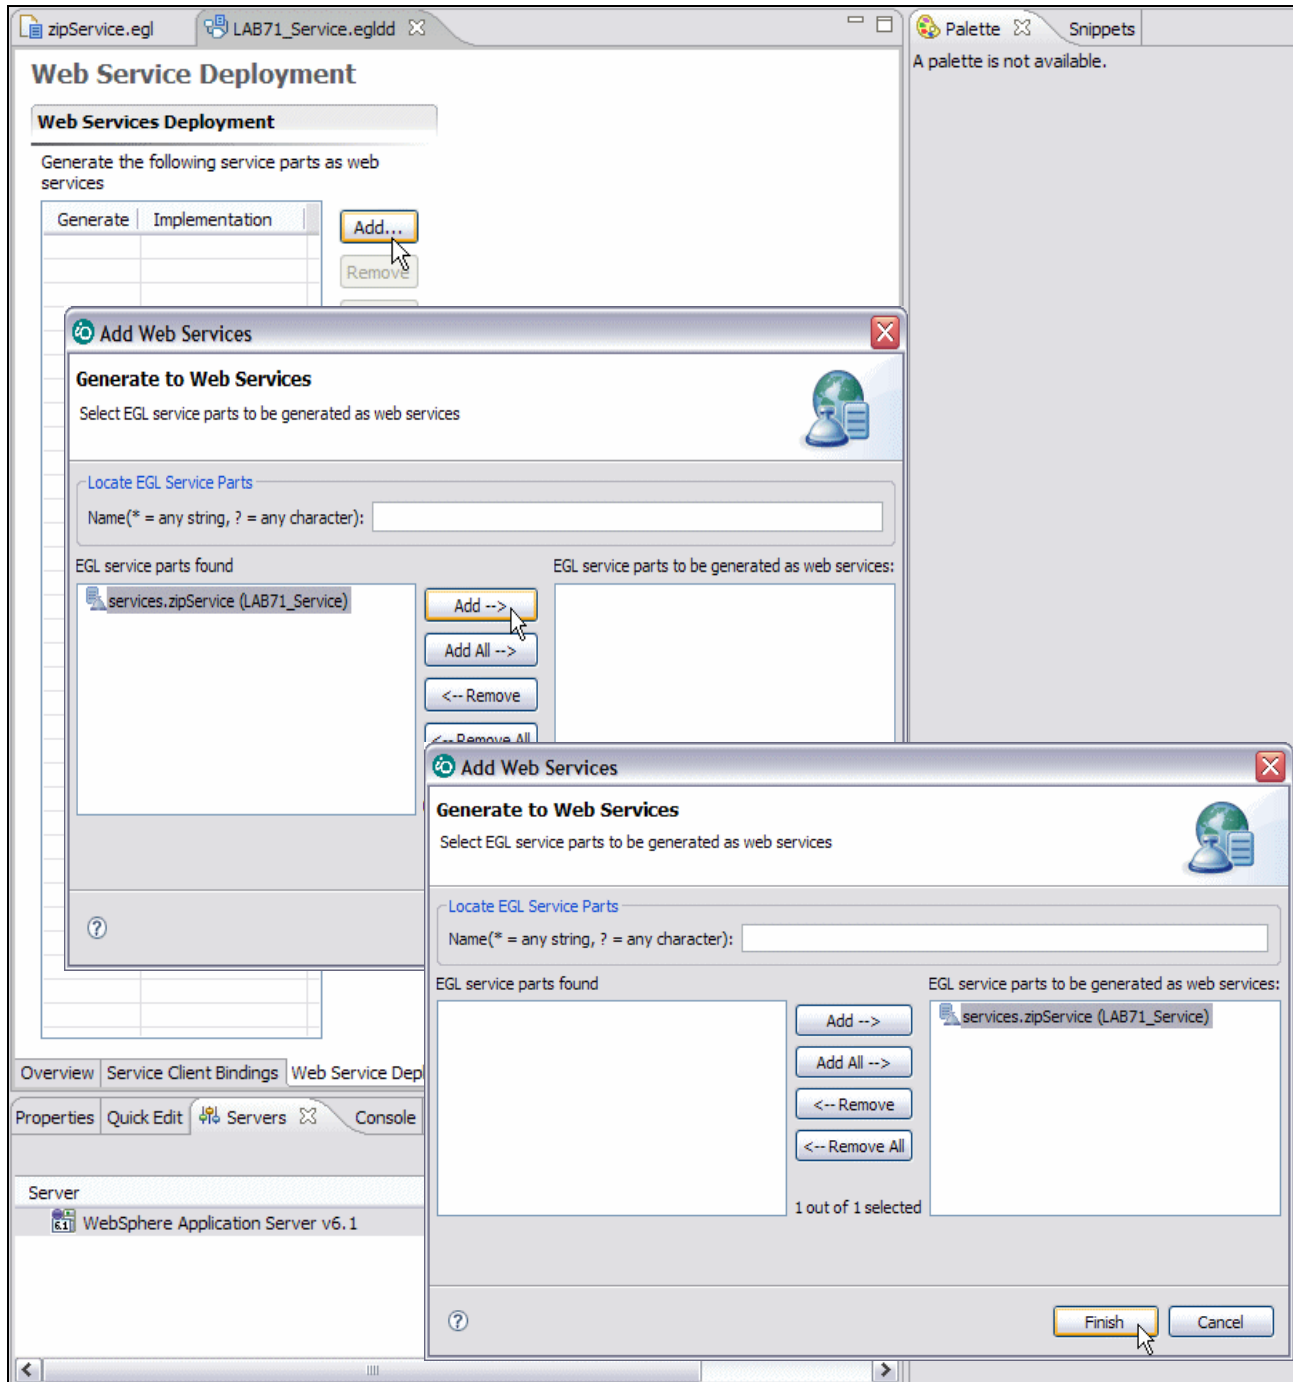


Figure 15-20 Add service to an EGL file

The Deployment descriptor has been automatically added to the build file (.egld). This is necessary to make the Deployment descriptor known to the EGL build environment.

10. Regenerate the project so that EGL can take this new information and apply it to the generated application. In the Project Explorer, right-click the project name and click **Generate**, as shown in Figure 15-21.

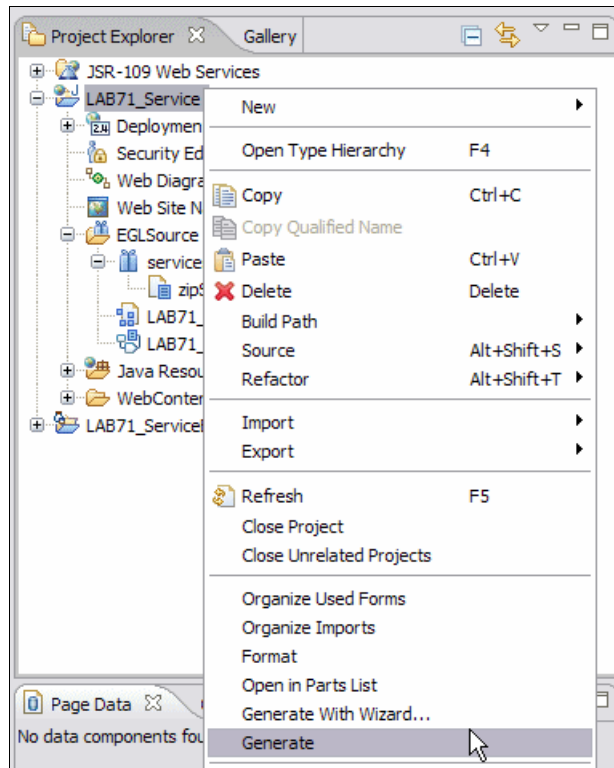


Figure 15-21 Regenerate EGL Project

Note: The Generate action generates the complete Web service, including the WSDL file. The WSDL is a standard file format that describes a Web service. You can hand over this WSDL file to people that want to use your Web service. It describes the service interface, the location of the service, and other characteristics of the service.

11. Test the Web service interactively
 - a. Click the **Servers** tab to select it. Right-click **WebSphere Application Server v6.1** and click **Start**, as shown in Figure 15-22 on page 391.

Once the server has started, the console will show Server server1 open for e-business, and the Servers tab will reflect started under the Status column.

Note: You can choose a local WebSphere Application Server installed on your computer, or a remote Websphere Application Server installed on a remote host.

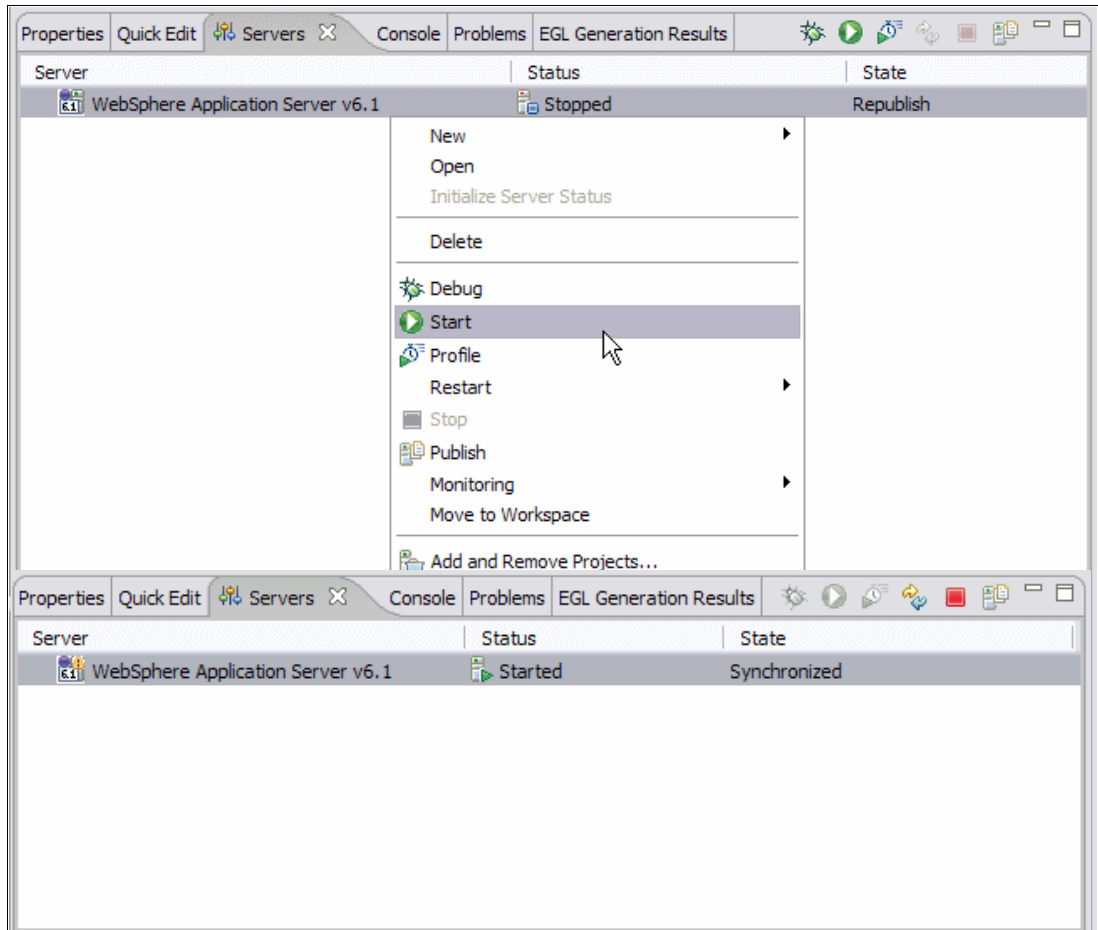


Figure 15-22 Start Websphere Application Server

- b. Add the project to the server. Right-click **WebSphere Application Server v6.1**, then click **Add and Remove Projects**, as shown in Figure 15-23 on page 392. In the Available Projects list, select your **ProjectNameEAR**, and click **Add**.
The project moves under Configured Projects. Click **Finish**.

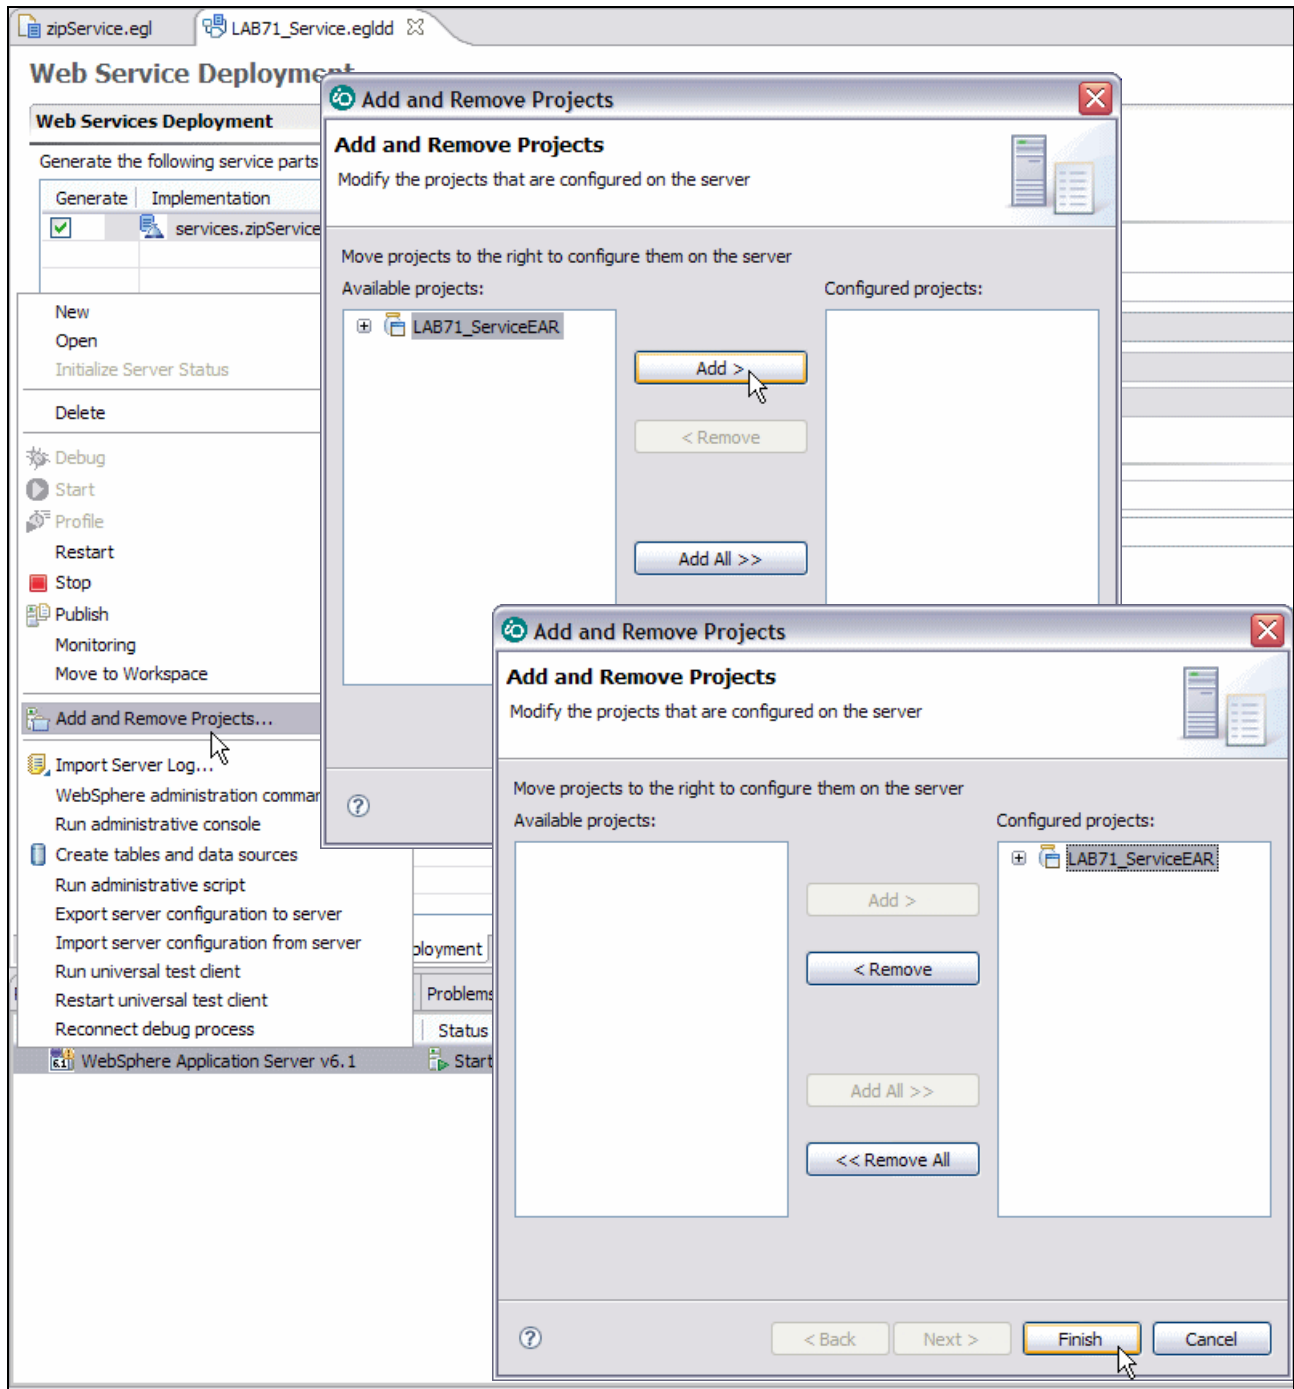


Figure 15-23 Add project to Websphere Application Server

- c. Rational Developer for i for SOA construction offers a useful interactive Web services test facility called the Web services Explorer.
- d. In the Project Explorer expand the following: **Your Project** → **WebContent directory** → **WEB-INF directory** → **WEB-INF directory** → **wsdl directory**. Right-click **YourServiceName.wsdl**, and select **WebServices** → **Test** → **Test with Web services Explorer**, as shown in Figure 15-24 on page 393.

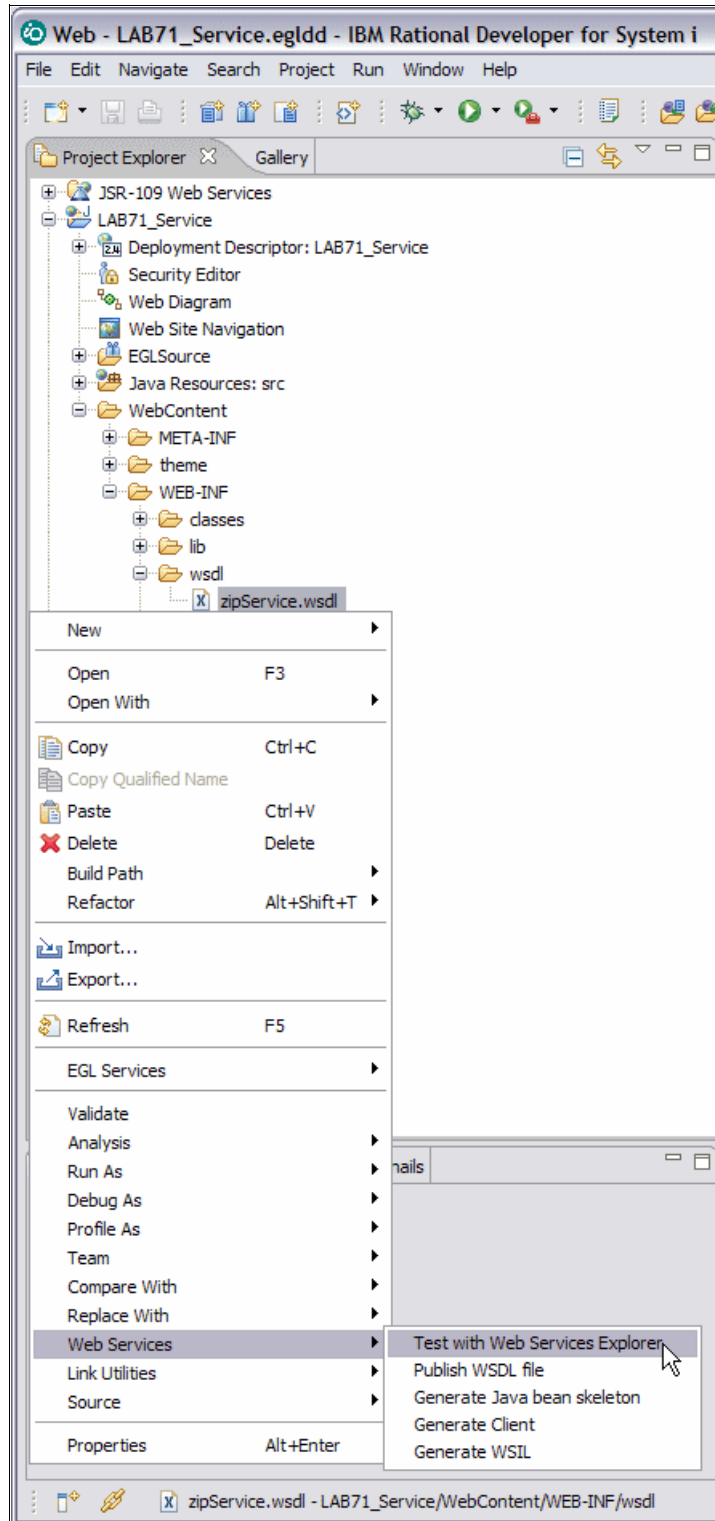


Figure 15-24 Start Web services Explorer

- e. The Web services Explorer will appear in the browser in the Content Area. Expand **YourServiceNameBinding**, and click the **caTestbk** method which is the service you are going to test. In the Actions dialog, enter values for testing your service, as shown in Figure 15-25 on page 394, and click **Go**.

Check the result from the Web service on the bottom Status dialog.

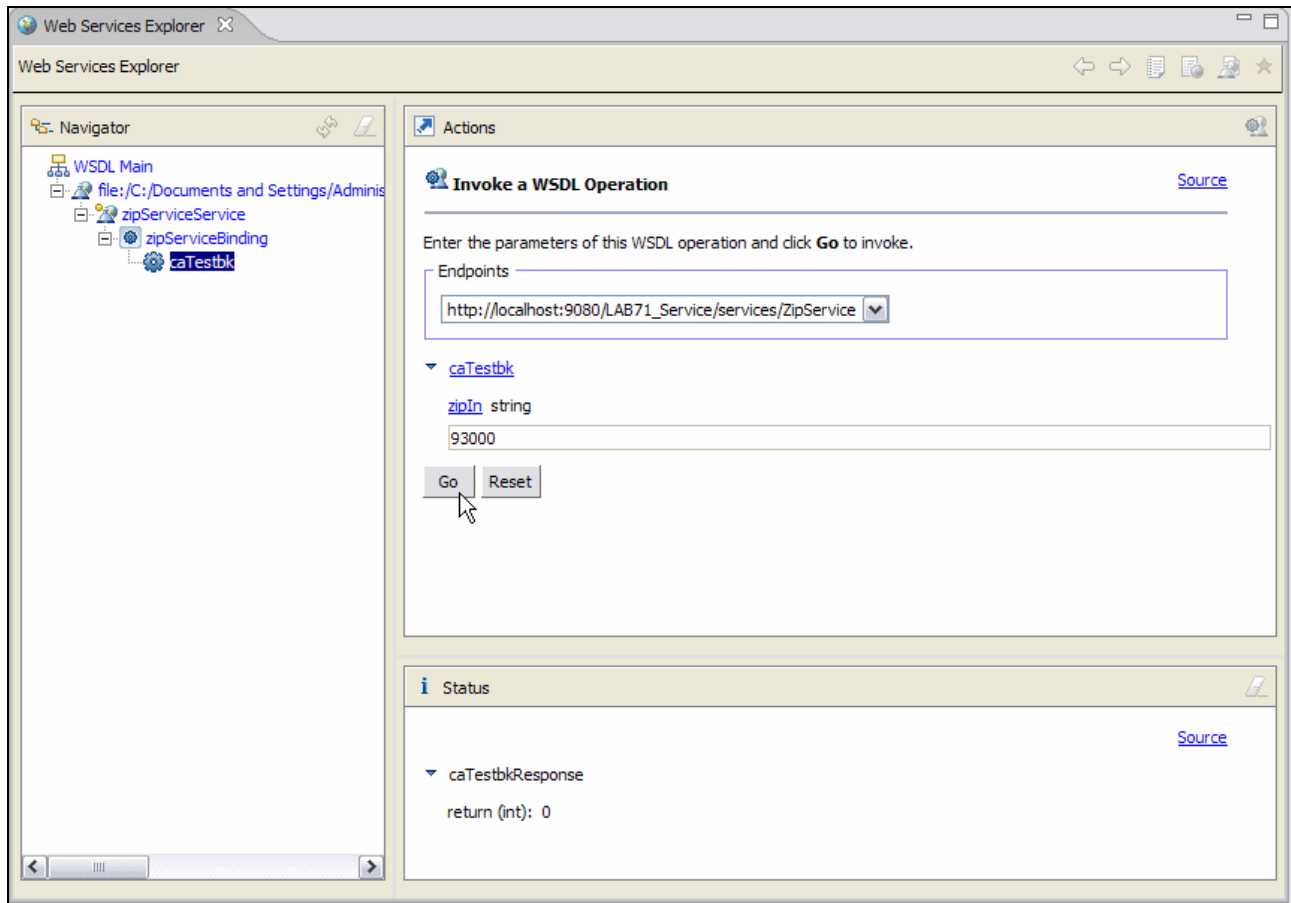


Figure 15-25 Test in Web services Explorer

For an example of calling IBM i programs from EGL using Rational Business Developer, refer to:

http://www.ibm.com/developerworks/rational/library/08/dw-r-ibmirbd/?S_TACT=105AGX28&S_CMP=EMSAND

15.6 Rational Host Access Transformation Services for 5250

IBM Rational Host Access Transformation Services (HATS) enables organizations to quickly modernize existing 3270 or 5250 green screen applications without impacting these proven, reliable applications. With HATS, applications can be transformed as Web, rich client, portlet, or mobile applications, or can be SOA-enabled as Web services.

You can make your 5250 applications available either as HTML through the most popular Web browsers or as a native client application running on an Eclipse Rich Client or Lotus Expeditor Client platform. You can also help improve the workflow and navigation of your host applications without access or modification to source code.

The HATS rules-based transformation engine makes it possible to extend your host applications to the Web or rich client platform. The power of HATS lies in its ability to accurately recognize the components of host panels and transform them in real time to a Web or rich client interface according to a set of predefined rules. You can modify the rules

according to the specific needs of your application. With HATS, you can add a variety of elements to your host panels, such as drop-down lists, hot links, tables, buttons, valid value lists, tabbed folders, and graphs.

The IBM Rational Host Access Transformation Services macro support allows programmed navigation through multiple terminal panels. Examples of productivity improvements include skipping and combining panels, pre-filling text entry fields, and storing and retrieving data in global variables. IBM Rational Host Access Transformation Services can give you such programmed access to a single host application, or can integrate panels from multiple host applications into a single interface.

HATS allows you to represent core business application logic and transactions contained in your green panel applications as programmed navigation through multiple terminal panels. These core business tasks can be reused as standard Web services that participate as an integral part of your business process integration plan. Use your core business services as building blocks to develop new internal applications or to integrate with applications outside your enterprise. Your host-based business tasks can then be included in your service-oriented architecture (SOA) solutions with IBM SOA Foundation products, such as WebSphere Process Server, WebSphere Enterprise Service Bus, and others.

15.6.1 WebFacing and HATS

WebFacing technology provides the ability to convert IBM i data description specification (DDS) display file source members into a Web-based user interface for existing 5250 programs.

Prior to HATS V7.1, the WebFacing technology was provided as part of a separate product. With HATS V7.1 the WebFacing technology is provided by installable HATS features. You can now install the IBM WebFacing Tool for IBM i and the Portal support for IBM WebFacing Tool for IBM i features as part of the HATS installation process.

The Portal support for the IBM WebFacing Tool for IBM i feature provides portal support for the WebFacing feature.

WebFacing information is also available in the HATS Information Center at:

<http://publib.boulder.ibm.com/infocenter/hatshelp/v71/index.jsp>

15.6.2 HATS interaction with WebFacing

HATS provides two methods for interacting with the WebFacing technology. One method allows a HATS Web application to use a connection to a WebFacing server instead of a Telnet 5250 server. The second method allows a HATS Web application and a WebFacing application to be linked together and interoperate as a single enterprise application that can use a single connection to a WebFacing server. Each method is described here.

HATS connection to a WebFacing server

HATS provides two connectivity options for connecting to 5250 host panel applications:

- ▶ 5250 (Telnet)
- ▶ 5250W (WebFacing)

There is no difference in the data flow between HATS and a Telnet 5250 server or HATS and a WebFacing server. Both connections provide access to 5250 applications, but the connection to the WebFacing server does not require online transaction processing (OLTP)

capacity (interactive CPW). In Power6 terminology this mean there is no requirement of an enterprise enablement (5250 OLTP). This significantly reduces the overall cost of using HATS with 5250 applications by removing the interactive feature requirement associated with TN5250 session connectivity.

To enable a connection to a WebFacing server, select the 5250W connection type from the Type field in the New Project wizard, the New Connection wizard, or the Connection editor. The default port for the WebFacing server connection is 4004.

You can enable existing projects to take advantage of the WebFacing connectivity option by selecting 5250W from the Connection editor. The WebFacing server connection can serve as the default connection or as a background connection in the project.

Connection to the WebFacing server requires the server process to be running on the 5250 host system. To start the WebFacing server, enter the following command from the operating system command line:

```
STRTCPSVR SERVER(*WEBFACING)
```

To stop the WebFacing server, enter the following command from the operating system command line:

```
ENDTCPSVR SERVER(*WEBFACING)
```

To ensure that the WebFacing server is running, enter the following command from the command line to list all WebFacing jobs:

```
WRKACTJOB JOB(QQF*)
```

The WebFacing runtime server job, QQFWFSVR, and the WebFacing virtual terminal (VT) server job, QQFVTSVR, should both be displayed. It is normal to see more than one QQFVTSVR job in the list of jobs.

You can also view and manage the WebFacing server-related jobs as follows:

- ▶ From a Windows PC workstation using System i Navigator, expand **Network** → **Servers** → **TCP/IP servers**.
- ▶ From a browser connected to IBM i Systems Director Navigator for i, expand **Network** → **Servers** → **TCP/IP servers**.

Figure 15-26 on page 397 contains panel captures using both of these interfaces. The upper window uses System i Navigator from a Windows client workstation. The lower window shows the last page of the list of TCP/IP servers using Systems Director Navigator for i.

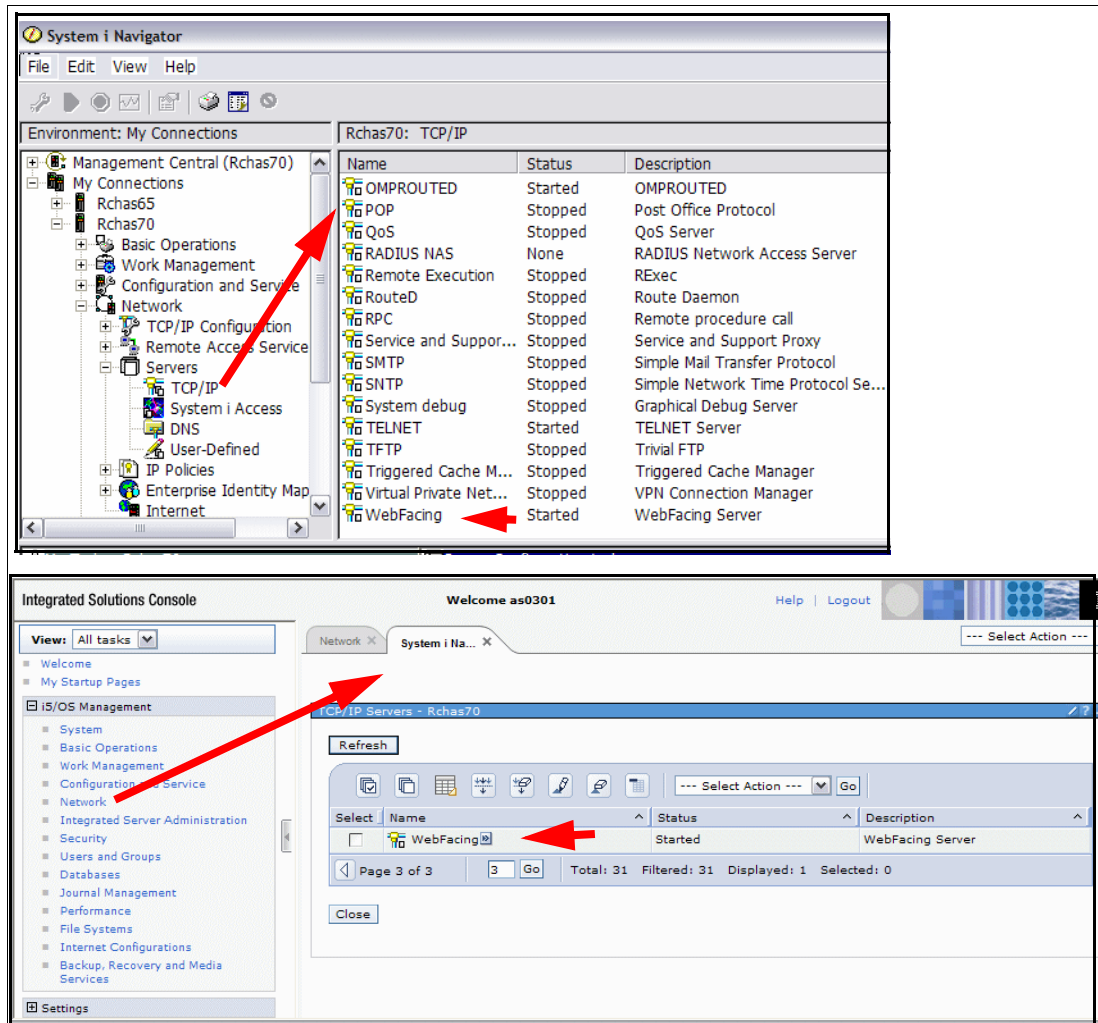


Figure 15-26 System i Navigator and Systems Director Navigator for i and WebFacing server

Note: Examine the system value QAUTOVRT. This value determines the number of VT jobs that can be auto started. The WebFacing runtime server relies on VT jobs. If the value is 0, no browser sessions can be launched. If necessary, set QAUTOVRT to *NOMAX or some value greater than 0. To change the value for QAUTOVRT, use the WRKSYSVAL system command to work with system values. WebFacing server connection jobs are still listed under the QINTER subsystem and marked as INT (interactive), but they are not counted as part of the interactive capacity of the server.

HATS interoperability with WebFacing applications

HATS interoperability with WebFacing applications provides the ability to perform data description specification (DDS) map transformation, as well as 5250 data stream transformation, in the same Web application. This function is particularly useful for those who have built a DDS source-based transformation application, using the WebFacing tool, that needs to interact with a host application that has been transformed using the HATS tool. This function is also useful for HATS users who want to take advantage of transformation, based on DDS maps, provided by the WebFacing tool.

Interoperability, in this context, means that a single HATS Web application interoperates with, or works alongside, a single WebFacing application. The combined Web application has the

capabilities of both tools, with some limitations, but the capabilities are cleanly separated within the Web application. To create a combined Web application, a single HATS Web project is linked with a single WebFacing project and packaged into a single enterprise application, or .ear file. This enterprise application contains the runtimes of both products, along with a common runtime component that is used to create and interact with a shared connection to the 5250 backend host.

For more information and a demonstration of this function, see the IBM Enterprise Modernization Demos at:

http://rational.dfw.ibm.com/atdemo/atdemo_wdht_recorded.html

Note: The use of this function requires that the HATS WebFacing feature is installed.

15.6.3 HATS V7.5 enhancements

The new features and enhancements in the latest version of HATS are:

Visual Macro Editor tool

The HATS Visual Macro Editor, or VME, is a tool for visually developing HATS macros. Macros are scripts that navigate through and interact with 3270, 5250, or VT host applications. Macros are used in HATS transformation applications and HATS Web services.

Prior to the VME, it was difficult for developers to see the relationships of the panels that made up the flow of the macro. The VME addresses this by visually showing the flow of the macro and enabling the developer to position panels in a visual canvas and create new flows. The VME combines many of the features of the HATS Macro Editor, Host Terminal, and Advanced Macro Editor into one tool, and simplifies offline development of macros.

It provides optional, automatic capture of panels as a macro is being recorded in the host terminal. It also allows flows to be copied between macros and provides drag-and-drop support for adding new panels. Although the VME is a replacement for the existing Macro Editor and Advanced Macro Editor, these editors are still available for use by a developer.

Features:

- ▶ Full visualization of the macro flow, including visualization of entry, exit, and transient panels
- ▶ Palette tools for adding new panels, adding actions, or wiring panels together
- ▶ Drag-and-drop new panels into the macro flow
- ▶ Associate panel captures with macro panels to facilitate configuring panel recognition and actions while offline
- ▶ Copy and paste panel flows between macros
- ▶ New *prompt all* and *extract all* actions to facilitate adding prompt or extract actions for every field on a panel
- ▶ Integrated host terminal, which provides for the easy addition of new panels and visualization of playback flow to enable a developer to see the path taken through the macro during development
- ▶ Visual wiring of existing macro panels (to define new paths in the macro)
- ▶ Copy and paste of the macro flow image for use in design documents, presentations, e-mails, and others

- ▶ Wizard to create new, blank macro (enables a developer to create a macro offline and assemble using existing panel captures)

Note: The VME was originally delivered as an unsupported technology preview in HATS V7.1. In version 7.5, the editor is fully supported and integrated into the HATS toolkit.

Support for latest runtimes

Support is now extended for deploying HATS applications to the following runtimes: IBM WebSphere Application Server V7.0, IBM WebSphere Portal V6.1, and Eclipse Rich Client Platform V3.4. HATS application continues to run under WebSphere Application Server v5.x and v6.x.

Support for latest development environments

Support is now provided for Eclipse 3.4-based IBM Rational development environments, including IBM Rational Application Developer V7.5.

Note: The HATS V7.5 toolkit will not install on V7.1-level Rational development environments, including IBM Rational Application Developer V7.1, IBM Rational Developer for System z V7.1, or IBM Rational Developer for i for SOA Construction V7.1.

Firefox support for WebFacing applications

WebFacing applications can now be accessed by users of the Firefox Web browser. Prior to this enhancement, applications could only be accessed using Microsoft Internet Explorer, thus limiting usage to clients running Microsoft Windows.

Coexistence

HATS V6, HATS V7, and HATS V7.1 can coexist on the same development machine with HATS V7.5. It is also possible for HATS V7.5 to be installed multiple times on the same system. If you have multiple IBM Rational Software Delivery Platform offerings installed on your system, HATS V7.5 can be installed into each of them.

Migration

HATS V4, HATS V5, HATS V6, HATS V7, HATS V7.1, HATS V4 LE, HATS V5 LE, or linked HATS/WebFacing projects can be migrated to HATS V7.5. If you are a Host Publisher V4.0 or V4.01 user, you can migrate entire applications, or just the Integration Objects. Migration is supported from V5.1.2 or later for WebFacing projects (both Web and portlet).

Support for visually editing HATS rich client transformations and templates is no longer available in the HATS Toolkit. Rich client applications can still be created and developed using HATS V7.5, but transformations and templates must be developed using the Java source editor provided by the IBM Software Delivery Platform. HATS transformations and templates can still be developed visually using a prior version of HATS, and then imported into HATS V7.5.

15.6.4 WebFacing V7.1 enhancements

Starting with Version 7.1, the IBM WebFacing Tool is installed as a part of Host Access Transformation Services (HATS). Following are new features in WebFacing:

Application bridge

The application bridge is of great value when enhancing your IBM i application by allowing you to implement part of it as a Web application such as Enterprise Generation Language (EGL), and modernizing other parts with WebFacing or HATS. It allows WebFacing to transfer control and data to and from other Web applications such as EGL, giving you the ability to combine HATS, WebFacing, and EGL into a single application.

For more information about Application bridge, refer to HATS 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/hatshelp/v71/index.jsp?topic=/com.ibm.hats.doc/doc/topics/twfaplinkage.html>

IBM i integrated Web application server

WebFacing applications can be deployed to the IBM i integrated Web application server. You can use the IBM i integrated Web application server if you are deploying WebFacing Web applications and have limited IBM i system resources that cannot optimally run the WebSphere Application Server.

For information on the IBM i integrated Web application server on 6.1 in this publication, see 12.6, “Integrated Web application server” on page 280.

Note: This configuration does not support high availability and failover.

15.7 Rational Team Concert for i

IBM Rational Team Concert for i (RTCi) extends recently announced Rational Team Concert to provide a collaborative software delivery environment that empowers IBM i development teams to simplify, automate, and govern software delivery. Dynamic project provisioning enables day one productivity, while real-time collaboration helps significantly reduce scrap and rework. Rational Team Concert extends the capabilities of the Team with integrated work item, build, software configuration management (SCM) and the collaborative infrastructure of the Jazz Team Server.

RTCi provides the following capabilities:

- ▶ Source control
- ▶ Change management and build
- ▶ Process management and governance
- ▶ Integration with Rational Developer for i (RDi)
- ▶ Specialized support for source control, change management, and builds of traditional language artifacts such as RPG and COBOL
- ▶ Support for multi-tier software development and application modernization efforts using RPG, Java, and EGL
- ▶ Support for IBM i native library file system and the root file system
- ▶ IBM i artifact builds and Native hosting on IBM i

15.8 WebSphere Development Studio Client for System i

WebSphere Development Studio Client for System i, as well as the non-ILE System i compilers and Application Development ToolSet components in WebSphere Development Studio for System i are stabilized. That means IBM will no longer add any enhancements to these components.

IBM future investment will focus on our strategic ILE compilers and IBM Rational application development tools. IBM recommends that customers use the ILE set of compilers. For workstation tools, IBM recommends that customers use the IBM Rational Developer for i product or the Rational Developer for i for SOA Construction.



IBM Systems Director Navigator for i

This chapter discusses the following topics:

- ▶ “IBM Systems Director Navigator for i” on page 404
- ▶ “Systems Director Navigator for i and Systems i Navigator” on page 405
- ▶ “Working with IBM Systems Director Navigator for i” on page 407

Note: PTFs SI35663, SI36093, and SI35833 will add significant new function to the Systems Director Navigator performance tasks. For detailed information about these functions, see:

End to end Performance Management on IBM i, SG24-7808

16.1 IBM Systems Director Navigator for i

IBM Systems Director Navigator for i is a Web console interface for IBM i administration, where you can work with the Web-enabled tasks of IBM i Navigator provided by IBM i. IBM Systems Director Navigator for i includes a number of welcome pages that allow you to quickly find the task that you want to perform.

Web-enabled tasks of IBM i Navigator (initially available in IBM i 5.4) offer a browser-based interface to a set of URLs that interface to a set of IBM i functions generally similar to those available using non-Management Central functions of the Windows client workstation System i Navigator. The IBM Systems Director Navigator for i is not just another interface to this set of URL addressable tasks. It is a robust Web console from which you can manage your IBM i partition.

The IBM Systems Director Navigator for i interface—the Web page layout, colors, point and click navigation—is the same as that used for console functions for other IBM products, such as the administration console for a WebSphere Application Server server and instance. This interface follows the IBM Integrated Solution Console architecture used by many IBM products.

For example, the initial log-on window looks the same for signing on to either IBM i using the IBM Systems Director Navigator for i or the WebSphere Administration Server administration console.

The IBM Systems Director Navigator for i interface is also very similar to that used by the suite of IBM Systems Director V6.1 products delivered 4Q 2008, which include:

- ▶ IBM Systems Director for Power Systems, V6.1 (AIX, IBM i, Windows, Linux)
- ▶ IBM Systems Director for Linux on System z, 6.1
- ▶ IBM Systems Director for x86, V6.1

This suite of IBM Systems Director V6.1 products allows you to effectively manage from a central location, discovering, visualizing, monitoring, and managing both physical and virtual systems in a heterogeneous environment.

However, each IBM Systems Director Navigator for i session works with only one IBM i partition.

This chapter focuses on the IBM Systems Director Navigator for i that is integrated into IBM i 6.1 and separate from the heterogeneous system support under the set of IBM Systems Director for multisystem previously listed.

The IBM Systems Director Navigator for i console contains the following task categories for IBM i management. As previously stated, these categories are similar to those that you see in the System i Navigator Windows client application.

- ▶ System (System status, System operator messages, History log, Disk status, Change password, and more)
- ▶ Basic operations (Messages, Printers, Printer output, and more)
- ▶ Work management (Active jobs, Server jobs, Output queues, Active subsystems, and more)
- ▶ Configuration and service (System values, Time management, Disk units, Disk pools, and more)
- ▶ Network (TCP/IP servers, Lines, DNS servers, and more)

- ▶ Integrated server administration (All virtual disks, Remote systems, Domains, and more)
- ▶ Security (Authorization lists, Cryptographic services key management, Intrusion detection, and more)
- ▶ Users and groups (Users, Create users, Groups, and more)
- ▶ Databases (Performance monitors, Health center, Databases, and more)
- ▶ Journal management (Journals, Create a journal, and more)
- ▶ Performance (Investigate data and collections, and more)
- ▶ File systems (Integrated file system, File shares, Create file share, and more)
- ▶ Internet configurations (Digital certificate manager, System i Navigator Tasks on the Web, and more)
- ▶ Cluster resource services (Display the list of nodes, Delete the cluster, and more)

The IBM Systems Director Navigator for i is a new topic collection for IBM i 6.1.

The key points for IBM Systems Director Navigator for i are:

- ▶ It is used to manage only the system in which it is installed.
- ▶ It provides all the current Web-enabled functions of Navigator for i (over 300 tasks).
All Web-enabled tasks are still URL addressable and can be accessed individually as servlets, or as portlets within the new Web console.
- ▶ It is included as part of the base operating system option 3 (5761SS1 0003).

Table 16-1 provides the list of IBM i products that need to be installed on the system before using Systems Director Navigator.

Table 16-1 Required IBM i products for IBM Systems Director Navigator for i to operate

Description	Licensed Program	Product Option
IBM HTTP Server for i5/OS	5761DG1	
Qshell	5761SS1	option 30
Portable App Solutions Environment	5761SS1	option 33
IBM Developer Kit for Java	5761JV1	
J2SE 5.0 32 bit	5761JV1	option 8
IBM Toolbox for Java	5761JC1	
Host Servers	5761SS1	option 12

16.2 Systems Director Navigator for i and Systems i Navigator

IBM Systems Director Navigator for i contains most of the tasks currently available in System i Navigator Windows client application. Systems Director Navigator for i offers these System i Navigator capabilities through a Web interface with no additional need to install the software on your computer.

In this section we highlight the main differences between the two products.

16.2.1 Features included in Systems Director Navigator for i

The following are System i Navigator features that are included in Systems Director Navigator for i:

- ▶ All Web-enabled functions of Navigator for i, organized into 13 different task categories, as shown in Figure 16-1, with over 300 tasks total.

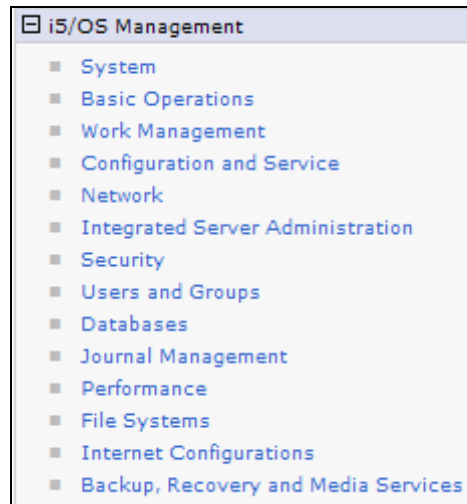


Figure 16-1 Task categories

- ▶ Functions from the following LPPs are available in the links listed in Figure 16-1 if they are installed:
 - PowerHA for i (System i High Availability Solutions Manager) 5761-HAS
 - High Availability Solutions Manager (solution-based) GUI
 - Cluster Resource Services (task-based) GUI
 - BRMS (Backup, Recovery, & Media Services) 5761BR1
 - IBM Performance Tools for i5/OS (DiskWatcher) 5761PT1 option 1
 - IBM Performance Tools for i5/OS (JobWatcher) 5761PT1 option 3

16.2.2 Features not included in Systems Director Navigator for i

System i Navigator features that are not included in Systems Director Navigator for i are:

- ▶ Management Central functions
- ▶ Use IBM Systems Director for multi-system tasks on the Web
- ▶ Database functions that involve graphics and charts
 - SQL Scripts
 - Visual Explain
 - Database Navigator
 - SQL Assist
- ▶ Configuration and Service links to hardware-related functions are restricted to disk configuration and management (includes disk pool (ASP and iASP) creation and management, disk add, remove, replace, and RAID parity management). The following

System i Navigator client disk-related functions that use graphics and charts are not available:

- Graphical view
 - Disk unit physical location
 - Disk unit/ASP capacity chart
 - Tower physical location.
- ▶ AFP Manager
- Resources
 - PSF configurations
 - Font Mapping tables

16.3 Working with IBM Systems Director Navigator for i

This section starts with how to connect your browser to the IBM i 6.1 or later partition using the Systems Director Navigator for i and discusses general navigation among the categorized i5/OS management tasks shown on the left pane.

Then we describe the basic capabilities under each high-level tasks category.

Important: The HTTP Server Administration server (referenced on various user interfaces as *ADMIN or ADMIN) must be running in order for the IBM Systems Director Navigator for i console to be available.

To work with IBM Systems Director Navigator for i, follow these steps:

1. Access the following URL from a Web browser where hostA is your IBM i partition TCP/IP host name:

`http://hostA:2001`

In our examples we connect to either RCHAS70 or RCHAS65.

If you get a window asking whether you will use the security certificate used by IBM i Systems Director Navigator, respond with a “Yes” and continue. The initial log-on window is then displayed, as shown in Figure 16-2. Enter your IBM i user profile information and click **Log in** or press the Enter key.

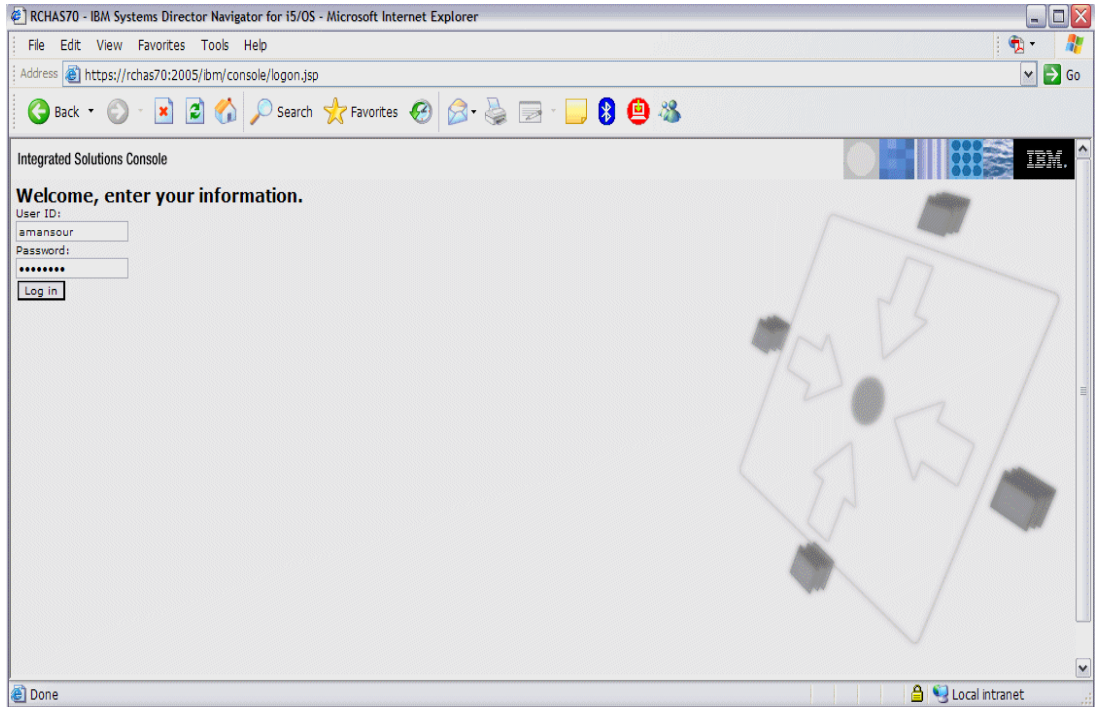


Figure 16-2 Systems Director Navigator Welcome panel

Note: IBM i 6.1 reroutes the URL <http://IBMiSystem:2001> to the secure session as indicated in:

<https://IBMiSystem:2005/ibm/console/logon.jsp>

2. To display a list of IBM i tasks from the IBM Systems Director Navigator for i console, expand **i5/OS Management**, as shown in Figure 16-3.

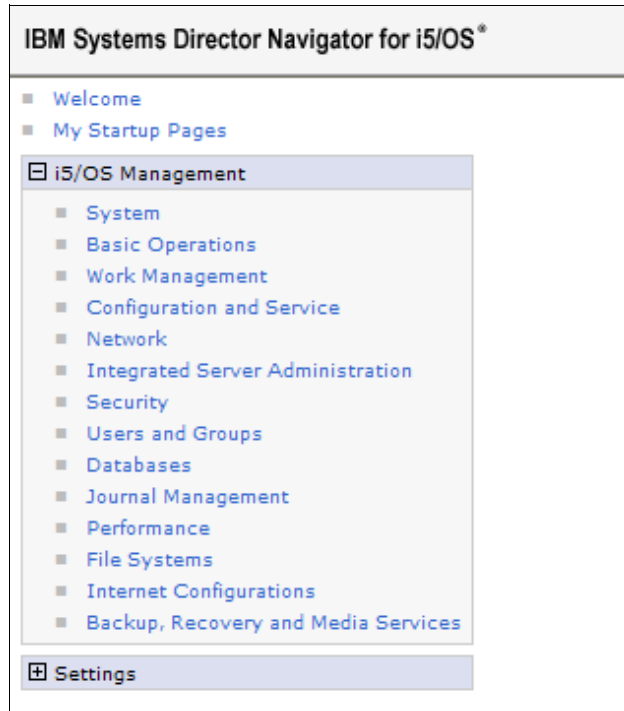


Figure 16-3 IBM i Management in Systems Director

Note: Each of the task categories has a welcome page that has links to the most common tasks. From that welcome window you can optionally click a button that shows all possible tasks.

3. Click **Show All Tasks** at the bottom of the welcome page to see a list of all the available tasks within that category.
4. Click the task you want to perform, and a new tab opens in the console.

- Use the tabs to navigate between open pages or to close a specific page. If you open more than five pages, you can scroll to access other tabs. We used the window showing the most common Network tasks in our example in Figure 16-4.

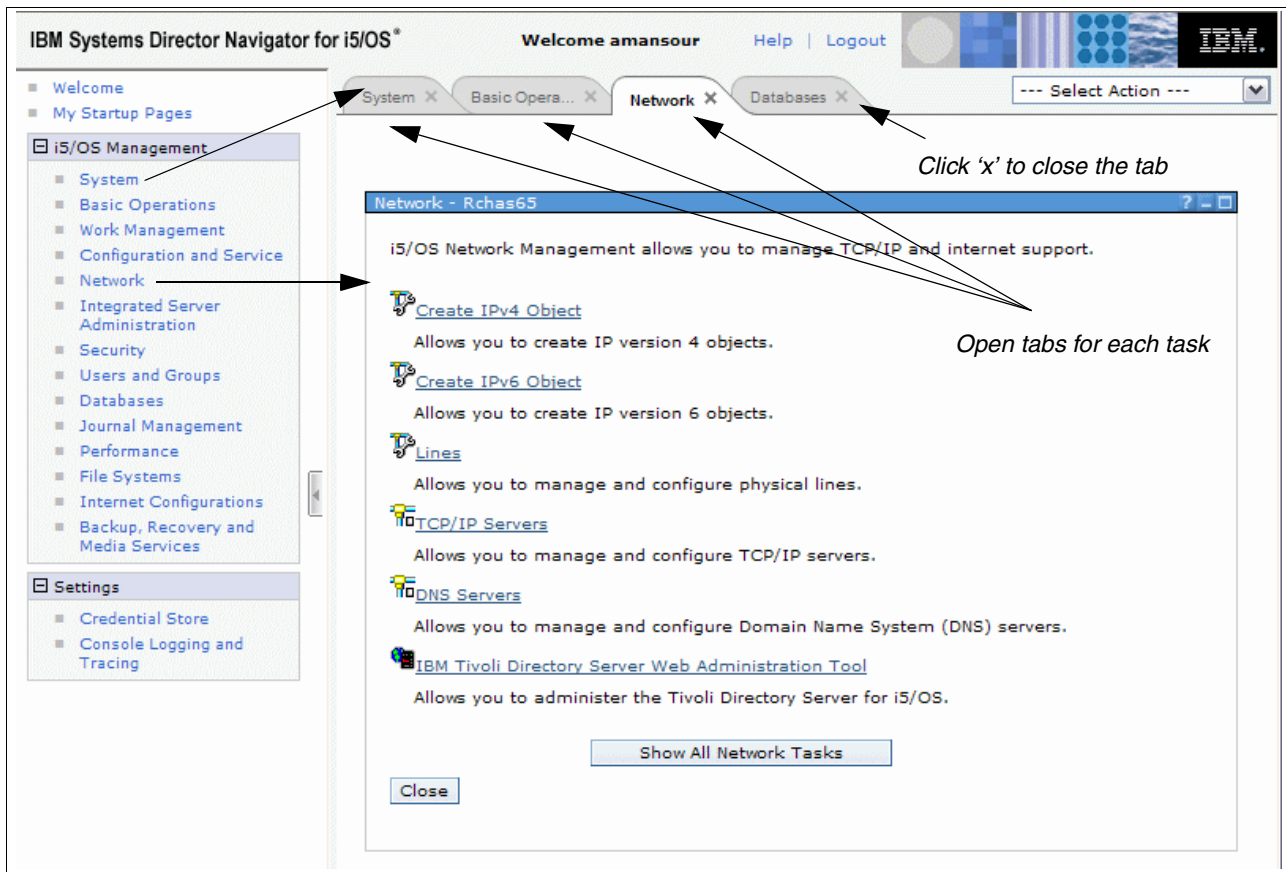


Figure 16-4 Navigation in Systems Director

The Select Action menu, shown in Figure 16-5 on page 411, provides the following actions that you can perform:

- ▶ **Close Page**
You can close the current page.
- ▶ **Add to My Startup Pages**
You can add the current page to the list of pages that are automatically displayed when you log into the console.
- ▶ **Manage Open Pages**
You can close one or more open pages and navigate to a page.
You can manage a large number of open pages where you can close some open pages and work with the other remaining open pages.

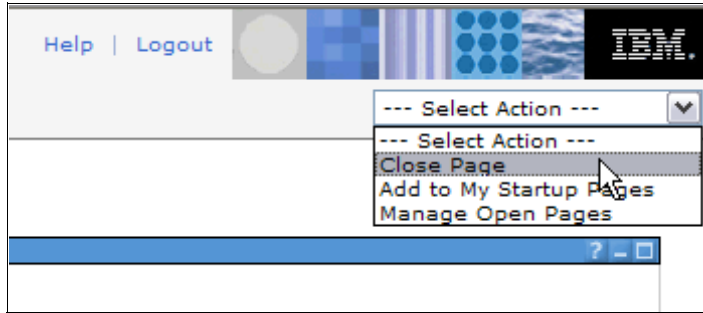


Figure 16-5 Select Action menu

My Startup Pages can be used to add or remove pages from the list of pages that are displayed every time you log into the console. You can select which page is opened as the default page, as shown in Figure 16-6. To open one of your startup pages, click the link for that page in the Page name column.

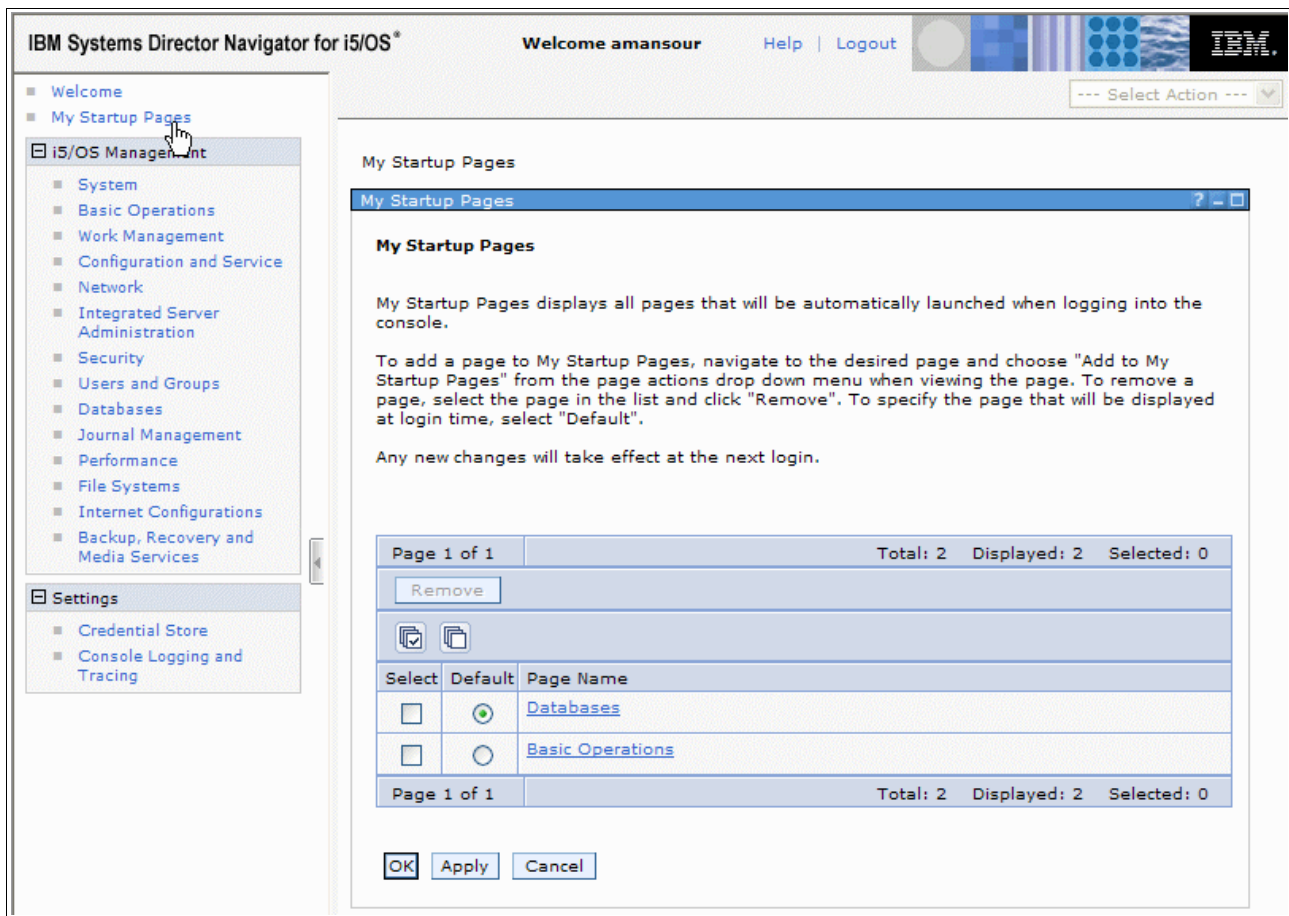


Figure 16-6 My Startup Pages

Note: If you want to use the URL Wizard to create custom links to System i Navigator tasks, you can access the System i Navigator tasks on the Web home page that was available in previous releases. You can click **i5/OS Tasks** from the IBM Systems Director Navigator for i console welcome page and select System i Navigator tasks on the Web.

16.3.1 i5/OS Tasks

In previous IBM i versions you could display the i5/OS Tasks page using:

<http://{IBMiSystem}:2001>

This URL will now open the Systems Director Navigator for IBM i. The Web-based tasks that used to appear on that page can now be accessed by clicking **i5/OS Tasks Page**, which appears on the IBM Systems Director for i Welcome Page, as shown in Figure 16-7.



Figure 16-7 i5/OS Tasks Page

16.3.2 Exploring IBM Systems Director Navigator for i

In this section we introduce some of the features available in Systems Director Navigator for i:

- ▶ System Page
 - Application Administration
 - System Status
- ▶ Basic Operations
- ▶ Work Management
- ▶ Configuration and Service
- ▶ Network
- ▶ Integrated Server Administration
- ▶ Security
- ▶ Users and Groups Menu
- ▶ Databases
- ▶ Journal Management
- ▶ Performance
- ▶ File systems
- ▶ Internet Configuration

System page

The welcome page provides links to the most commonly used System tasks. To view a complete list of available tasks, click **Show all System Tasks**, as shown in Figure 16-8.

System operations include viewing system status, messages, disk status, and passwords.

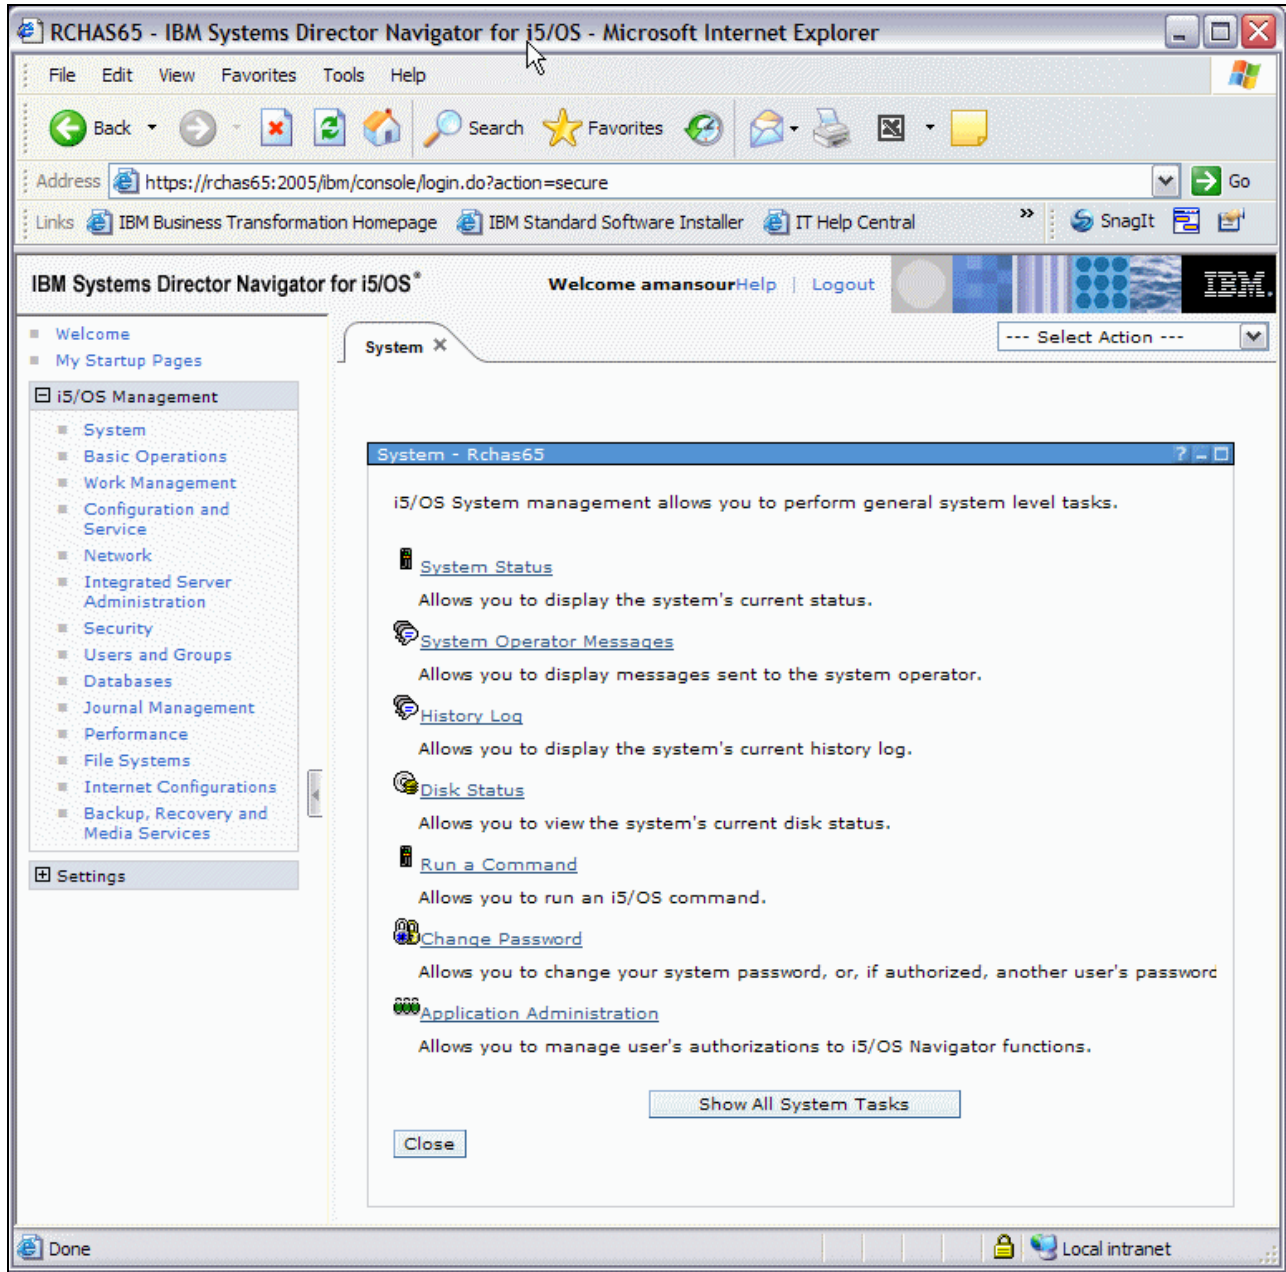


Figure 16-8 System page

Application Administration is a central system that is used to manage many of the properties of the system. It allows administrators to configure the system and act as an administration system. Typically, a network has only one system acting as an administration system.

This set of functions corresponds to the available functions in System i Navigator for Windows for application administration.

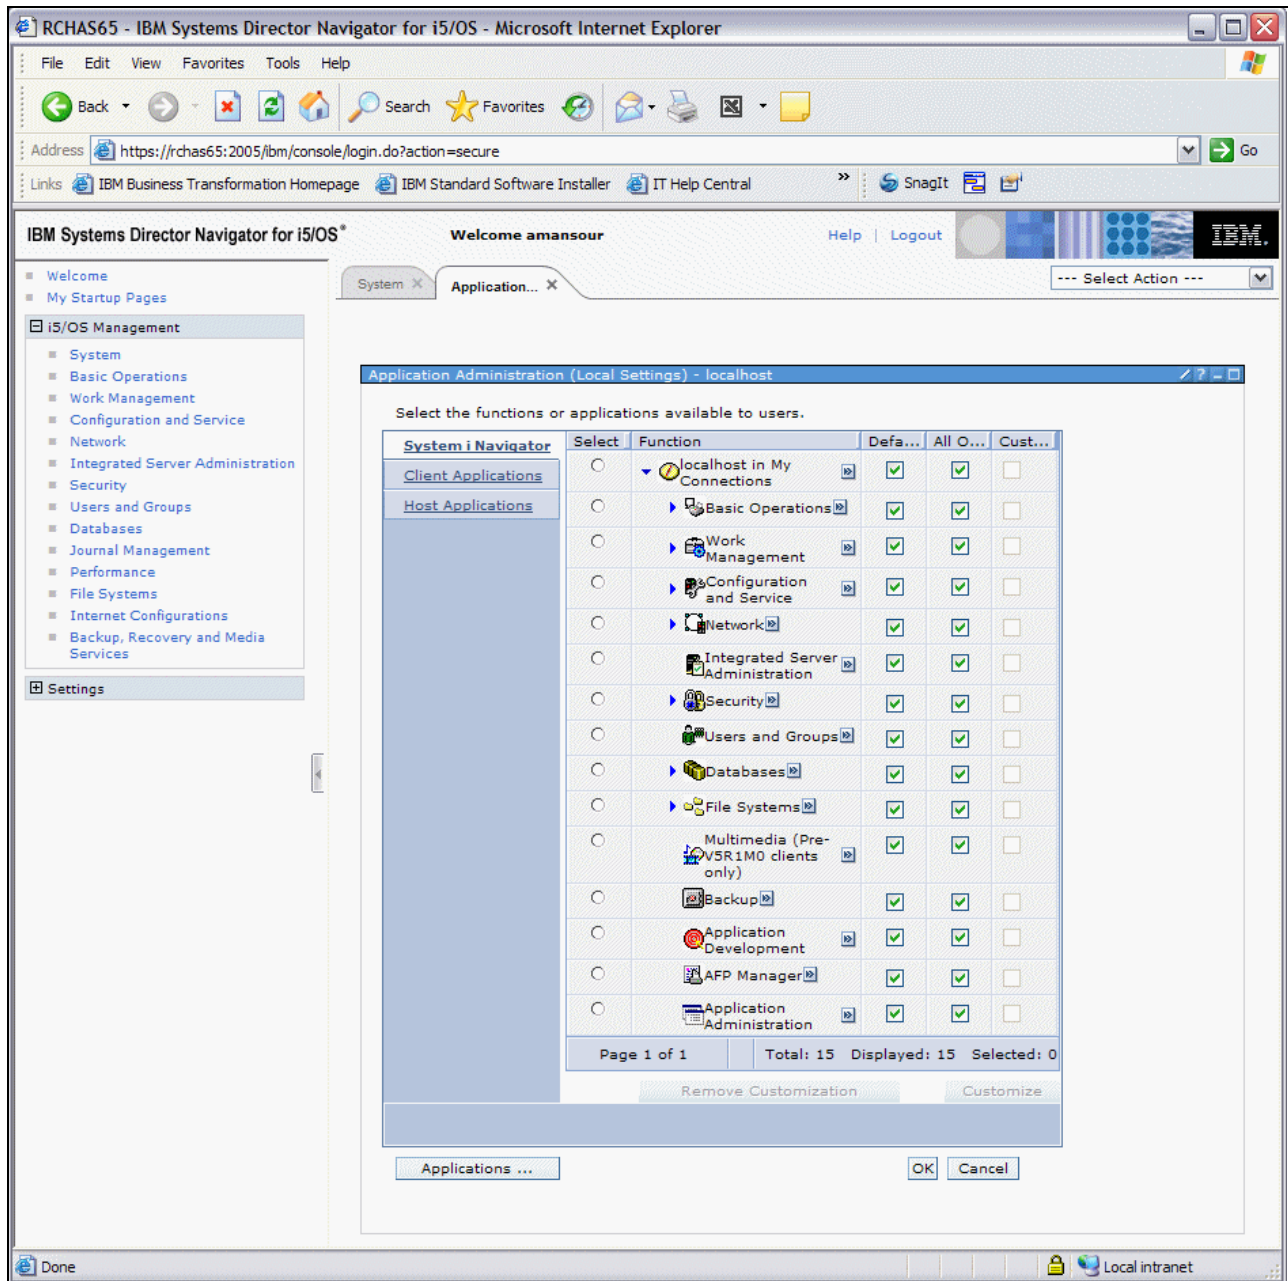


Figure 16-9 Application Administration

Note: You must have security administrator (*SECADM) and all object (*ALLOBJ) special authorities to work with advanced settings on an administration system. This differs from other settings in Application Administration, which only require security administrator (*SECADM) special authority to make changes.

System Status

System Status task allows you to view general system status, current jobs, processors status, memory status, disk space usage, and status of used addresses, as shown in Figure 16-10 on page 415.

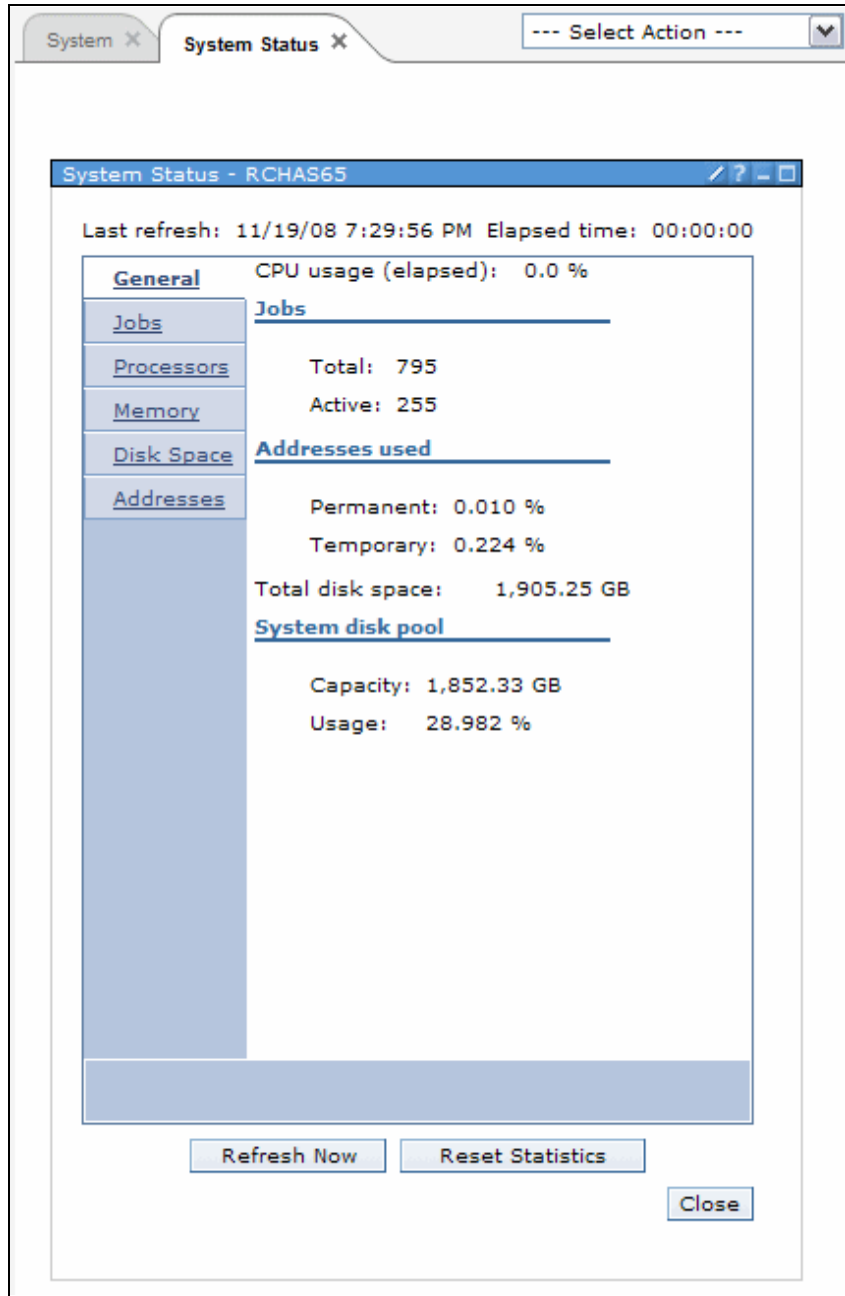


Figure 16-10 System Status

Basic Operations

The welcome page of Basic Operations shown in Figure 16-11 on page 416 provides links to the most commonly used Basic Operation tasks. To view a complete list of available tasks, click **Show all Basic Operation Tasks**.

Basic Operations allow you to easily manage lists of messages, printer output, and printers.

From IBM Systems Director Navigator for i, you can display a list of messages. You can view all messages, messages for the system operator, or QSYSMSG messages. You can also specify which columns of information you want to display in the list and in what order you want the columns to be displayed.

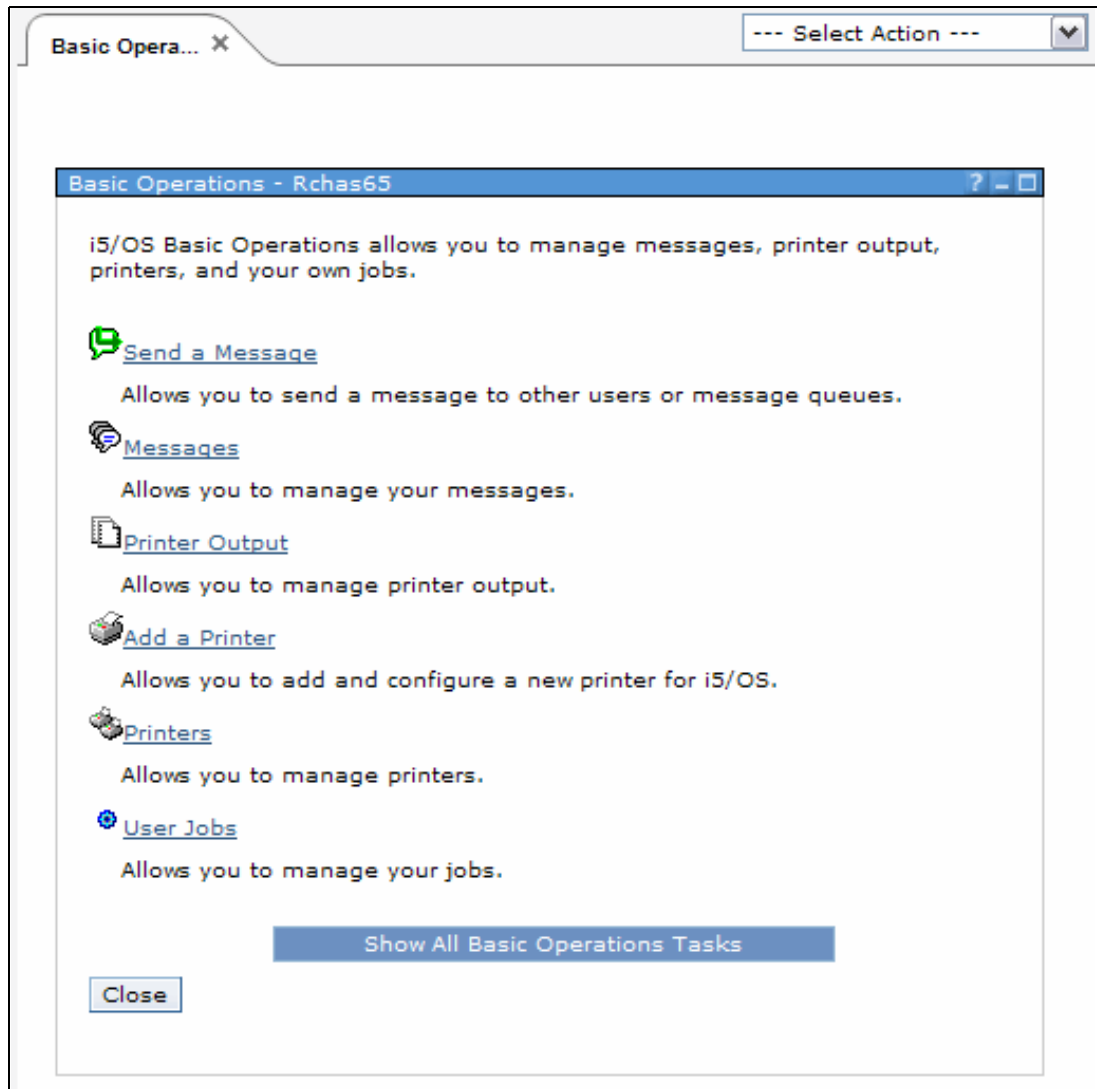


Figure 16-11 Basic Operations

Work management

The work management welcome page provides links to the most commonly used Work Management tasks, as shown in Figure 16-12 on page 417. To view a complete list of available tasks, click **Show all Work Management Tasks**.

Work management functions control the work performed on the system. The work management environment supports all interactive and batch work. In addition, work management contains the functions you need to distribute resources for your applications so that your system can handle your applications.

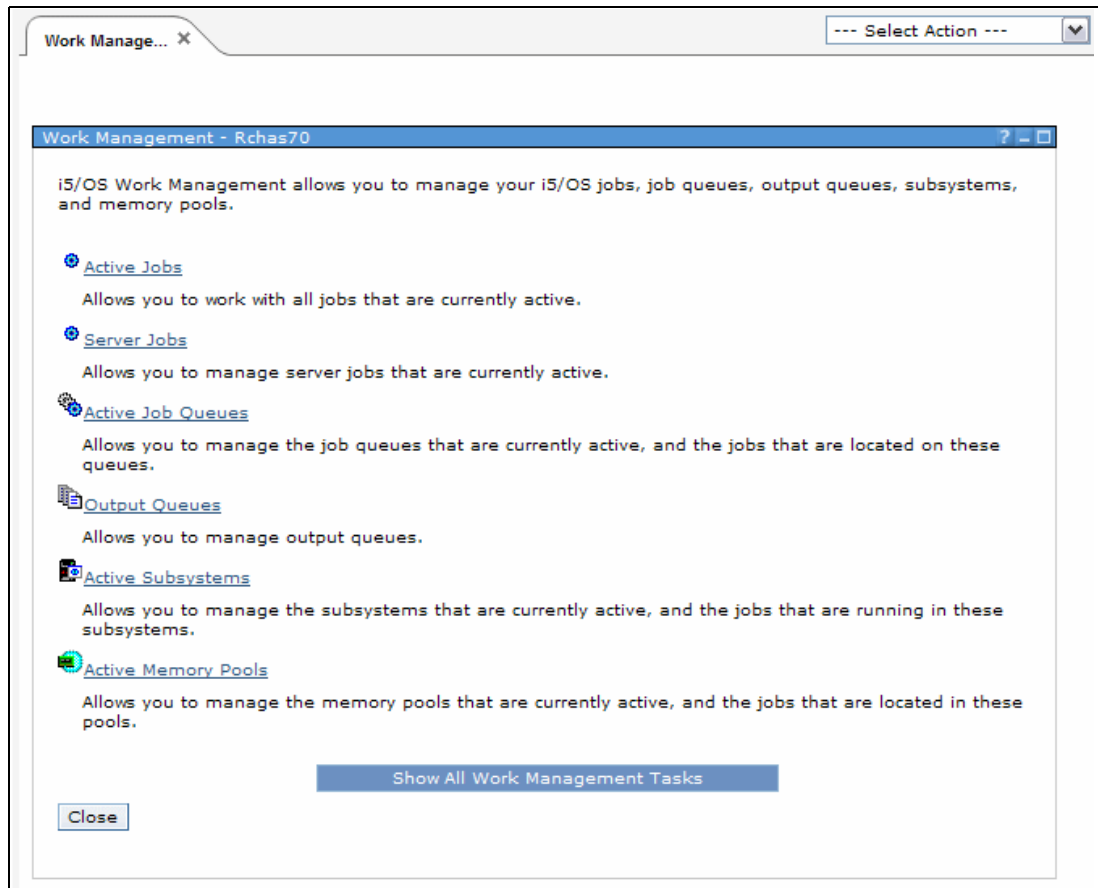


Figure 16-12 Work Management

Configuration and Service

The Configuration and Service welcome page, shown in Figure 16-13 on page 418, provides links to the most commonly used Configuration and Service tasks. To view a complete list of available tasks, click **Show all Configuration and Service Tasks**.

Configuration and Service provides functions to help you manage both hardware and software on your system. You can view and change the system values, and use disk management functions to view and manipulate large disk configurations.

You cannot view other configured hardware, such as processor and memory features, LAN adapters, tape, or optical devices.

The disk functions—viewing the disk configuration, adding and removing disks, creating and managing disk pools (independent ASPs and dependent ASPs)—are powerful but require special authorization to IBM I System Service Tools.

In order to access any of the IBM Systems Director Navigator for i disk functions, the IBM i user must have signed on with a user profile and password that is also defined to the System Service Tools component of IBM i. While the IBM password may be entered in either uppercase or lowercase, the corresponding SST password must have been specified in uppercase.

For example, the IBM i user as0311 signed successfully entering either as0311pwd or AS0311PWD. Selecting one of the disk functions successfully starts if the SST password is

specified as AS0311PWD. For more information in this publication, see 2.8, “Service Tools user profile and IBM i profile” on page 21.

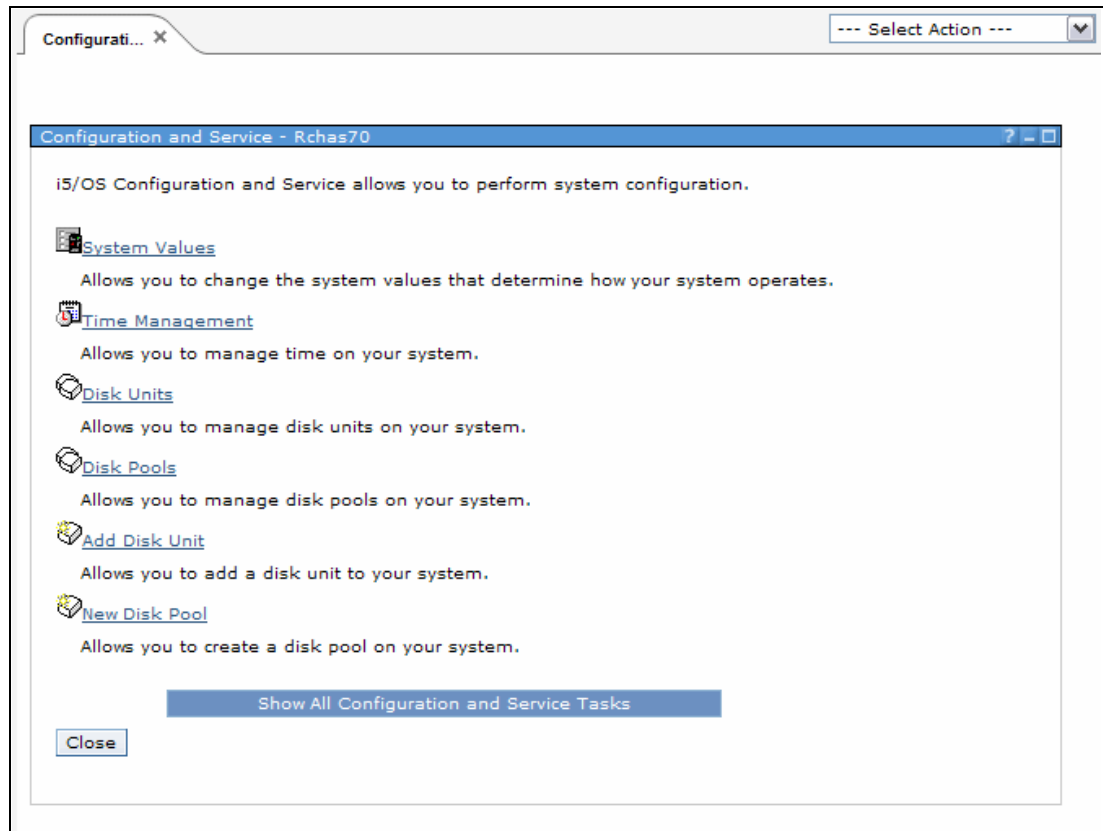


Figure 16-13 Configuration and Service

Network

The Network welcome page, shown in Figure 16-14 on page 419, provides links to the most commonly used Network tasks. To view a complete list of available tasks, click **Show all Network Tasks**.

Your understanding of networking technologies is a vital part of your company's total e-business solution. Using IBM Systems Director Navigator for i, you can work with TCP/IP configuration, remote access services, networking servers, and IP Policies.

You can use TCP/IP Configuration to manage and monitor your TCP/IP network, using TCP/IP Configuration wizards for configuring both IPv4 and IPv6 lines, interfaces, and routes. The Network page also allows you to work with Domain Name System (DNS) servers.

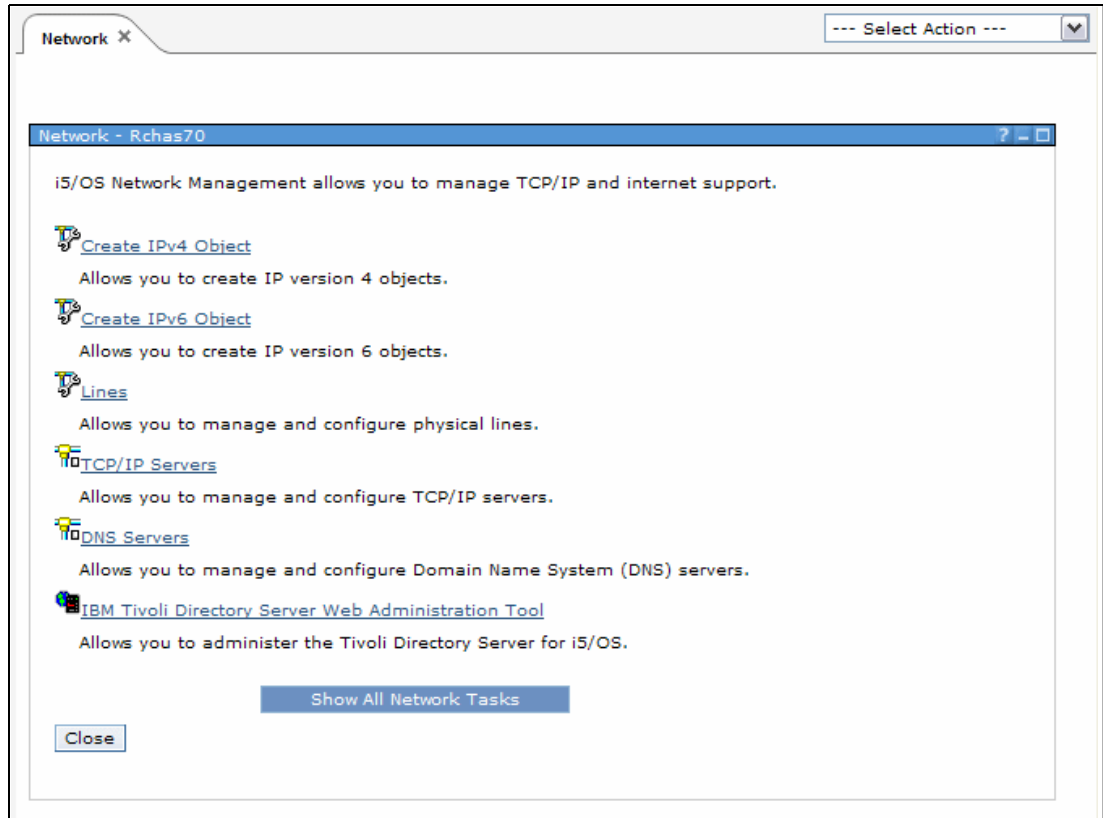


Figure 16-14 Network welcome page

There is a wide range of Network tasks under IBM Systems Director Navigator for i Show All Network Tasks. A complete set of these tasks is shown in 8.1.5, “Systems Director Navigator for IBM i Network interfaces” on page 220.

Examples of using one of these tasks are shown in 2.1, “Intrusion detection and prevention” on page 6.

Integrated Server Administration

The Integrated Server Administration welcome page, shown in Figure 16-15, provides links to the most commonly used Integrated Server Administration tasks. To view a complete list of available tasks, click **Show all Integrated Server Administration Tasks**.

IBM i Integrated Server Support combines the ease-of-use of a Windows, Linux, or VMware server with the dependability and power of i, giving you the flexibility to enroll and manage users and groups, greater accessibility and protection for your data, and simplified administration.

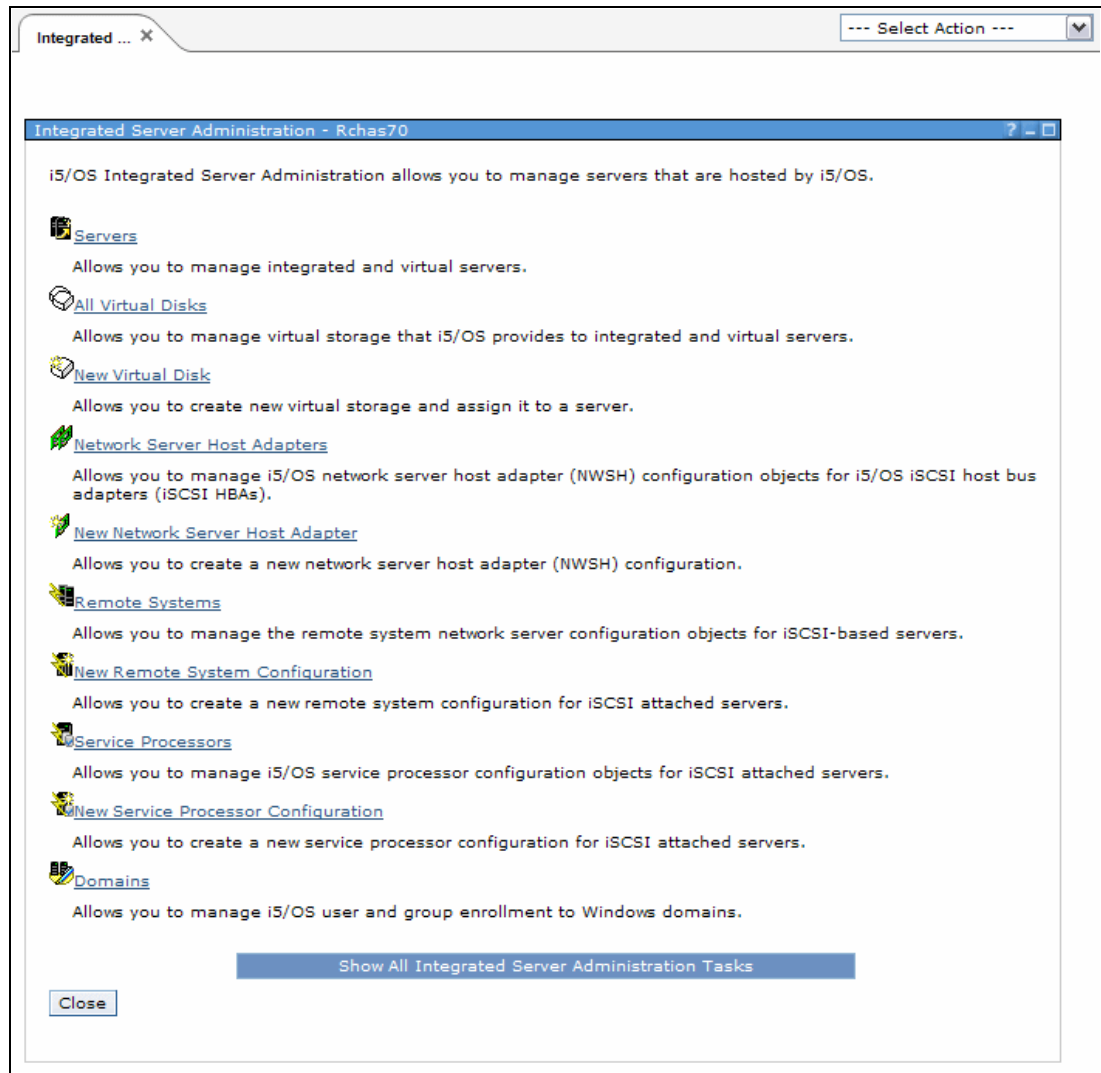


Figure 16-15 Integrated Server Administration

Security

The Security welcome page, shown in Figure 16-16 on page 421, provides links to the most commonly used Security tasks. To view a complete list of available tasks, click **Show all Security Tasks**.

Security functions allow you to plan and implement security on your system by performing functions related to:

- ▶ Authorization lists
- ▶ Object permissions
- ▶ Cryptographic Services Key Management
- ▶ Intrusion detection system

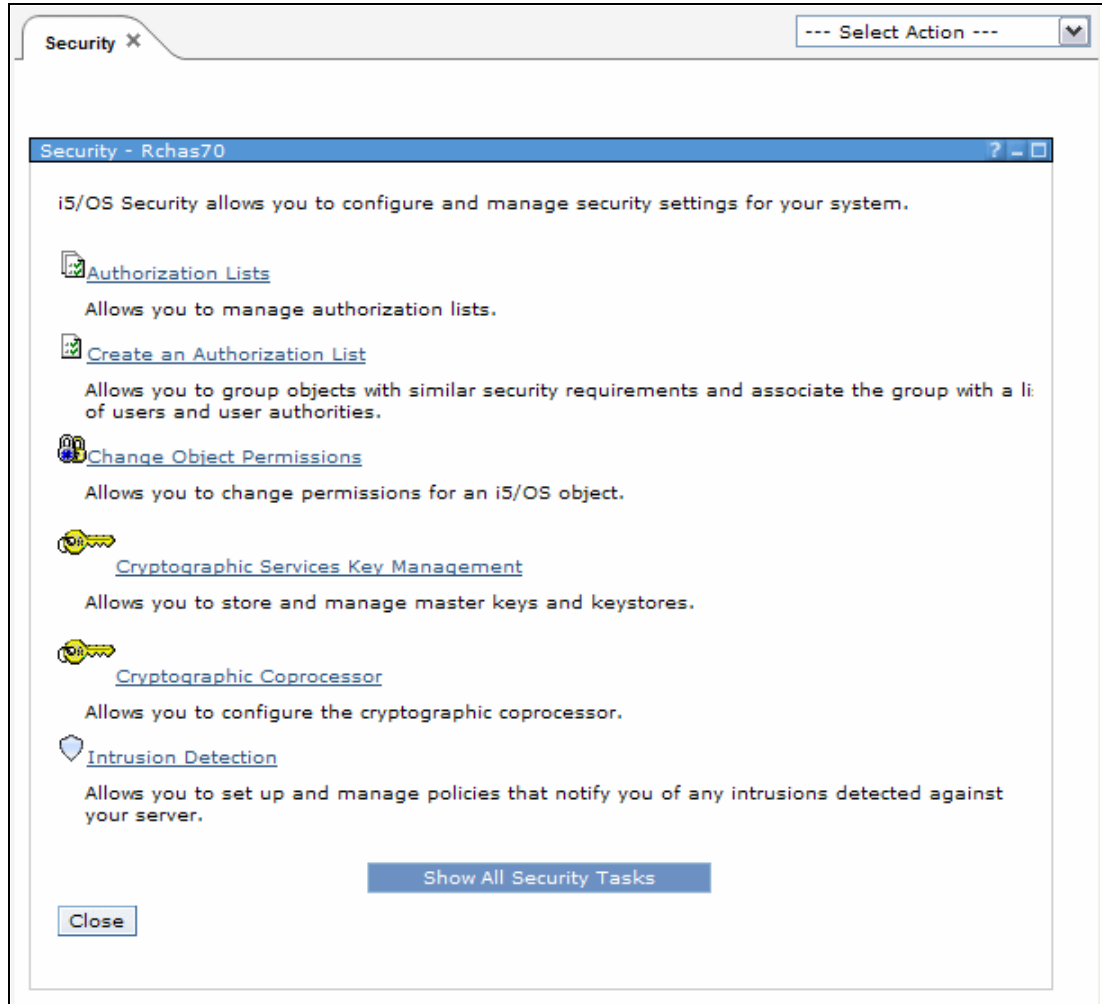


Figure 16-16 Security welcome page

Users and Groups

The Users and Groups welcome page, shown in Figure 16-17, provides links to the most commonly used Users and Groups tasks. To view a complete list of available tasks, click **Show all Users and Groups Tasks**.

Users and Groups displays a complete list of the users and groups configured on the system. You can use this functional area to manage users and groups.

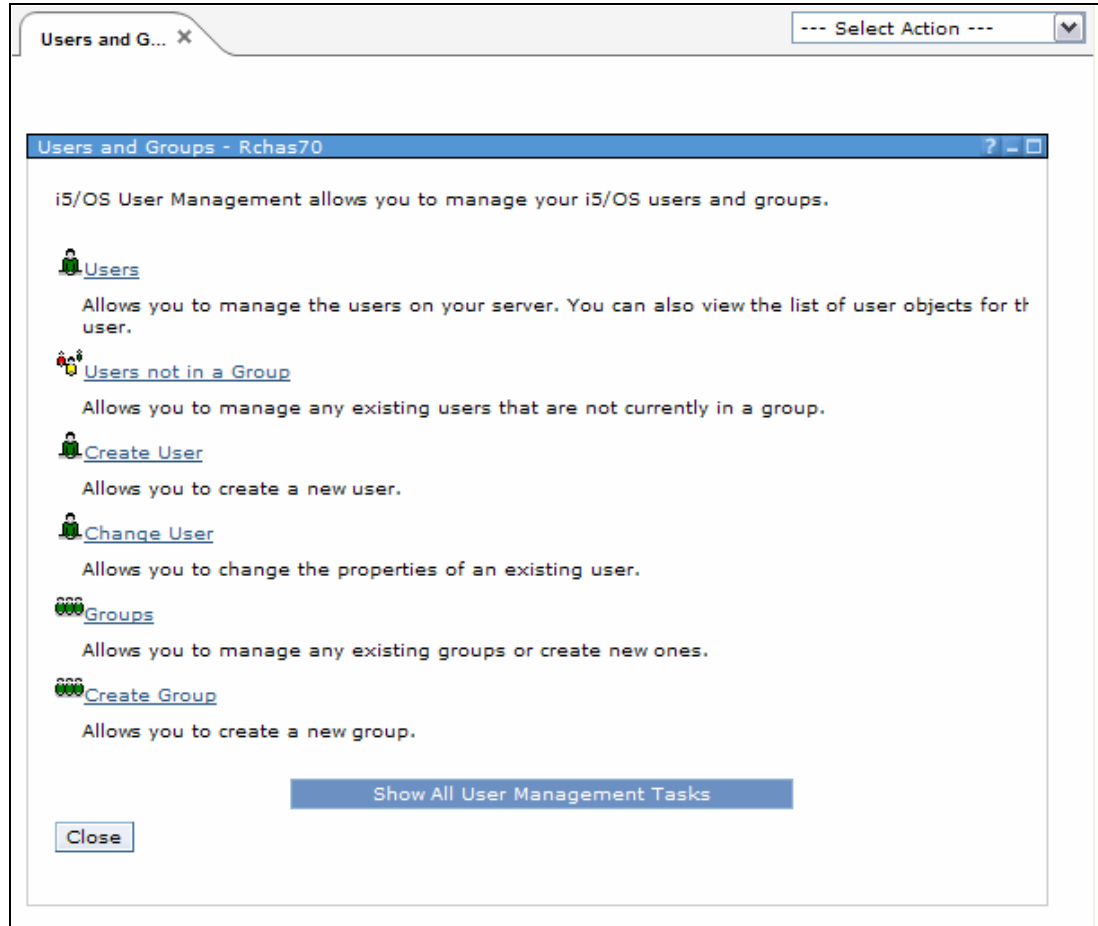


Figure 16-17 Users and Groups welcome page

Databases

The Databases welcome page, shown in Figure 16-18, provides links to the most commonly used database tasks. To view a complete list of available tasks, click **Show all Database Tasks**.

Database is a graphical interface that you can use to perform many of your common administrative database operations. Most operations are based on Structured Query Language (SQL), but you do not need to fully understand SQL to perform them.

Using the database function of IBM Systems Directory, you can access and modify objects, work with performance monitors, as well as access the Health center.

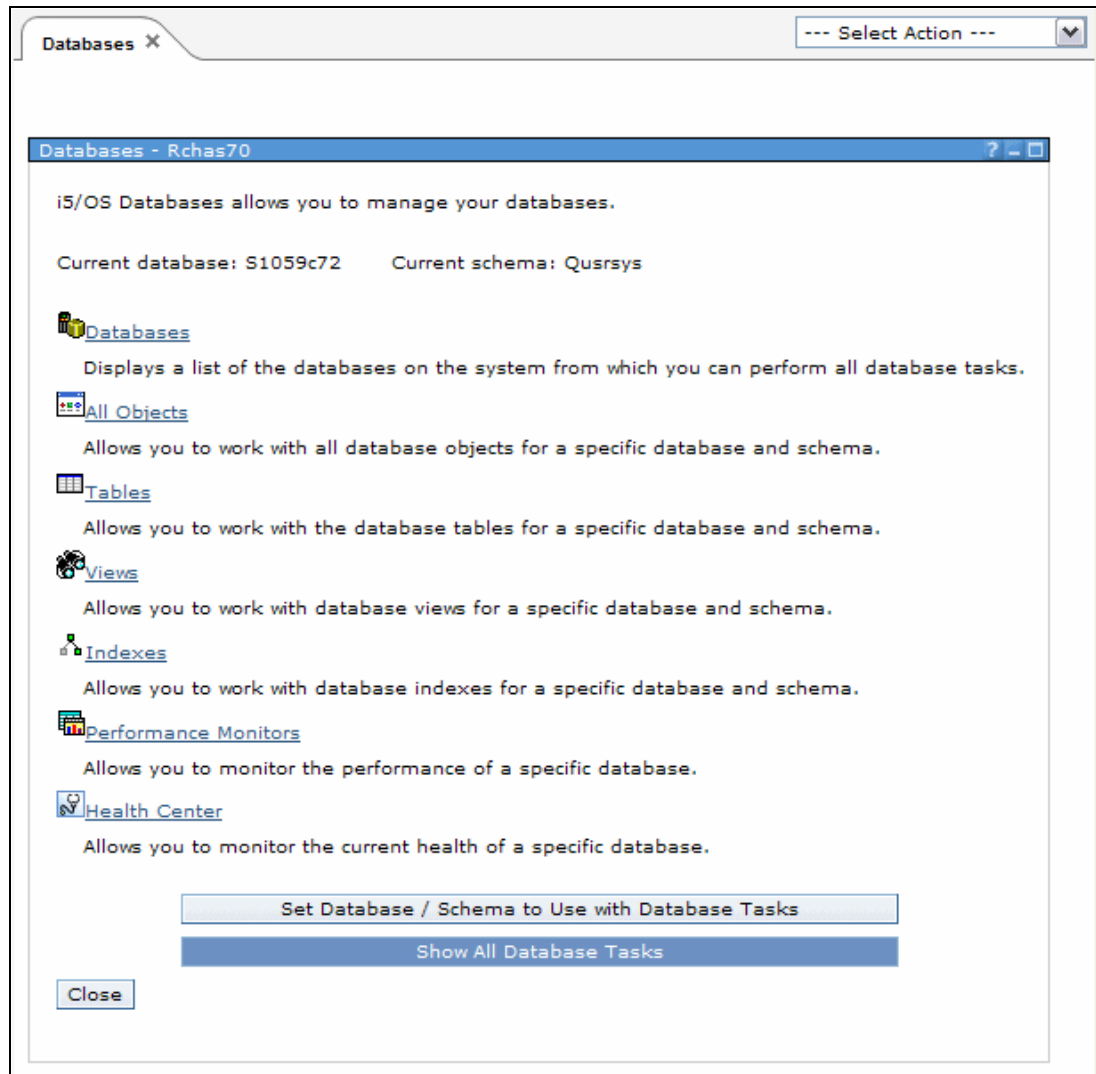


Figure 16-18 Databases welcome page

Journal Management

The Journal Management welcome page, shown in Figure 16-19 on page 424, provides links to the most commonly used Journal Management tasks. To view a complete list of available tasks, click **Show all Journal Management Tasks**.

A journal is an object that records changes to the file by sending information to the journal receiver. This information includes the file name and library it belongs to, the job ID, user ID,

workstation ID, program name, date and time, type of activity, and relative record number of the affected record. You can also record information before and after a change takes place, giving you a more complete picture of the change being made.

You can work with journals and journal receivers using the IBM Systems Director Navigator for i interface.

For more information about Journal Management, refer to IBM Information Center for IBM i 6.1 at:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaki/rzakiki ckoff.htm?tocNode=int_98514

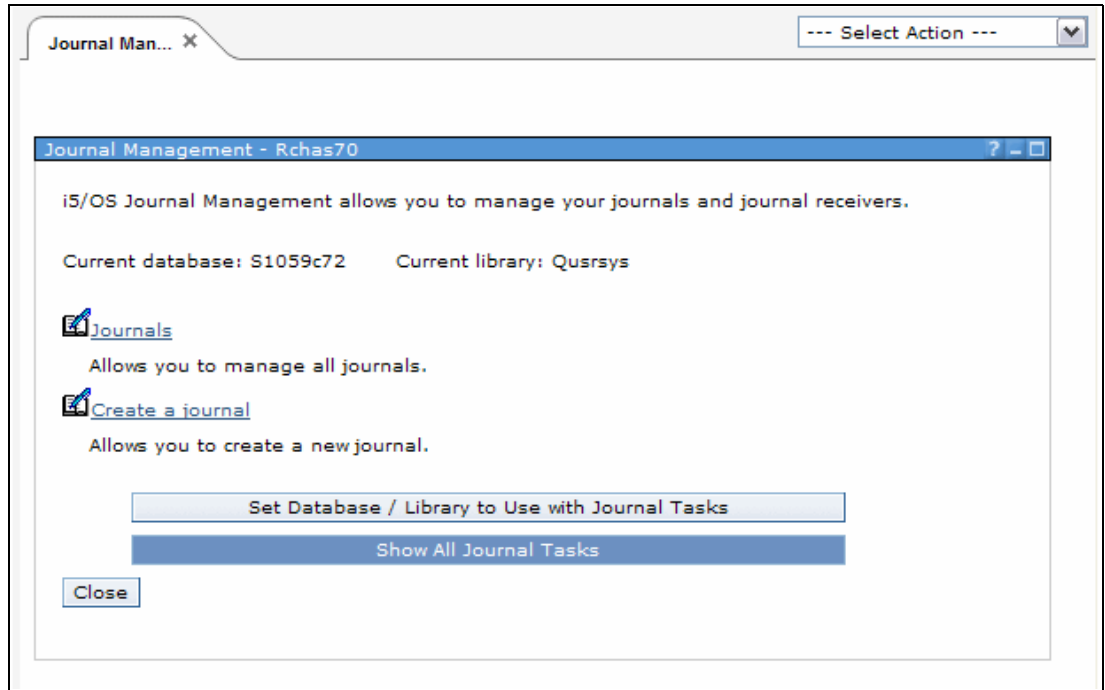


Figure 16-19 Journal Management welcome page

Performance

The Performance welcome page, shown in Figure 16-20, provides links to the most commonly used Performance tasks. To view a complete list of available tasks, click **Show all Performance Tasks**.

You can collect performance data in a collection by using a collector. To start collecting data, you must configure and start a collector. Once some data is collected into a collection, you can view the collection to investigate the data.

For more information about performance, refer to Chapter 6, “Performance” on page 125.

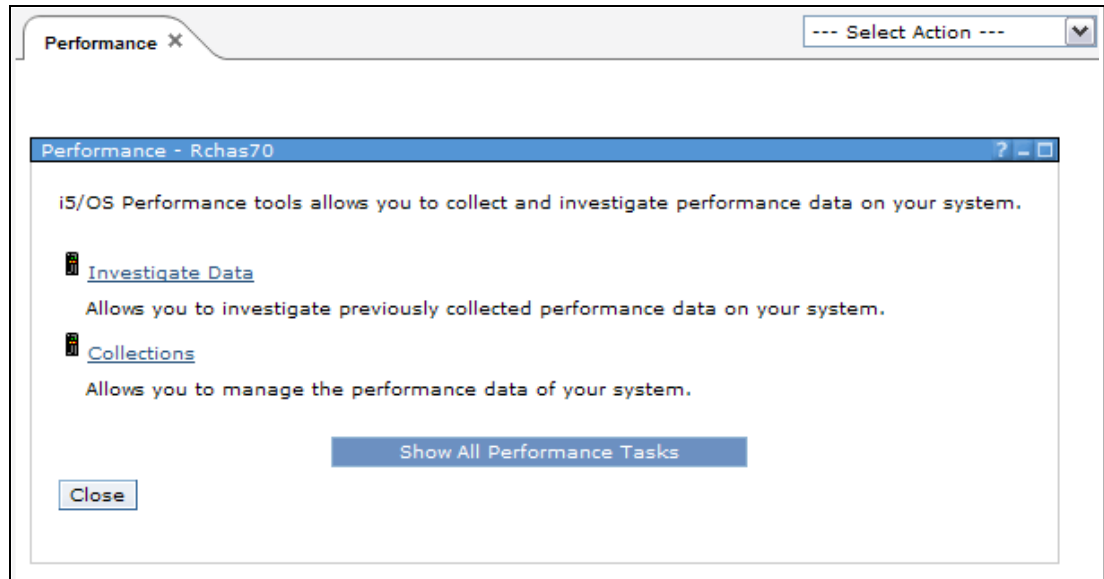


Figure 16-20 Performance welcome page

File Systems

The File Systems welcome page, shown in Figure 16-21, provides links to the most commonly used File Systems tasks. To view a complete list of available tasks, click **Show all File Systems Tasks**.

File Systems are a part of IBM i that support stream input/output and storage management similar to personal computer and UNIX operating systems. They also provide a similar hierarchical directory structure.

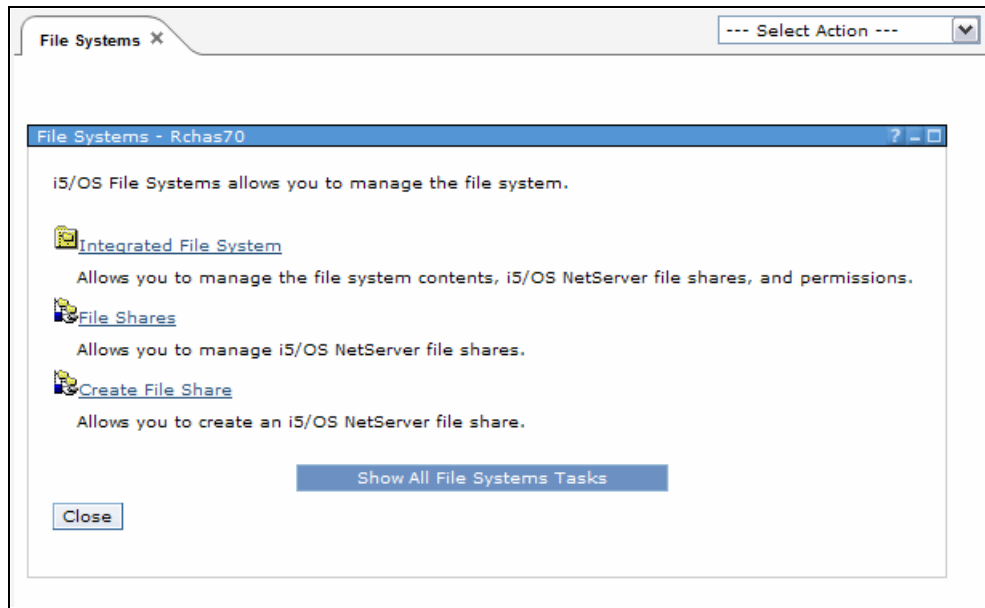


Figure 16-21 File Systems welcome page

Internet Configurations

The Internet Configurations welcome page, shown in Figure 16-22, provides links to the most commonly used Internet Configurations tasks.

From the Internet Configurations page you can work with IBM i Internet configuration tasks: Configure your HTTP servers, access System i Navigator Tasks on the Web, work with Digital Certificates, and more.

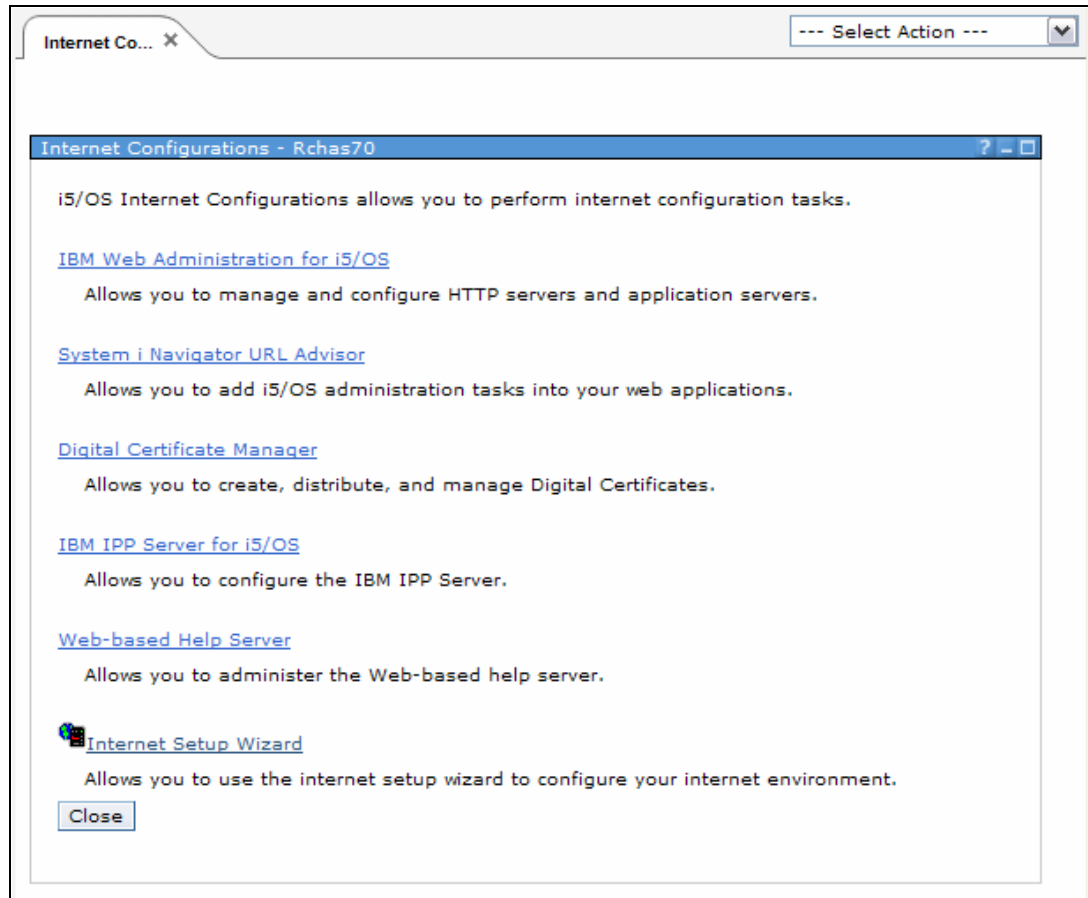


Figure 16-22 Internet Configurations



System i Access products

This chapter discusses the following topics:

- ▶ “Windows” on page 430
- ▶ “Linux” on page 442
- ▶ “System i Access for Web” on page 442
- ▶ “Wireless” on page 446

Note: In this chapter, we refer to both 6.1 and V6R1. 6.1 refers to the operating system level (IBM i) while V6R1 refers to the level of System i Access.

17.1 Windows

New System i Access for Windows administration functions are available for this release.

With the installation of IBM i 6.1 you can manage your environment through new System i Access for Windows functions that have been added to the database providers and through other product enhancements. These include:

- ▶ Communications functions that support the Internet Protocol version 6 (IPv6)
- ▶ Data Transfer functions that support 128-byte table names and a new download option to open files
- ▶ PC5250 that supports Personal Communications 5250, version 5.9
- ▶ .NET Data Provider supports ADO .NET 2.0
- ▶ Support of the Microsoft Installer Package
- ▶ ODBC, OLE DB, Data Transfer enhancements
- ▶ Distributed on DVDs

The Windows 98 (all editions), Windows ME, and Windows NT® operating systems are not supported with the V6R1 System i Access for Windows product.

17.1.1 .NET Data Provider

Many of the following enhancements are based on IBM i 6.1 operating system enhancements.

- ▶ Supports the ADO.NET 2.0 base class model.
- ▶ Running in 64-bit mode.
- ▶ Distributed transactions using the System.Transactions namespace.
- ▶ Decimal float and datalink data types (IBM i 6.1).
- ▶ Multi-row blocked inserts.
- ▶ Improved control over sending and receiving LOB data in blocks.
- ▶ SQL query storage limit (IBM i 6.1).
- ▶ New SQL special client information registers.
- ▶ Enhanced support for optimistic updates using row change time stamp.
- ▶ Integration with Visual Studio 2005 to enhance design-time experience when writing applications.
- ▶ Ability to specify the IsolationLevel for autocommit (IBM i 6.1).

Note: The C/C++ Database APIs (Optimized SQL APIs) are no longer being enhanced or supported. See the Programmer's Toolkit for other technologies that you can use for database access.

17.1.2 OLE DB

OLE DB enhancements include:

- ▶ Supports setting a maximum storage on SQL queries.

- ▶ Supports using decimal floating point data.
- ▶ Handles errors from numeric data.
- ▶ Truncates trailing blanks from input data.
- ▶ Sets and passes client information to and from the database host.
Includes client accounting, program IDs, user IDs, application names. and workstation names.
- ▶ Capability to handle stored procedures compiled with differing date and time formats.
- ▶ Improved support for default values (Extended Indicator support).
 - Enables an application to specify that the column's default value should be used to insert data, instead of passing data in the parameter itself.
 - Enables an application to use the indicator variables associated with the host variables to pass special values designating DEFAULT and unassigned. These new special values can be used in UPDATE and INSERT SQL statements, in both static and dynamic forms.

17.1.3 Basic operations

Basic operations changes include:

- ▶ New printer configuration wizard to configure a new printer and associated objects using the following interfaces:
 - Windows client System i Navigator interface
 - IBM Systems Director Navigator for i interface
- ▶ Support for configuring and managing printer shares extended to the IBM Systems Director Navigator for i interface.
- ▶ Printers can be deleted one of the following:
 - Windows client System i Navigator interface
 - IBM Systems Director Navigator for i interface

See 9.3, “Adding a printer with System i Navigator” on page 239 for more details.

17.1.4 Work management

Work management changes include:

- ▶ Job properties (General tab)
The General dialog now includes Disk pool information for the library that contains the job description. This dialog also includes a new possible value for the Function field which indicates the user-specified function set with the Change Current Job (QWCCCJOB) API.
- ▶ Date/Time
The Date/Time dialog now includes Year offset and Day of week information.
 - The Year offset information specifies the number of years that the current year in the calendar system used differs from the current Gregorian calendar year.
 - Day of week specifies the day of the week, which is determined by the job's current date and current time.

See Figure 18-3 on page 457 for further details.

- ▶ **Server jobs**
The Server dialog now includes the client port number used by a server job.
- ▶ **Job Queue**
The Job queue dialog now includes Disk pool information for the library that contains the job queue.
- ▶ **New support to allow you to view transaction and permission information using the IBM Systems Director Navigator for i interface.** See Chapter 16, “IBM Systems Director Navigator for i” on page 403.

17.1.5 Network

Many of the networking functions have been updated to accept and use IP version 6 information similar to how these functions use IP version 4 information:

- ▶ **Virtual Private Networking**
 - Use IP version 6 to create a VPN with the following connection types:
 - Host-to-host
 - Host-to-gateway
 - Gateway-to-gateway
 - VPN connections support IP version 6 features including address, range, subnet, and host name.
 - All VPN wizards were updated to accept the new IP version 6 ID types.
- ▶ **Remote Access Services**
Remote Access Services under IBM i adds a function that allows you to select whether you want to wait for a dial tone before using the modem.
- ▶ **Digital Certificate Manager**
Under V6R1, Digital Certificate Manager has been updated to allow administrators to easily check the expiration dates of the Certificate Authority (CA) certificates in their certificate stores. This helps you take corrective action before an expiration occurs.
- ▶ **Enterprise Identity Mapping**
Under V6R1, the Enterprise Identity Mapping (EIM) configuration wizard has been updated to allow you to select the directory server instance you wanted to use to create the EIM domain on when one or more directory server instances already existed on the system.
- ▶ **Universal Connection wizard** has been updated to allow you to specify an IBM-assigned customer number and contract number.

The customer number is used in various business and service transactions with IBM. The contract number is used to look up all customer-purchased services under the identified contract.

17.1.6 TCP/IP servers

The TCP/IP servers have received the following enhancements:

- ▶ **Simple Mail Transfer Protocol (SMTP)**
Allow Simple Mail Transfer Protocol (SMTP) to send e-mail through a secure connection and authenticate with a user ID and encrypted password using Secure Sockets Layer (SSL) and Transport Layer Security (TLS), as shown in Figure 17-1 on page 433.

- ▶ Open Shortest Path First (OSPF)
 - IBM i 6.1 includes new support for OSPF. A new TCP/IP server has been added called OMPROUTED to support the OSPF (Open Shortest Path First) routing protocol.
 - Each router or system that participates in an OSPF network will have the same link state database, which is used to calculate the distance to other systems using the Shortest Path First (SPF) algorithm.
- ▶ Post Office Protocol (POP)

In IBM i 6.1, you can allow Post Office Protocol (POP) to use encryption to secure POP data streams (including passwords). Encryption is provided with either the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) as shown in Figure 17-2 on page 434.
- ▶ NetServer

Support for configuring and managing NetServer file and printer shares is available in the browser-based IBM Systems Director Navigator for i interface.
- ▶ IBM Tivoli Directory Server for i5/OS

Support for configuring and managing multiple instances of the directory server on the system.

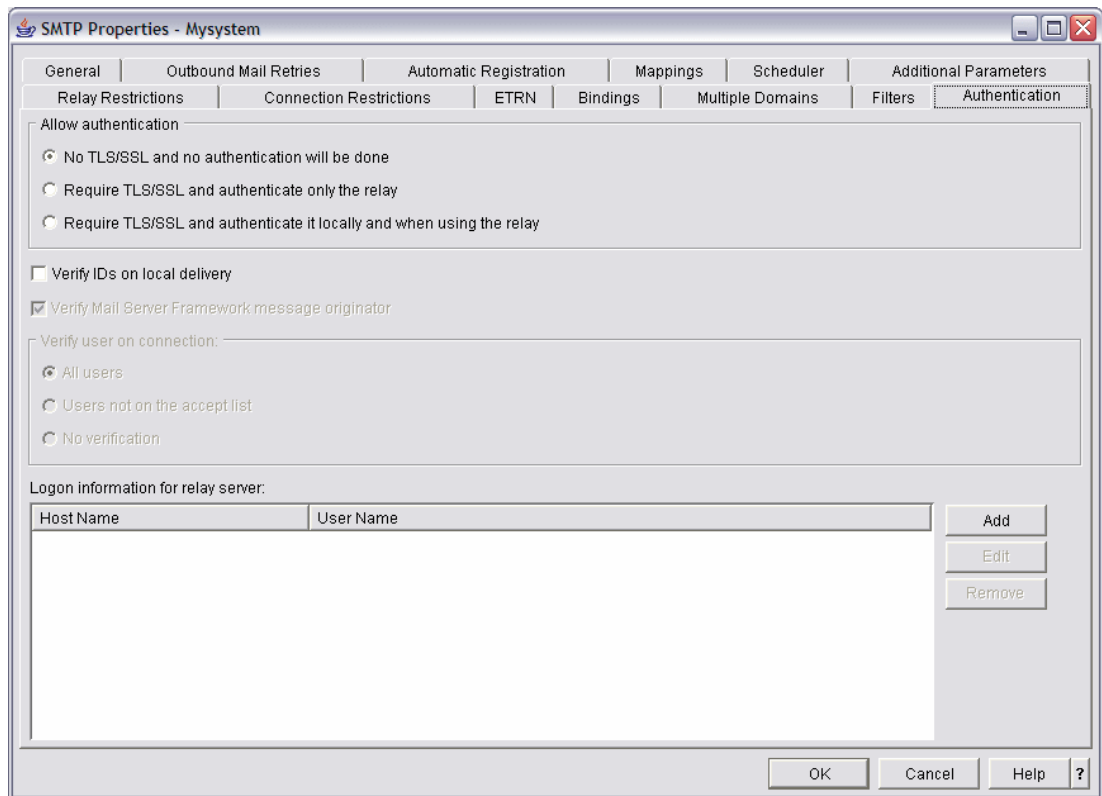


Figure 17-1 SMTP server properties

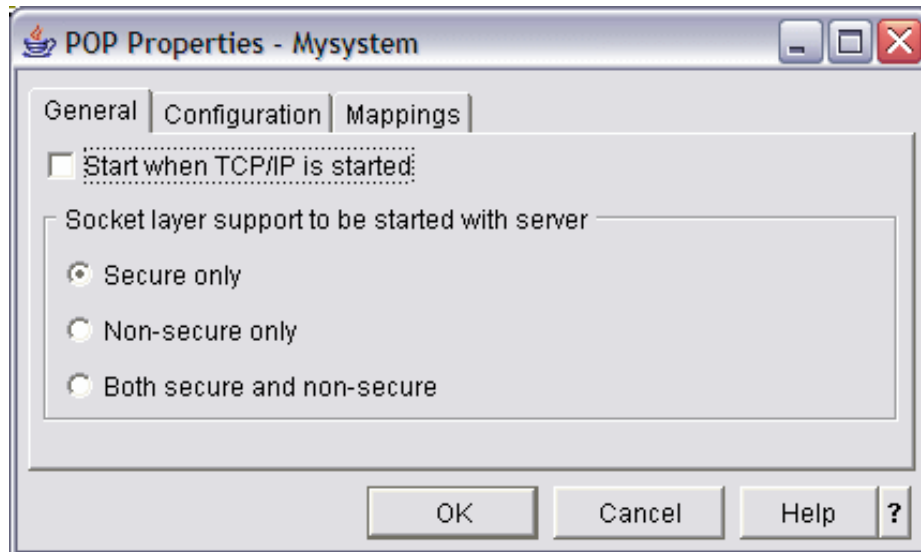


Figure 17-2 POP server properties

17.1.7 Integrated server administration

Added managing and configuring integrated servers running Linux or Windows operating systems to the IBM Systems Director Navigator for i interface.

Supported tasks include:

- ▶ Starting and stopping integrated servers
- ▶ Managing integrated server disks
- ▶ Managing integrated server configuration objects
- ▶ Managing user enrollment and running remote commands for integrated Windows servers

17.1.8 Security

Security changes include:

- ▶ Cryptographic Services Key Management (see Figure 17-3 on page 435)
 - Master key creation and management.
 - The “Manage keystore files” option allows you to view, add, create, delete, translate, or display the properties of the keystores.
 - Allows you to view, create, extract, export, or display the properties of the key records.
 - Master Key Management interfaces added to both Windows System i Navigator and browser-based IBM Systems Director Navigator for i.

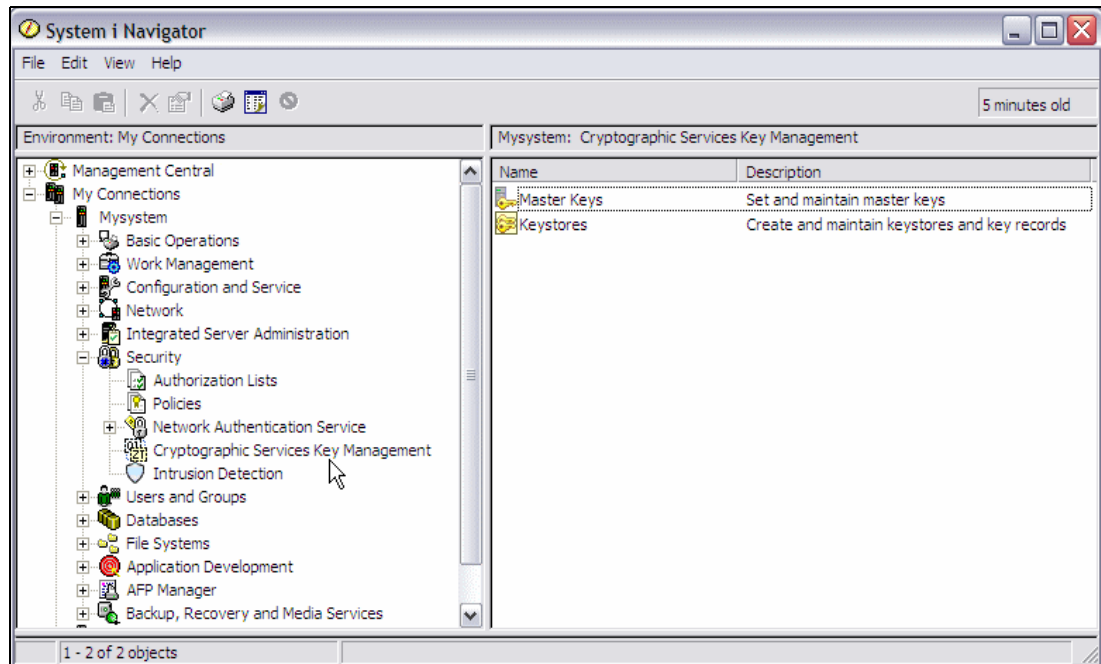


Figure 17-3 Cryptographic Services Key Management

- ▶ TCP/IP intrusion detection
 - Greatly expanded intrusion detection support.
 - Policies defined using the Windows System i Navigator and IBM Systems Director Navigator for i interfaces.
 - See 2.1, “Intrusion detection and prevention” on page 6 for further details.
- ▶ Permissions

Permissions and resource security interfaces are now available through IBM Systems Director Navigator for i.

17.1.9 Users and groups

IBM Systems Director Navigator for i interface supports creating and managing IBM i users and groups. New IBM i 6.1 user profile parameters can be displayed and changed. This includes support, for example, to prevent a user from changing a password immediately.

17.1.10 Structured query language

Structured query language (SQL) changes include:

- ▶ SQL performance monitors enhancements
- ▶ New simplified compare two monitors output tool
- ▶ Save Analyze Results to a spreadsheet
- ▶ SQL plan cache can show active jobs in the cache
- ▶ Run SQL scripts:
 - Check your SQL statements for any non-standard statements that might impact porting SQL statements.

- View Visual Explain refresh results while a statement is running.
- Enhanced index abilities.
- Interface to support derived key expressions and sparse index.

Note:

SQL performance monitors: A new simplified compare tool allows you to quickly compare two monitors and see if further comparison is warranted. You can save your analyze results to a spreadsheet. Additionally, there are new filtering options available to further refine your monitoring selections.

SQL plan cache: SQL plan cache functionality has been enhanced to allow you to show a list of active jobs in the plan cache and to edit the properties of your plan cache.

Health Center: Environmental limits and Activities pages are now available through the Health Center interface. These pages allow you to run SQL scripts to better monitor the health of your database.

17.1.11 Integrated file system

The changes to the Integrated File System include:

- ▶ Folder Attribute Information context menu has been added to allow you to retrieve storage and attributes information used by the Integrated File System objects (see Figure 17-4).

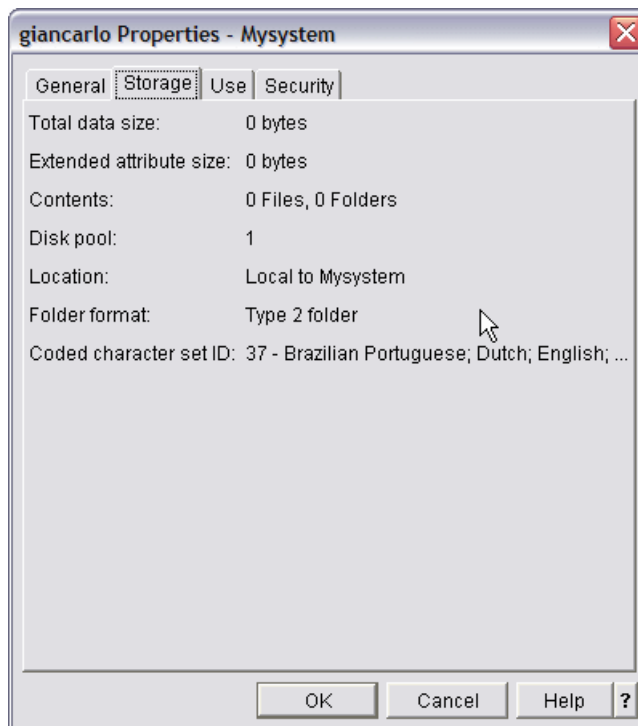


Figure 17-4 Folder Storage Size

- ▶ Check in and check out functions available for folders.
- ▶ Dynamic Mount Information page has been added so that you can view the list of currently mounted file systems (see Figure 17-5).

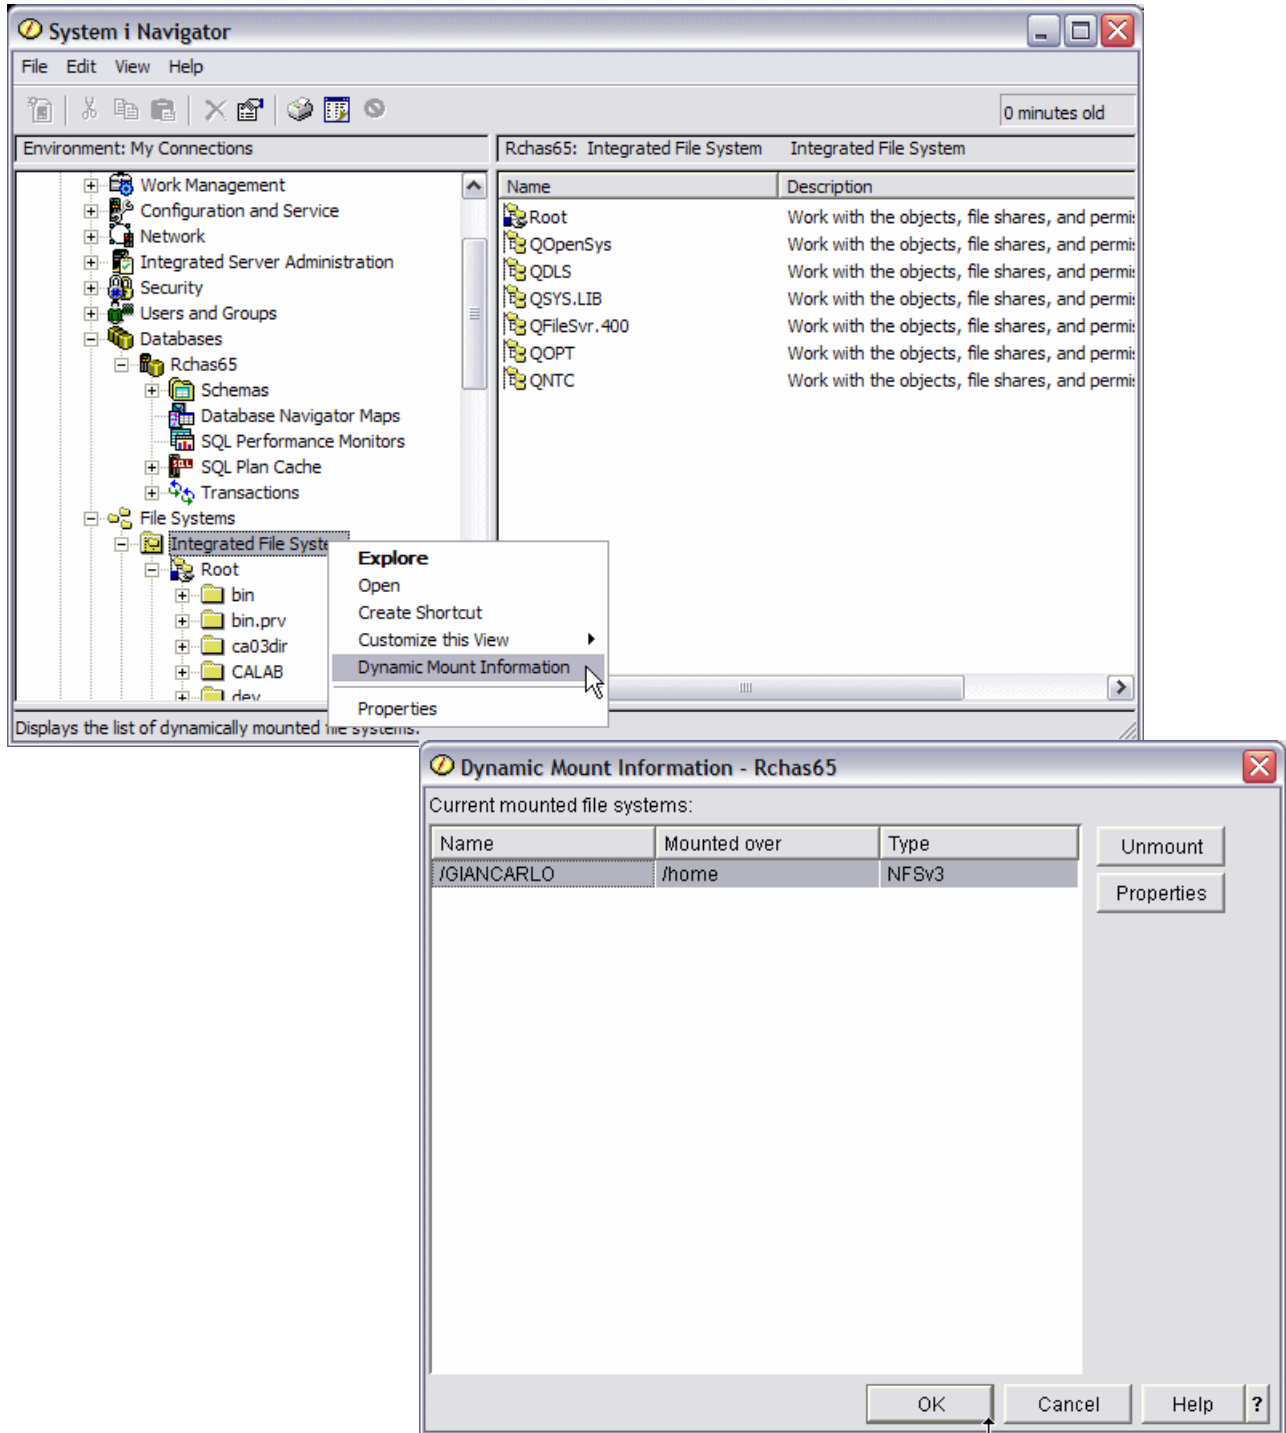


Figure 17-5 Dynamically mounted NFS file system

17.1.12 Network File System

The changes to the Network File System include:

- ▶ Using System i Navigator tasks on the Web you can also specify the NFS version and security options that you want to use on the Advanced page when you are working on systems running IBM i 6.1 or later.

- ▶ A new export option enabling you to specify an external name for the NFS version 4 export, while locally the export has another name (see Figure 17-6).

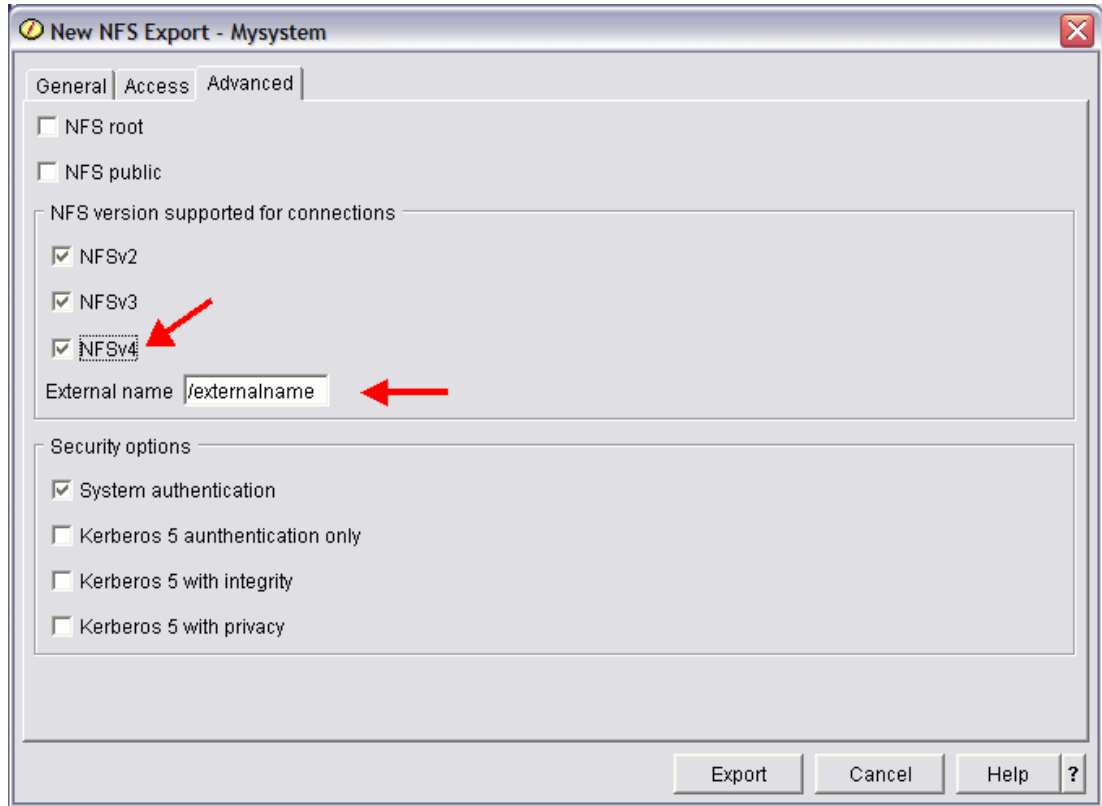


Figure 17-6 NFS Export using NFS v4 allows to use an external name

17.1.13 Backup Recovery and Media Services

BRMS 6.1 enhancements available using System i Navigator include:

- ▶ Backup to tape and IBM i virtual tape encryption
- ▶ Backup to optical devices
- ▶ Simplified Backup and Restore wizards
- ▶ Media Policies (previously hidden)
- ▶ More HTML reports
- ▶ Media library manager (define which systems are attached to media libraries)
- ▶ Changed the terms Backup Policy and Archive Policy to Backup Control Group and Archive Control Group, respectively
- ▶ Copy Backup and Archive Control Groups to other systems
- ▶ New view for Backup Control Group and Archive Control Group Properties
- ▶ In general, all IBM i 6.1 BRMS enhancements are added to both of the following:
 - Windows-based client workstation System i Navigator interface
 - Browser-based IBM Systems Director Navigator for i interface

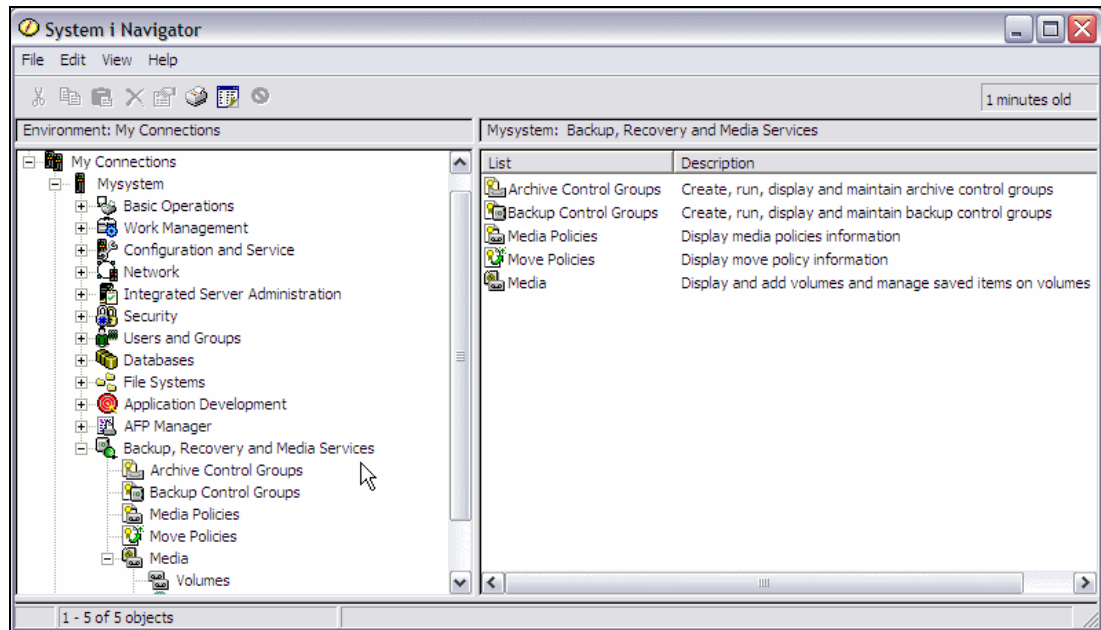


Figure 17-7 Backup Recovery and Media Services

17.1.14 Management Central

This section lists enhancements to the functions under Management Central.

► Performance

There are now performance enhancements under the Windows interface. However, IBM i 6.1 contains a rich new set of Performance Management enhancements using the performance functions under IBM Systems Director Navigator for i. These include:

- Collection Services
- Disk Watcher
- Job Watcher graphs for performance analysis
- Manage Collection Services
- Disk Watcher
- Job Watcher collections

For more information about performance, refer to Chapter 6, “Performance” on page 125.

► Fixes inventory

- You can now have different levels of fix groups on your system.
- Ability to send distribution requisites along with your prerequisites and corequisites.
- Additional statuses for fix groups, allowing you to better determine the status of your fix group.
- Compare and update fix groups on multiple systems as well as a single system.

► Monitors

You can view a 24-hour time window in your system monitor graph history window.

- ▶ Users and groups

Additional support is added to allow you to better manage users. For example, you can now prevent a user from changing a password immediately (mirrors new IBM i password control capabilities).

17.1.15 Configuration and service: Time management and time zone

The changes in time management and zones include:

- ▶ Support for time Alternate name

This is a descriptive name that provides additional information about the time zone description.

- ▶ Ability to change the Year offset for a time zone object

This is the number of years that the current year in the calendar system used with this time zone differs from the current Gregorian year.

- ▶ Support to specify, in minutes, the amount of time for Daylight Saving Time shift when Daylight Saving Time has been enabled.

For more information on IBM i 6.1 time-associated enhancements, see 18.9, “Time management” on page 456.

17.1.16 Configuration and service: Hardware

The changes in hardware management include:

- ▶ Disk units

- Support for parity hot spares configuration.

- Disk tasks are also accessible using IBM Systems Director Navigator for i (see Figure 17-8 on page 441).

- Support for cross-site mirroring functions when High Availability Solutions Manager (5761-HASM) has been installed.

- Support for (creating) encrypted disk pools when the Encrypted ASP Enablement option (5761-SS1 Option 45 has been installed.

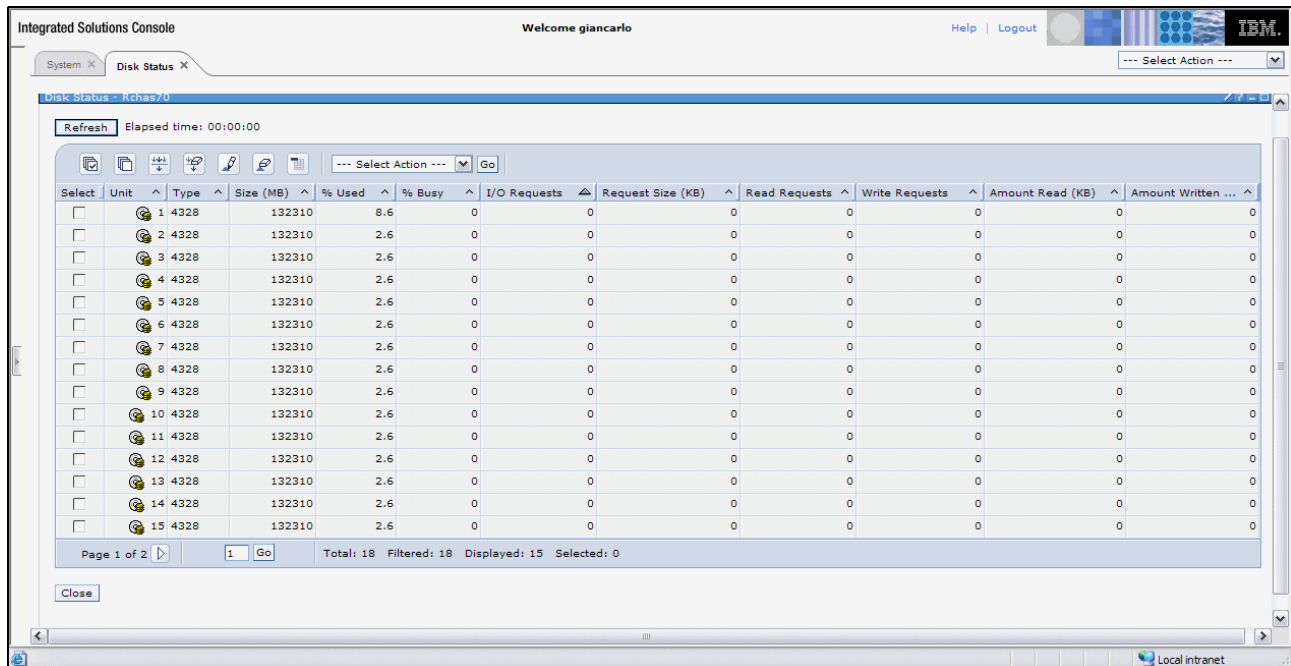


Figure 17-8 Disk Status

- ▶ Tape devices
 - Create a virtual tape device using System i Navigator.
 - Update the drive firmware for a standalone tape device.
 - On the Create Image Catalog -Options page, you can specify details about virtual volumes that you want added to the image catalog. On the Create Image Catalog, you can now specify the catalog ASP threshold.
 - On the Image Catalog Properties pages, you can specify the current directory for the image catalog, and whether to import the catalog and volume information from an image catalog directory. Also, you can now specify the catalog ASP threshold.
 - The virtual volumes list now has a new column for Allocated size of the volume.

17.1.17 Installation and setup

There is new and changed installation and set-up information for this release of System i Access for Windows.

- ▶ The Windows Installer technology is being used.
- ▶ Support for Windows Vista has been added.
- ▶ Previously, the Tailored Installation Wizard let administrators control which System i Access for Windows components you could install. This ability has been replaced by Windows Installer Transforms.
- ▶ Silent Installations have been replaced by User Level Interface.
- ▶ You can use Public Properties to control a user's environment. For information, see Public Properties.
- ▶ Service Packs are implemented by using minor upgrades. See “Preparing to install service packs on multiple PCs” for more information.
- ▶ EZ-Setup and the ADO Wizards for Visual Basic have been removed.

For complete information about Installation and Setup, see the IBM i Information Center at the following Web address and select the path **Connecting to System i** → **System i Access** → **System i Access for Windows** → **Installation and setup**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

17.2 Linux

The System i Access for Linux (5761-XL1) offering is the latest in the System i Access (5761-XW1) family of products. It offers Linux-based access to your IBM i operating environment. The System i Access for Linux product can leverage business information, applications, and resources across an enterprise by extending the system's resources to the Linux client. The product allows you to establish a 5250 display emulation session or to access the IBM DB2 for i database, using an ODBC driver, all from a Linux client.

The following new capabilities are included with the V6R1 release:

- ▶ Remote Command Unicode API support
The remote command Unicode interface `cwbRC_RunCmdW` now flows Unicode data to IBM i. IBM i must be at 6.1, or later.
- ▶ IPv6 support for the System i Access toolkit
The ODBC, Remote Command, and other utilities are supported over an IPv6 network. 5250 sessions are not supported over IPv6.
- ▶ ODBC enhancements are equivalent to the Windows ODBC driver.

System i Access for Linux is shipped with the System i Access for Web (5761-XH2) product and is available as a download by accessing the System i Access for Linux homepage:

<http://www.ibm.com/systems/i/software/access/linux/>

17.3 System i Access for Web

System i Access for Web (5761-XH2) is a key offering in the System i Access Family (5761-XW1). It offers Web browser-based access to IBM i resources on System i systems. System i Access for Web enables users to leverage business information, applications, and resources across an enterprise by extending the IBM i resources to the client desktop through a Web browser (see Figure 17-9). Use the following URL to access the System i Access for Web main page and replace *host* with your system name:

`http://host:37000/webaccess/iWAMain`

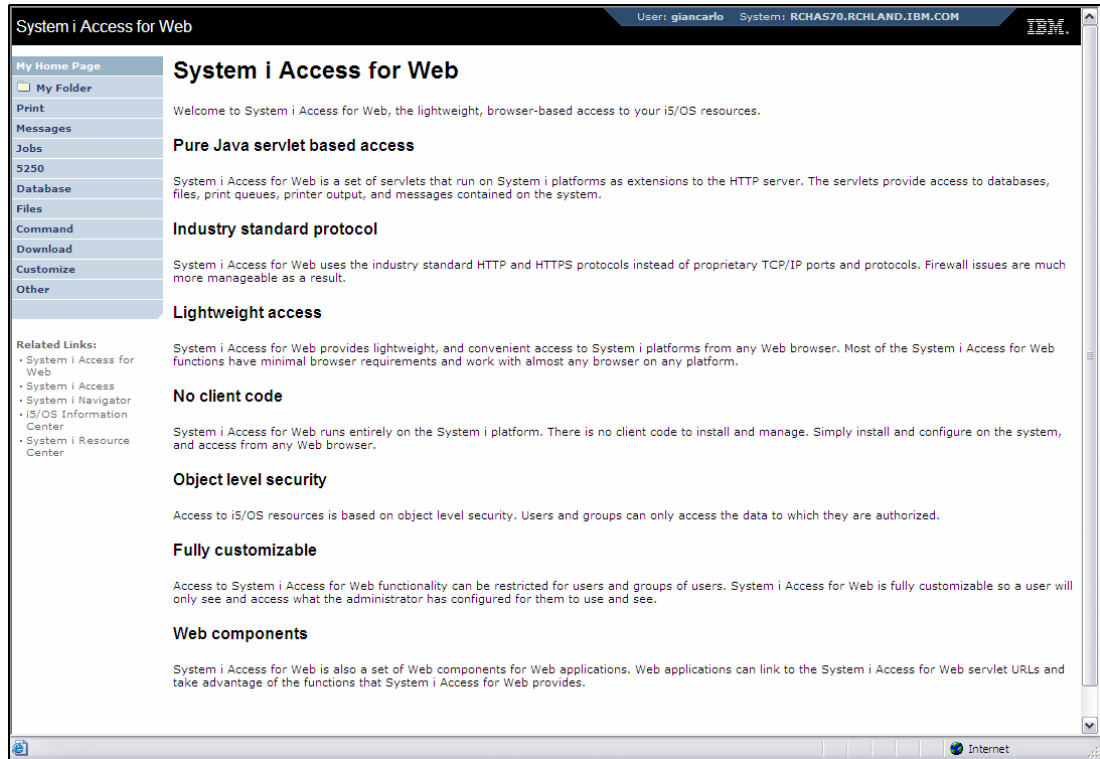


Figure 17-9 System i Access for Web

System i Access for Web has the following advantages:

- ▶ It is server-based.
- ▶ It is implemented using Java servlet and portlet technology.
- ▶ It uses industry standard protocols: HTTP, HTTPS, and HTML.
- ▶ It is lightweight, requiring only a browser on the client.
- ▶ It provides a 5250 user interface and access to IBM i resources, such as the System i database, integrated file system, printers, printer output, jobs, batch commands, and messages.

The following information describes the new support in the Web application server environment.

17.3.1 Database function

The following enhancements have been made to the database support:

- ▶ SQL Wizard
 - The wizard has been enhanced to handle building statements which join data from multiple tables.
- ▶ Additional supported file types
 - Run SQL added support for two new output file types: “Scalable Vector Graphics Chart (.svg)” and “Text – Delimited (.txt)”. Scalable Vector Graphics Chart provides a way to store SQL results in a bar, area, line, or pie chart. Text – Delimited is a delimited text file format in which you specify the delimiter to use. Copy Data to Table also added support for the “Text – Delimited (.txt)” file format.

- ▶ Excel date and time format

Run SQL and Copy Data to Table updated the Microsoft Excel XML (.xml) support to read and write date and time values using the Excel date and time formats.
- ▶ Transfer saved requests

My Requests has support for a new “Transfer request” action, which provides a way to transfer ownership of a request to and from a group, of which the current user is a member. By transferring a request to a group, multiple users can share in the management of the request.
- ▶ Create automated task

My Requests has support for a new “Create automated task” action, which provides a way to create a Java archive (.jar) to run the database request. The Java archive can be invoked by a task scheduling function to run the database request at a specified time.
- ▶ Import/Export requests

My Requests now supports exporting requests to a file and importing requests from a file. This provides users the ability to copy requests to other Web application server environments, to other users, and to other systems running System i Access for Web.
- ▶ DECFLOAT

DB2 for i added support for Decimal Floating Point (DECFLOAT) columns in IBM i 6.1. System i Access for Web updated the following database functions to handle DECFLOAT columns:

 - Insert Record
 - Update Record
 - Find Record
 - Run SQL
 - SQL Wizard
 - Copy Data to Table
 - Import Request
- ▶ Print function

The print support has been enhanced as follows:

 - Delete multiple printer output files

Users can now select multiple printer output files for deletion at one time.
 - Default PDF output option

A new preference was added to specify the default PDF output destination to use when the PDF output settings page is bypassed. The default destination can be set to:

 - Browser
 - Mail as Attachment
 - Personal Folder
 - Output Queue
- ▶ Single sign-on using Windows domain login

System i Access for Web added the ability for a Windows workstation user to log on to a Windows domain and have System i Access for Web use those credentials to authenticate to IBM i.
- ▶ List sorting

Most lists displayed by System i Access for Web now support sorting column data in ascending, descending, and unsorted order.

► Automating tasks

System i Access for Web added the ability to automate certain System i Access for Web tasks. Tasks which can be automated can also be scheduled to run at a specified time, using a task scheduling function.

► Removal of ASF Tomcat support

The ASF Tomcat Web application server environment is not supported by System i Access for Web in V6R1. WebSphere Application Server and the integrated Web application server are the supported Web-serving environments for System i Access for Web in V6R1. The CFGACCWEB2 configuration command provides a way to migrate System i Access for Web user data to a supported Web-serving environment.

See Figure 17-10 to view a list of the database options.

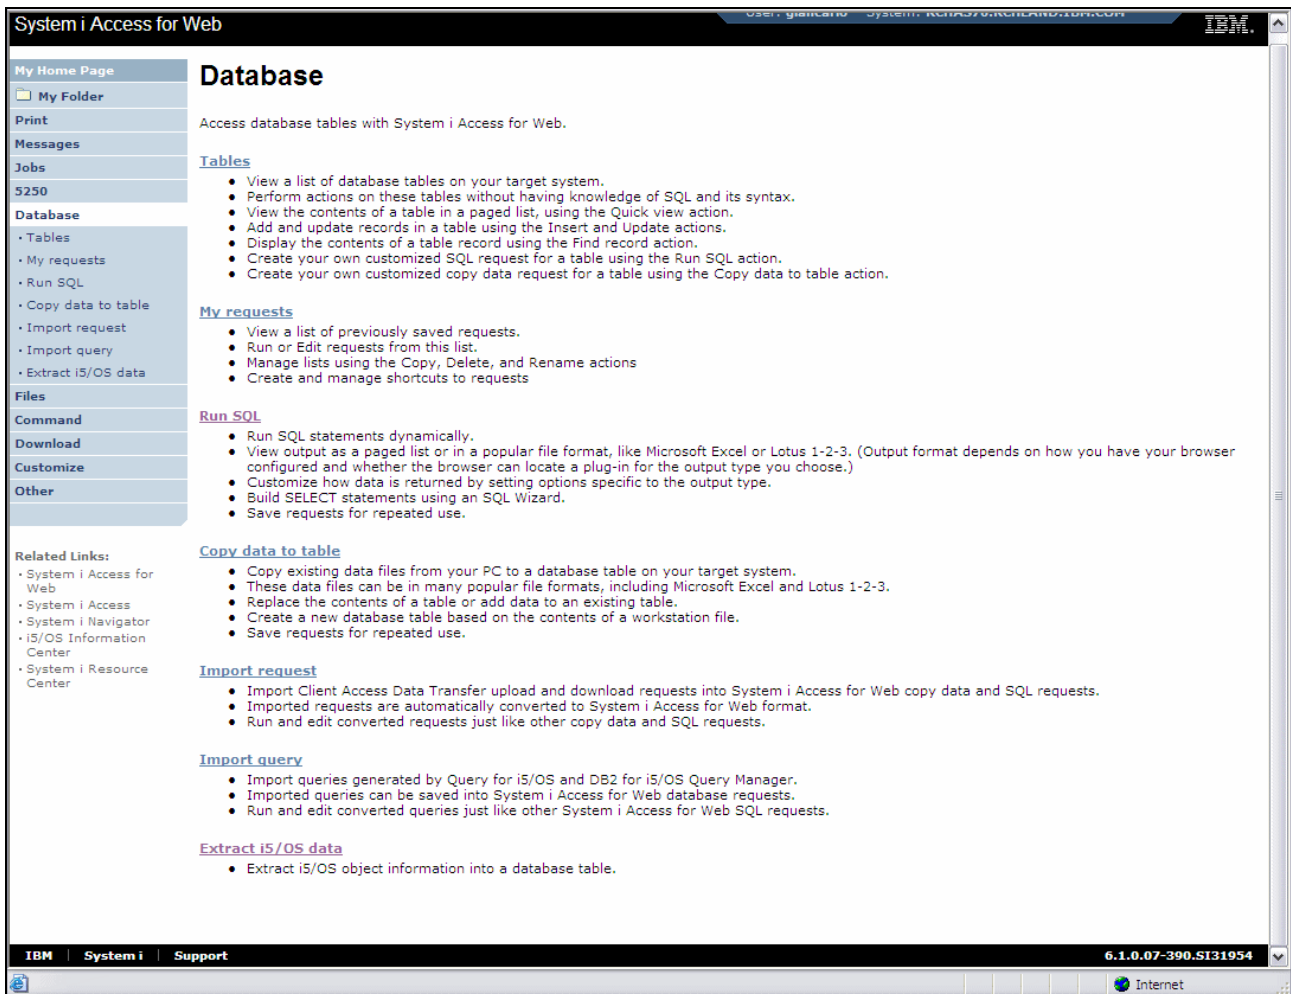


Figure 17-10 System i Access for Web: Database Options

17.3.2 New portal environment support

The new support in the portal environment is as follows:

Portal application (Standard): System i Access for Web added a new portal application offering that implements the Java Specification Requests (JSR) 168 Portlet Specification. A portal application offering, which implements the IBM Portlet API, is also still provided.

17.4 Wireless

Beginning in V6R1, System i Navigator for Wireless can run using only the Lightweight Infrastructure Web application server as the servlet engine. Previously, you could choose from an Apache Software Foundation (ASF) Jakarta Tomcat servlet engine or a WebSphere Application Server to run the servlet.

For complete information about System i Access for Wireless, see the IBM i Information Center at the following Web address and select the path **Connecting to System i** → **System i Navigator** → **System i Navigator for Wireless**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

17.4.1 Version and release changes in the ODBC driver behavior

The following list describes some of the important changes for V6R1:

- ▶ 128-byte cursor names
- ▶ SQL query storage limits
- ▶ ODBC application and QZDASOINIT system jobs association
- ▶ Decimal Floating Point (DECFLOAT) data type
- ▶ Capability to handle stored procedures compiled with differing date and time formats
- ▶ Ability to specify the IsolationLevel for autocommit
- ▶ Improved support for default values (Extended Indicator support)
 - Enables an application to specify that the column's default value should be used to insert data, instead of passing data in the parameter itself.
 - Enables an application to use the indicator variables associated with the host variables to pass special values designating DEFAULT and unassigned. These new special values can be used in UPDATE and INSERT SQL statements, in both static and dynamic forms.

17.4.2 Developing System i Navigator plug-ins

With the plug-in feature of System i Navigator, you can integrate your system administration tasks and client/server programs into a single application environment.

You can use plug-ins to consolidate third-party applications and specialized functions written in C++, Visual Basic, or Java into the System i Navigator interface. Use this topic collection to help you understand what plug-ins are, how to create or customize them, and how to distribute them to your users.

New field in setup.ini

The information section of the setup.ini file has a new field called EclipseHelp. This field indicates whether the plug-in application uses the Eclipse platform to develop the help.

For complete information about Developing System i Navigator plug-ins, see the IBM i Information Center at the following Web address and select the path **Connecting to System i** → **System i Access** → **System i Access for Windows** → **Introduction** → **Installation and setup** → **Setting up the PC** → **Integrating System i Navigator plug-ins**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

17.4.3 Tasks on the Web

IBM Systems Director Navigator for i is a new Web console interface for System i administration where you can work with the Web-enabled tasks of System i Navigator, and perform a subset of System i Navigator tasks through an Internet Web browser. The interface might look slightly different, but the tasks you can perform are the same.

The following System i Navigator functions are available to work with on the Web:

- ▶ System (System status, System operator messages, History log, Disk status and more)
- ▶ Basic Operations (Messages, User jobs, Printers, Printer output, and more)
- ▶ Work Management (Active jobs, Server jobs, Output queues, Subsystems, and more)
- ▶ Network (TCP/IP configuration, Remote access services, Servers, IP policies, Internet, NetServer, and more)
- ▶ Configuration and Service (System values, Time management, Disk status, History log, Change password, and more)
- ▶ Integrated Server Administration (All virtual disks, Remote systems, Domains, and more)
- ▶ Security (Authorization lists, Cryptographic services key management, Intrusion detection, and more)
- ▶ Users and Groups (Users, Create users, Groups, and more)
- ▶ Database (Performance monitors, Health center, Databases, and more)
- ▶ Journal Management (Journals, Journal receivers, Create a journal, and more)
- ▶ Performance (Investigate data and collections, and more)
- ▶ File Systems (Integrated file system, File shares, Create file share, and more)
- ▶ Cluster Resource Services (Displays the list of nodes, Deletes the cluster, and more)

For more information about these and other tasks, see the IBM i Information Center at the following Web address and select the path **Connecting to System i** → **System i Navigator** → **System i Navigator tasks on the Web**.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>



Miscellaneous IBM i 6.1 changes

This chapter describes changes in IBM i 6.1 that were not included in the preceding chapters. They are still important changes and hence they have been consolidated into this chapter. The topics are:

- ▶ “System values” on page 450
- ▶ “Integrated file system” on page 452
- ▶ “Virtual optical and virtual tape support” on page 453
- ▶ “Service tools language” on page 454
- ▶ “Watch for event function” on page 454
- ▶ “Power Down System command change” on page 455
- ▶ “Maximum capacities” on page 455
- ▶ “Quiesce I/O for an Auxiliary Storage Pool” on page 455
- ▶ “Time management” on page 456
- ▶ “Models not supported” on page 462
- ▶ “Operations Console” on page 459
- ▶ “Common Information Model” on page 462

18.1 System values

We introduce new or changed system values.

18.1.1 Security auditing level

In the character-based interface, you can specify *AUDLVL for the QAUDCTL system value. By specifying *AUDLVL, you can use any of the auditing actions in the QAUDLVL system value. In addition, you can specify *AUDLVL2 for the QAUDLVL system value. You can use this special parameter (*AUDLVL2) to specify more auditing actions. If the QAUDLVL system value does not contain the value *AUDLVL2, the system ignores the values in the QAUDLVL2 system value. The default value has changed.

In System i Navigator, you can select what actions to audit without differentiating between QAUDLVL2 and QAUDLVL. There is no limit on how many actions you can audit. Figure 18-1 on page 451 shows the selection.

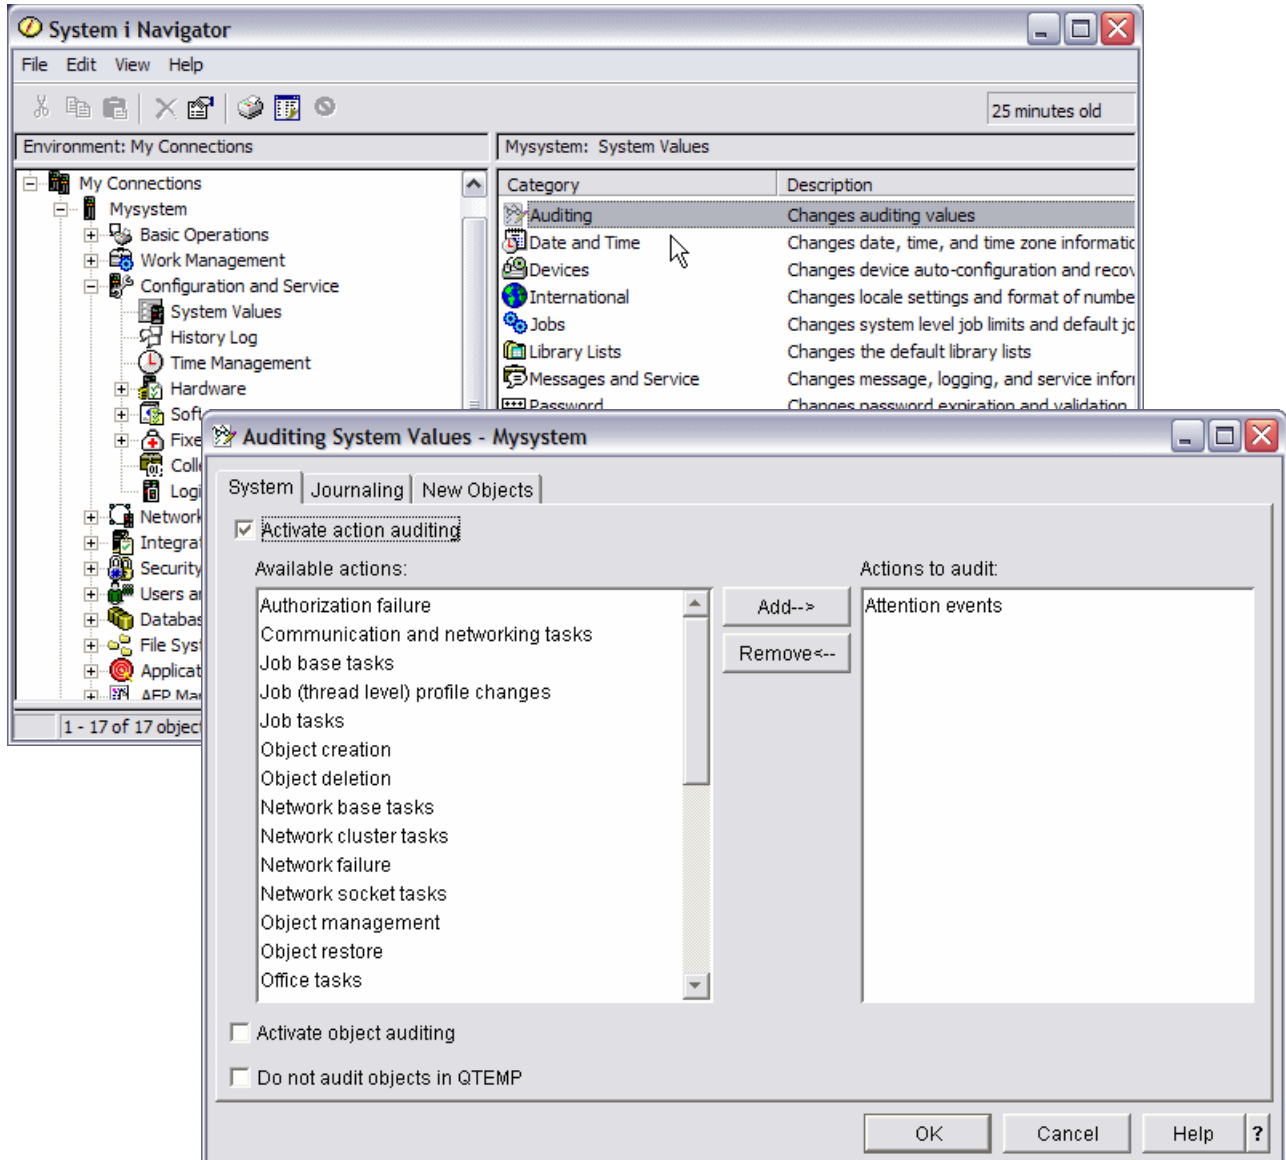


Figure 18-1 Auditing System Values

IBM i 6.1 also provides new auditing actions:

- ▶ Job base tasks (*JOBBAS)

Use this option to audit job base functions; for example, job start and stop data. This option is not supported on systems running IBM i 5.4, or earlier.

- ▶ Job tasks (*JOBDTA)

Use this option to audit actions that affect a job. Use job tasks to monitor who is running batch jobs. The Job tasks option is composed of the Job base tasks (*JOBBAS) value and the Job (thread level) profile changes (*JOBCHGUSR) value. If you specify both of these values, you get the same auditing as if you specified the Job tasks option. This option is not supported on systems running IBM i 5.4, or earlier. The default value has changed to 30.

- ▶ Job (thread level) profile changes (*JOBCHGUSR)

Use this option to audit changes to a thread's active user profile or its group profiles.

18.1.2 Additional active jobs

The recommended setting for the Additional active jobs (QADLACTJ) system value is changed from 100 to be between 10 and 50.

18.1.3 Additional total jobs

The recommended setting for the Additional active jobs (QADLTOTJ) system value is changed from 100 to be between 10 and 50.

18.1.4 Active jobs

The default setting for the Active jobs (QACTJOB) system value is changed from 20 to 200.

18.1.5 Total jobs

The default setting for the Active jobs (QTOTJOB) system value is changed from 30 to 200.

18.1.6 Uninterruptible power supply delay time

The default setting for the Active jobs (QUOPSDLYTIM) system value is changed from Power down system, system calculates delay time to 200.

18.1.7 Security related system values

New and changed system values that control security are documented in 2.4, “New security system values” on page 13 and 2.5, “Changed security system values” on page 18.

18.2 Integrated file system

New functions and enhancements to the Integrated File System (IFS) are as follows:

18.2.1 Collecting and analyzing folder attributes with System i Navigator

As an alternative to the Retrieve Directory Information (RTVDIRINF) and Print Directory Information (PRDIRINF) commands, System i Navigator now provides a function to collect and analyze the attributes of objects in the integrated file system.

For further information, refer to “Collecting and analyzing folder attributes with System i Navigator” in the IBM i Information Center.

18.2.2 Enhanced check-in and check-out functions

Besides checking in and checking out a single file, you can now check in and check out all eligible objects within a folder.

For further information, refer to “Checking in objects and Checking out objects” in the IBM i Information Center.

18.2.3 Automatic conversion to support additional characters in names

Shortly after IBM i 6.1 is installed, the system begins converting directories in file systems that are not case sensitive to support new Unicode characters and casing rules for the names.

For further information, refer to “Converting names to support additional characters” in the IBM i Information Center.

18.2.4 Support for IPv6

The following file systems, which used to support only IPv4 addresses in releases earlier than IBM i 6.1, now support both IPv4 and IPv6 addresses:

- ▶ i5/OS NetClient file system (QNTC)
- ▶ i5/OS file server file system (QFileSvr.400)
- ▶ Network File System (NFS)

18.2.5 Unicode enablement of some CL commands

Some CL commands are now Unicode-enabled.

For further information, refer to the table *Integrated file system commands* in “Accessing using CL commands” in the IBM i Information Center.

18.2.6 Dropping support for the QNetWare file system

As of IBM i 6.1, the QNetWare file system and its related features are no longer supported in the integrated file system.

18.3 Virtual optical and virtual tape support

The following sections describe some of the enhancements to virtual optical and virtual tape support.

18.3.1 Image catalog entry for virtual tape

In addition to the previously existing status values, mounted, loaded, unloaded, and error, there is a new status value, *next*. Next means this image catalog entry is the next one mounted when a command specifying a volume of *MOUNTED is issued to use the virtual tape device. For example, if the Image Catalog is loaded, but there is no entry in the image catalog that is mounted, if VOL(*MOUNTED) is specified the next time that the catalog is used, the *next* entry will be used.

18.3.2 Independent ASP support for virtual tape

An independent disk pool can be used to share virtual images between systems. The following CL command can be used to create a new image catalog to access the virtual volumes stored in a directory in an independent disk pool:

```
CRTIMGCLG IMGCLG(MYCATALOG) DIR('/MYIASPNAME/') TYPE(*TAP) IMPORT(*YES)
```

18.3.3 Size ranges for image files

Optical images can be between 48 and 16,000 megabytes.

Tape images can be between 48 and 1,000,000 megabytes.

18.4 Service tools language

With IBM i 6.1, the system service language can now be a different language from the operating system language.

For systems running IBM i 6.1, or later, you can set the system service language (the dedicated service tools (DST) language) to a language different from the operating system language.

18.5 Watch for event function

The Watch function is enhanced with two features:

The occurrence of Product Activity Log (PAL) entries can be watched through the Watch event for function.

A Display option is added to the Work with Watches display, where you can view the details of the Watch sessions. A sample of the output is shown in Figure 18-2.

```
Display Watch
System: RCHAS70
Session ID . . . . . : SRVMON0000   Started:
Started by:                               Date . . . . . : 08/03/09
Job name . . . . . : QZRCSRVS       Time . . . . . : 11:08:21
User . . . . . : QUSER              Call exit program:
Number . . . . . : 004094           *WCHEVT
Watch program . . . : QPDETWCH
Library . . . . . : QSYS
Origin . . . . . : QSCSWCH
Run priority . . . . : 25

Message
Message queue   Library   Job   User   Number
CPI93B4 QSRVMON   QSERVICE

Bottom

Press Enter to continue

F1=Help  F3=Exit  F11=Comparison data  F12=Cancel
```

Figure 18-2 Display Watch

For more information about the Watch function, refer to the IBM i Information Center or *i5/OS Diagnostic Tools for System Administrators: An A to Z Reference for Problem Determination*, SG24-8523.

18.6 Power Down System command change

The default value for the CONFIRM parameter on the PWRDWNSYS command has changed for IBM i 6.1. The default value for the Power Down System (PWRDWNSYS) command is to display a confirmation. You can change the default, using the Change Command Default (CHGCMDDFLT) command, by specifying *NO for the Confirm (CONFIRM) parameter (CONFIRM(*NO)).

Note: This function is valid only if you have 6.1 installed.

18.7 Maximum capacities

Maximum capacities provides users and administrators with system limits for many IBM i functions. The purpose of this information is to aid in planning and managing how system limits can affect your system operations. To view changes to maximum capacities, refer to the tables in Infocenter by navigating to **Availability** → **Maximum Capacities** → **Updated system limits** → **What's new for 6.1**.

18.8 Quiesce I/O for an Auxiliary Storage Pool

A new command is available, Change ASP Activity (CHGASPACT), which enables you to control the activity of an auxiliary storage pool at a storage pool level, either:

- ▶ System pool 1 (*SYSBAS) and all other existing dependent (basic) storage pools 2 to 32, or
- ▶ A named independent auxiliary storage pool

You can:

- ▶ Suspend database transactions and database and Integrated File System file change operations for the system and configured basic auxiliary storage pools (ASPs) or an independent ASP.
- ▶ Resume suspended transactions and file operations for the system and configured basic auxiliary storage pools (ASPs) or independent ASPs that are currently in a suspended state.
- ▶ Force to disk changed pages associated with the partition and configured basic ASPs or independent ASPs.

You can suspend an ASP in order to facilitate *vary on* after performing a flash copy or geographic mirror detach followed by a vary on. The vary on will go through recovery steps that will rollback incomplete transactions, rebuild access paths, update database records of journaled files and take other recovery actions.

Switch-over has tracking of changes that occur while the switch is in progress.

When specifying force (*FRCWRT), the ASP will remain active during this operation. Forcing changed data to disk should reduce the time it takes to vary off an ASP or to switch an ASP from one system to another.

When specifying *SUSPEND the activity for the auxiliary storage pools specified for the ASP device (ASPDEV) parameter is suspended and the specified ASPs will be quiesced. A 100% quiesce is not required. Initiation of new transactions or new file (database or integrated file system) operations are temporarily halted.

The Suspend Time Out (SSPTIMO) value is valid only when OPTION(*SUSPEND) is specified. Specify in seconds the time allowed for the suspend to complete. Suspension continues even after the time out expires.

The Suspend Time Out action (SSPTIMOACN) specifies the action to take in relation to the specified time out value in SSPTIMO:

- *CONT Suspension should continue, even if the ASP was not able to be quiesced within the specified time.
- *END The suspend operation should be cancelled if the ASP was not able to be quiesced within the specified time.

After you have quiesced (suspended) a disk pool you must use the resume (*RESUME) function to re-enable operations using objects stored in the suspended storage pool.

When a *SUSPEND command is issued, the following steps occur:

1. Flush main store to disk (pinned pages will not be flushed, so not *all* of main store gets flushed). Pinned pages contain data or instructions the system requires in main storage.
2. Suspend database transactions (running under commitment control). This halts all new transactions and lets existing transactions run until they get to a commit boundary. If the IBM i cannot get transactions to a commit boundary within the specified timeout, IBM i flags the transaction as *not quiesced* and suspend will be allowed to continue.
3. Suspend database non-transaction operations. The currently running operations are allowed to complete and new operations are halted.
4. Flush main store again (to get the items that were still in-flight while halting the transactions and operations).
5. At this point, we are suspended. Database activity should be quiesced and at commit boundaries. Non-database activity may still be occurring within the ASP, however.
6. You can then perform whatever action you wanted the quiesce for (for example, Detach/FlashCopy).

Note: A new API, Control ASP Access (QYASPCTLAA) can be used instead of the CHGASPACT command.

18.9 Time management

Within the time management component of System i Navigator, you can work with the time zone and time adjustment functions. With these functions, you can choose a time zone for your system to use and adjust the system time, as shown in Figure 18-3.

In addition, you can identify the software to use as the time maintenance application that synchronizes your system time with an external time source.

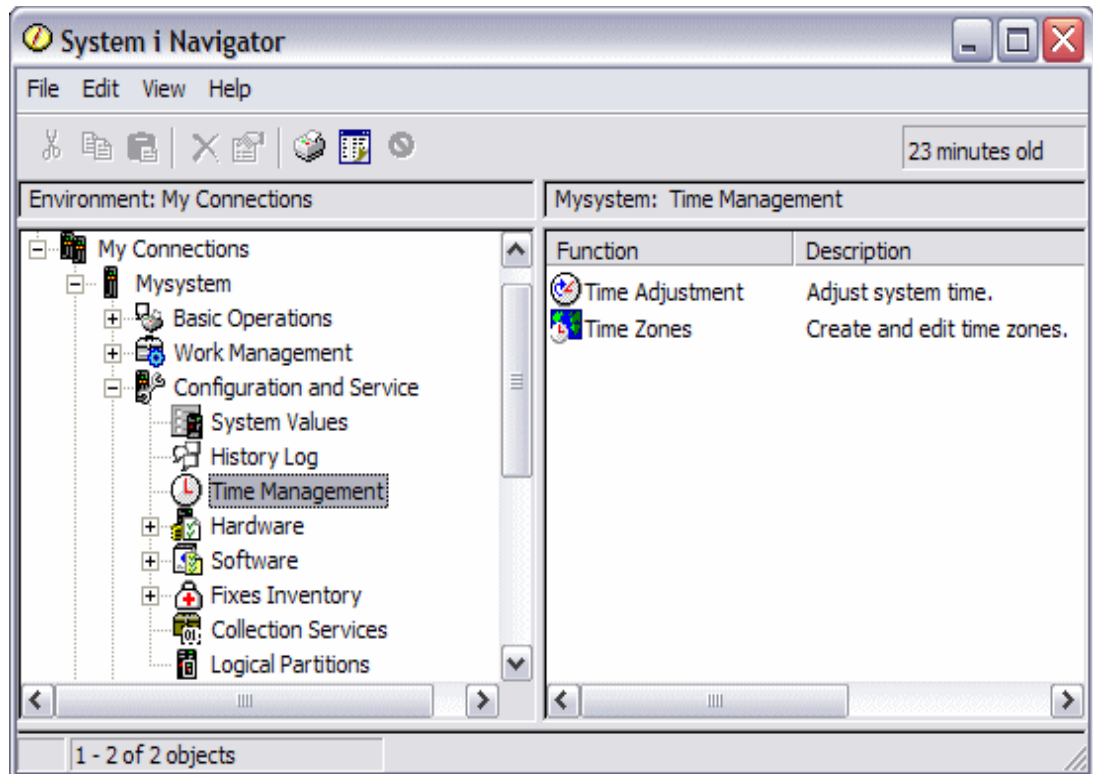


Figure 18-3 Time Management

With the time zones function, you can specify a time zone for your system to use. Many time zone descriptions are provided with the system and you can also create your own. Time zone descriptions include whether your system should observe daylight saving time (DST). If your system observes DST, the system automatically updates the system time when DST starts and ends.

The time adjustment function allows you to adjust the system time by manually starting a time adjustment. In addition, you can identify a time maintenance application to use to adjust the system clock to keep it synchronized with an external time source.

In addition, you can work with time management using IBM Systems Director Navigator for i, as shown in Figure 18-4.

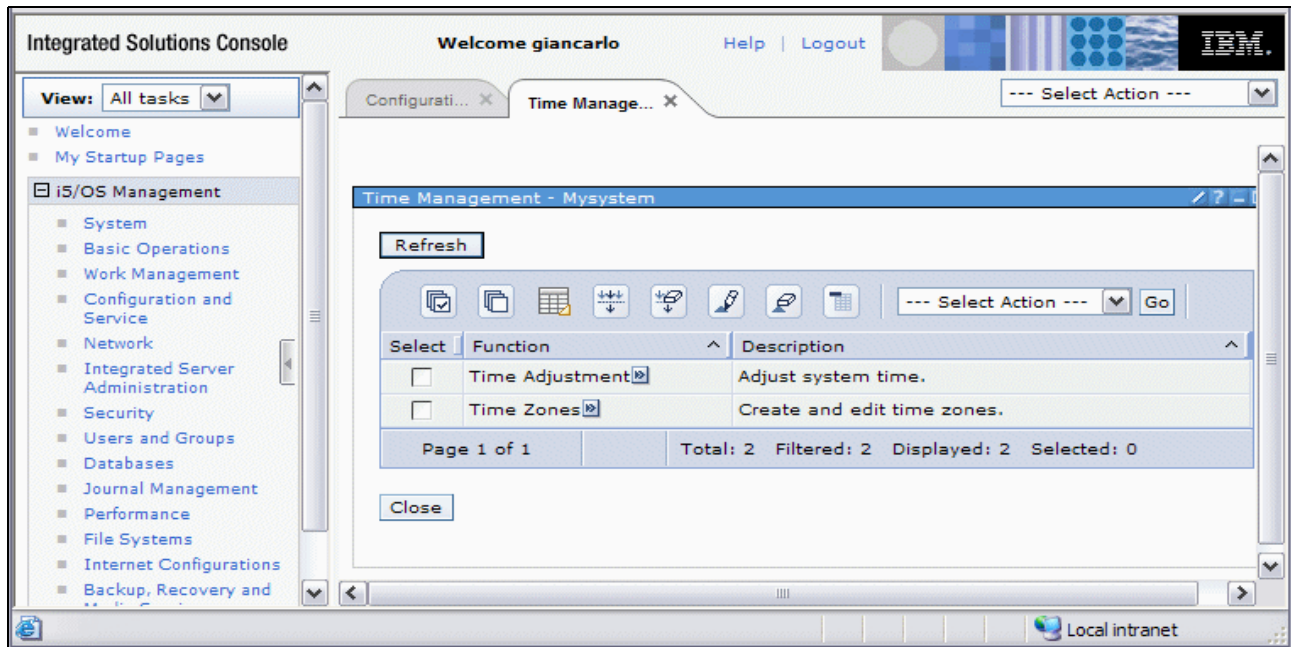


Figure 18-4 IBM Systems Director Navigator for i - Time management

18.9.1 New time zone values

The operating system provides several new time zones in this release.

You can specify a time zone for each system in your business environment. The time zone setting specifies the offset from Coordinated Universal Time (UTC) and whether to observe daylight saving time (DST). If you specify to observe DST, the system automatically updates the system time when DST begins and ends. You do not need to manually set the system time when DST begins and ends.

In a time zone description, the new *Year offset* value indicates the difference in years between the current Gregorian year and the current year in the calendar used with the time zone for the system. For systems observing a non-Gregorian calendar, this Year offset value is applied to the UTC year to calculate the local system year. For example, if the time zone description specifies a Year offset value of -57, then the local system year is the current Gregorian year minus 57.

In a time zone description, the new *Daylight saving time shift* value represents the number of minutes that local time moves forward when DST starts, or the number of minutes that local time moves backward when DST ends.

In a time zone description, the *Alternate name* value specifies the corresponding Olson database entry used by Java. You can specify a maximum of 128 characters for this value.

Note: The Year offset value, the Daylight saving time shift value, and the Alternate name value are not supported on systems running IBM i 5.4, or earlier.

18.9.2 New time zones

You can use one of the IBM-supplied time zones, or create your own time zones. If you want to edit one of the supplied time zones, you need to create a new time zone based on the

existing time zone. This prevents your changes from being lost when you install the operating system in the future. When the operating system is installed, the system-supplied time zones are overwritten. All IBM-supplied time zone identifiers begin with the character Q, and the full and abbreviated names of time zones are retrieved from the messages in the QCPFMSG message file. If you create a time zone (whether it is new or based on an existing time zone), the time zone identifier should not begin with the character Q. All time zone identifiers that begin with the character Q are reserved for IBM.

Use the time zone (QTIMZON) system value to specify the time zone for your system to use. To use the time zone system value, you must be using i5/OS V5R3, or later.

For detailed information about new time zones, refer to the time zone table in the IBM i Information Center, selecting the path **Systems Management** → **Time Management** → **Time Management concepts** → **Time zone**.

18.10 Operations Console

The Operations Console acts as a system console for you to access and administer your systems. IBM facilitates interaction with your systems by providing management consoles that can be accessed through terminals and PCs. The Operations Console is an installable component of the System i Access for Windows licensed program. Using Operations Console, you can access and control the console and control panel functions either locally or remotely through one or many PCs, which facilitates many administrative functions.

The following sections describe enhancements in IBM i 6.1 and also some information about upgrades.

18.10.1 Simplified console setup for network connected consoles

The initial console setup process is more automated in this release. During the setup, the system automatically detects some of your system settings, for example the system serial number or your IP address if the system is capable of providing it.

The traditional setup wizard is still available for use, if the system is not able to be contacted.

18.10.2 Auto-create service tools device IDs

By default, IBM i 6.1 sets the value of the Auto-create service tools device IDs option to 10. You can see and control this value through a Dedicated Service Tools (DST) session and Work with Service Tools Security Data – Option 11.

This default value gives you 11 connections to be automatically managed as QCONSOLE is an already IBM-supplied QCONSOLE default ID with up to 10 additional service tools device IDs that can be automatically created as needed.

These display as QCONSOLE00 or QCONSOLE01, through whatever value is set. We recommend you have at least one spare service tools device ID so you need not remember to increase the value if you try to add an additional connection. These service tools device IDs are listed in the Work with Service Tools Device IDs window so you can keep track of how many are created. The device ID QCONSOLE is already provided and can be used normally.

If you manually create a service tools device ID with a name of QCONSOLE02, for example, you are expected to assign that device ID to a specific PC workstation running Operations

Console. The automatic created device ID function can be used for new connections where the automatically created device ID is assigned to the new connection. If you have an environment of both manually created device IDs and automatically created device IDs, you are responsible for managing the IDs.

Note: You can change the default value of the *Autocreate service tools device IDs* option under DST. For example, you can turn it off.

If you upgraded from a prior release of Operations Console on the PC, your current IBM i configurations are retained and continue to work just as they did before. However, if the PC and system do get out of synchronization, instead of issuing an error message to resynchronize them, this function automatically attempts to recover the issue. You may still need to review the list of device IDs to ensure everything is correct.

When you create a local console on a network configured connection, you need to enter the name of the service interface (service host name.) The name you use depends on the circumstances for the connection. For example, if this is a new system and this is the first console to connect, then the name might be whatever you want to call it or it might be what a DHCP server gave it.

18.10.3 Automatic management of the access password

By default, you cannot set an access password when you create an Operations Console configuration. An access password is automatically generated and supplied to the connection process at the appropriate time. You do not need to know this password as it is automatically managed by the system. However, if you do want to change the access password, you can do this on the Properties page. Changing this password causes the automatic process of supplying the password to be disabled. You need to enter your access password for any connection that requires an access password.

The sign-on process requires a valid DST user ID and password to be entered to authenticate the user.

18.10.4 Automatic discovery of systems

When Operations Console starts on the PC workstation, a search for qualified systems begins. A qualified system is any system running IBM i 6.1. It does not need to have Operations Console as the console type.

If the PC gets a response from the network for any qualified systems, Operations Console creates a configuration for that system. It also preserves any pre-existing configured connections. A qualified system must have a value set for the Auto-create service tools device IDs option, which has the default value of 10 unless you have used DST to change it. If this value is zero, the system might not report correctly and therefore might not show up in this auto-generated list of systems. This auto-discovery function occurs any time Operations Console begins on the workstation or when performed on-demand by holding down the Alt key, the Shift key, and then pressing w. This causes Operations Console to begin a new search for eligible systems.

Note: If a system shows up in the list of configured connections, it does not guarantee that the system can be connected to. In rare cases, there are conditions in which the system is able to satisfy the requirements to report but cannot meet all requirements for a connection.

18.10.5 Planning your Operations Console installation or upgrade

If you are upgrading to IBM i 6.1 and want to replace an existing console with Operations Console, upgrade the system before migrating the console. This prevents any conflict between the existing console and the Operations Console. For instructions on upgrading your operating system, see the IBM i information center and select **i5/OS and related software** → **Installing, upgrading, or deleting i5/OS and related software** → **Upgrading or replacing i5/OS and related software**.

You must comply with the following before upgrading or installing your software (operating system, Licensed Internal Code) to IBM i 6.1:

- ▶ For all upgrades and installations, you need to establish a connection between the system and Operations Console PC using the 11111111 (eight 1's) service tools user ID. The default password for this user ID is 11111111; however, this password might have been changed following a previous installation. This default user ID ensures a successful re-authentication of the client connection to the system. When you receive the operating system release upgrade, the shipped service tools user IDs (except 11111111) are expired. To re-authenticate the client connection to the system, enter the service tools user ID of 11111111 (eight 1's) and either the default password of eight 1's or the password you might have previously created for this user ID. This is especially important for automatic installations.
- ▶ It is recommended that you update System i Access for Windows to V6R1 before you upgrade the operating system. For more information, see the "Installing System i Access for Windows" topic in the IBM i Information Center.

Note: Failure to comply with the above actions might prevent the console from working correctly during the upgrade or installation.

During a manual IPL of the system, and if no console has been specified before, you will receive two additional panels to confirm the setting of the console type. The first will require an F10 to accept your current console type and the second will show that a value did not previously exist (a zero will be present for the old value) and the new value will be shown. Pressing Enter will exit and set the console type automatically. The IPL will then continue to the IPL or Install the System panel. This condition is most likely to occur during the installation of a new partition but might happen on your first manual IPL of IBM i 6.1; for example, the A-mode IPL following the restore of Licensed Internal Code during the upgrade or install when a console value of zero is found.

18.10.6 Migrating to Operations Console before your server model upgrade

If you are going to be using Operations Console on your new server model (migrating from a different console type), it is important that you configure the new Operations Console PC before the beginning of the server model upgrade. At the point in the upgrade instructions where console functions are required on the new server model, you will be able to perform any required functions without the need for your current console device. The Operations Console features matching the connectivity you plan to use should be specified as part of the order for your new server model.

18.10.7 New models

If you are using a Power system model 5xx that uses a Hardware Management Console (HMC), you can change from HMC to Operations Console or from Operations Console to

HMC. For details on how to do this, see the “Managing consoles, interfaces, and terminals” and “Changing consoles, interfaces, and terminals” topics in the IBM Systems Hardware Information Center.

Embedded Ethernet ports on POWER5 models 5xx (except 595) will now be the default port to be used for Operations Console (LAN) on systems that are not HMC-managed. System i Model 825 is the only model using an embedded port for Operations Console.

18.10.8 Switching console control to another user

If you use Operations Console or HMC console and run a 5250 emulation console, you can use an option on the Select Console menu in DST or SST, called “Allow console recovery, and the console can be taken over by another console.” This option allows users of local consoles on a network to take control of the console from another user, if needed.

The “Allow console recovery and console can be taken over by another console” option also controls a new recovery function that allows the recovery of the console without the loss of data or the job. For details about what this option does, see the “Operations Console” topic in the IBM i Information Center.

18.11 Models not supported

As of IBM i 6.1, the following hardware models are no longer supported:

- ▶ 270
- ▶ 820
- ▶ 830
- ▶ 840

18.12 Common Information Model

The Common Information Model (CIM) is a standard that is developed by the Distributed Management Task Force (DMTF). DMTF is a consortium of major hardware and software vendors (including IBM) that is a part of the Web-Based Enterprise Management (WBEM) initiative.

Changes to CIM

CIM has moved out of the operating system to be ported with the IBM Universal Manageability Enablement for i5/OS licensed program (5722-UME). The CIM server runs in i5/OS Portable Application Solutions Environment (i5/OS PASE).

The following CIM functions are supported:

- ▶ The following Common Information Model Object Manager (CIMOM) functions are supported:
 - Secure Sockets Layer (SSL)
 - Common Manageability Programming Interface (CMPI)
 - Out-of-process provider
 - Service Location Protocol (SLP)
 - Distributed Management Task Force (DMTF) schema 2.14
 - Internet Protocol version 6 (IPv6)
 - Server Profile support (OpenPegasus start support Server Profile from 2.6.0)

- ▶ The following CIMOM functions were removed:
 - Kerberos support
 - Obsolete configuration properties
 - tempLocalAuthDir
 - enableHttpLocalconnection
 - kerberosServiceName
- ▶ The following CIMOM functions were changed to fixed properties:
 - home
 - daemon
 - slp
 - repositoryDir
- ▶ The -q option is no longer needed for the **cimconfig** and **cimprovider** commands in PASE. Some new properties were added to the basic startup properties and advanced startup properties for CIMOM. The **cimmof** and **ssltrustmgr** commands are available in this release.
- ▶ The following properties were changed to fixed properties, and you cannot change their values:
 - daemon
 - repositoryDir
- ▶ The following providers are supported:
 - Storage Management Initiative Specification (SMI-S) host bus adapter (HBA) and host discovered resources (HDR) CIM providers.
 - Providers that are inherited from the operating system.
 - Hardware inventory and network management providers.

You can use the IBM Universal Manageability Enablement for i5/OS licensed program (5722-UME) on i5/OS client partitions, with limitations. Physical hardware information is not accessible to i5/OS partitions that are managed by Integrated Virtualization Manager.

You can use IBM Systems Director Navigator for i to start and stop CIMOM.

CIMOM has a new configuration property, *idleSessionTimeout*.

For further information about any of these changes, refer to the “Common Information Model” topic in the IBM i Information Center.

If you upgrade from IBM i 5.4 to IBM i 6.1, the IBM Universal Manageability Enablement for i5/OS licensed program is installed by default. The CIM function that is supplied with IBM i 5.4 is disabled.

If you use the system-based CIM function in IBM i 5.4, you must start the CIM server before you configure CIMOM after the upgrade operation. The CIM server starts a repository and configuration migration process. During the migration, symbolic links are created; therefore, you cannot use the **cimconfig**, **cimmof**, and **cimprovider** commands.

Before the migration is completed, the CIM server is not available to process CIM requests; do not use any client programs.

The time the migration requires depends on the size of the repository, processor speed, and system utilization.

If you are already using the CIM function of the IBM Universal Manageability Enablement for i5/OS licensed program on your IBM i 5.4 operating system, you can continue with CIMOM configuration directly after you upgrade your system to IBM i 6.1.

Repository migration

The IBM Universal Manageability Enablement for i5/OS licensed program provides a repository that includes a set of files that contain the CIM class definitions, instances of classes, and provider registration instances. The repository is stored in the UserData directory (/QOpenSys/QIBM/UserData/UME/Pegasus).

The CIM repository is migrated from Distributed Management Task Force (DMTF) CIM schema V2.7 or V2.9 to DMTF CIM schema V2.14. During the migration, the CIM server is not available to process CIM requests. If you stop the server job during the migration, loss of data might occur.

The following objects are not migrated:

- ▶ CIM provider register information
- ▶ Static instances of metric definition in the repository (Providers dynamically collect information and implement the same functions as these metric instances.)

The CIM server log (by default, located at /QOpenSys/QIBM/UserData/UME/Pegasus/logs) might have the following messages related to the repository migration:

- ▶ Message PGS10100

This message is written into the CIM server log when the migration begins.

PGS10100: The CIM server is starting to restore the repository and then migrate the repository from an earlier version. This will take several minutes, during which the server will not be available. Stopping the server job might result in the loss of data.

- ▶ Message PGS10101

This message is written into the CIM server log when the migration is completed without errors.

PGS10101: The Common Information Model (CIM) repository has been migrated successfully.

Configuration migration

The CIM server migrates the configuration file of the operating system CIM server. The configuration properties are migrated and their values are not changed. However, the following configuration properties are not migrated:

- ▶ logdir, home, daemon, slp, repositoryDir, tempLocalAuthDir, and kerberosServiceName
- ▶ enableHttpsConnection, sslClientVerificationMode, and httpAuthType if the value of property httpAuthType in the operating system CIM server is Kerberos. Property enableHttpConnection is set to false after the migration.
- ▶ enableHttpLocalConnection

The migrated configuration properties are not checked for validity. If the configuration properties of the operating system CIM server are not set to function properly, this might prevent the new CIM server from starting and functioning successfully.



Upgrading IBM i

When upgrading IBM i from a previous release, there are special considerations for upgrading to IBM i 6.1. This chapter discusses these considerations. It also contains information about new features and enhancements to installing, upgrading, maintaining IBM i, and distributing software.

19.1 Installing or upgrading IBM i

The following sections describe special considerations when upgrading to IBM i 6.1.

19.1.1 Models supported

To determine whether your hardware is supported by IBM i 6.1, refer to the *Upgrade planning IBM i, i5/OS and OS/400 level mapping* Web page at:

<http://www.ibm.com/systems/support/i/planning/upgrade/osmapping.html>

19.1.2 Planning statements

Planning statements provide insight into IBM's current plans, directions, and intent, and are subject to change or withdrawal without notice. Statements pertaining to IBM i 6.1 can be found on the *Upgrade planning Planning statements* Web site, at:

<http://www.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>

19.1.3 Program Temporary Fixes for upgrading to IBM i 6.1

As you plan for your IBM i installation or upgrade, ensure that you review Information APAR II14310 for IBM i 6.1 in the preventive service planning information, selecting **Technical Databases** → **Authorized Program Analysis Reports (APARs)** at the following Web site:

<http://www.ibm.com/systems/support/i>

19.1.4 Pre-upgrade verification tool

The IBM pre-upgrade verification tool for IBM i runs on a Windows client and checks the IBM i data to ensure that all the necessary requirements are completed before beginning an upgrade. This tool verifies key pre-upgrade steps. Links to download the tool can be found in the "Optional: Running the IBM Pre-Upgrade Verification tool" topic in the IBM i Information Center.

19.1.5 Object conversions

To adequately prepare your system for the upgrade to the new release, you will need to allow extra time beforehand to analyze your system and possibly readjust your programs for conversions. The length of time required to analyze your system depends on the type, number, and size of your libraries and directories. You also need to plan in advance for possible down times.

Conversions occur during a software upgrade to IBM i 6.1 or during normal operations after the upgrade.

Note: If any of these preparations are normally done by a third party, contact your application or service provider for advice and assistance.

You will need to install planning PTFs for the release from which you are upgrading (either i5/OS V5R3 or IBM i 5.4) to assist you in the conversion process. These PTFs do not convert your objects but rather enable you to run the Analyze Object Conversion (ANZOBJCVN)

command to assist you in preparing for the upgrade. After you run the ANZOBJCVN command, you receive information to help you determine the effect that object conversions will have on your system, such as identifying potential conversion difficulties, if any, and supplying estimates of the time required for program conversion. The following IBM i Information Center topics provide additional information:

For an overview and time line on how the conversion tasks fit into the overall upgrade process, see “Software installation process.” Also review the information “Installing PTFs and analyzing your system for object conversions” and the topics under “Times for conversions” to determine how the conversions might affect your system.

19.1.6 Spooled file conversions

With IBM i 6.1, spooled file operations are processed more efficiently than in previous releases due to conversions that can be done either during the release upgrade or after the upgrade.

After the conversion, long-running operations, such as displaying a list of all the spooled files on the system, should not affect another job's ability to perform a spooling operation such as creating, deleting, or printing a spooled file.

Before you upgrade to the new release, it is recommended that you analyze your system before you upgrade by following the directions in Information APAR I114306 and *IBM i5/OS Program Conversion: Getting Ready for i5/OS V6R1*, REDP-4293. The information APAR is in the preventive service planning information about the *Support for IBM System i* Web page. Select the link for technical databases, then select the link for Authorized Problem Analysis Reports (APARs).

The planning PTFs for program object conversion that are described in the Redpaper™ publication include tools to collect and report the number of spooled files on your system and the time to convert them for 6.1. The times can vary depending on the workload of the system, and post-IPL or post-vary-on operation times could be longer than estimated.

Several options for converting your spooled files are available. The options are documented in the IBM i Information Center topic “Spooled file conversions.”

Note: Only newly created spooled files or spooled files that have been converted can be seen and used after the upgrade. Until the conversion is complete, unconverted spooled files appear to not exist. If a data area is used to direct the conversion, delete the data area after the conversion has occurred.

Spooled files restored to the IBM i 6.1 release are automatically converted. It is possible that the time for the spooled file conversion process can be lessened by saving and deleting the spooled files before you upgrade from IBM i 5.4 and then restoring them after you have 6.1 installed.

19.1.7 Licensed program releases and sizes

The size of Extended Base Directory Support (5761SS1 option 3) has increased to 2315 MB.

19.1.8 Media labels and their contents

Additional Fonts (5761SS1 option 43) and IBM Portable Utilities for i5/OS (5733SC1) are now on the B29xx_02 media.

19.1.9 Optical media for IBM i 6.1

The optical media for the IBM i 6.1 operating system is DVD. If you use physical media to install your IBM i 6.1 release, you must have a DVD device. Previous typical releases required approximately 13 CD-ROMs. The IBM i 6.1 release is available on three optical media. Some individual licensed programs are on CD-ROM. For a description of the 6.1 media, see “Media labels and their contents” on the IBM i Information Center.

19.1.10 Virtual tape devices for release upgrades

In addition to using a tape image catalog and a virtual tape device to install PTF fixes or single licensed programs, you can now also use these to perform an IBM i software upgrade.

19.1.11 Installing using virtual optical devices and remote virtual image files

This special type of virtual optical device supports virtual image files on a remote system in a network. An image directory identifies a network path on the central system that contains the virtual image files that are prepared for use with a target system. For more information, see “Virtual optical storage using the Network File System” in the IBM i Information Center.

19.1.12 Sharing optical resources

Starting in IBM i 6.1, optical resources from an IBM i logical partition can be used to perform a software upgrade, to install PTFs (fixes), or to install single licensed programs on another logical partition that is on the same system. Supported devices for sharing optical resources include CD-ROM devices, DVD devices, and virtual optical device types 632B-002 and 632B-003. All these devices are considered a 632C-002 device type to the logical partition that is used to install or upgrade. The 632C-002 device type requires that you manually change the media in the logical partition that owns the device.

19.1.13 Electronic software delivery

When you order IBM i licensed programs, you can select electronic software delivery as the method of delivery. If you select electronic software delivery, you receive information on how to access your ordered licensed programs from the Entitled software support Web site. Associated publications and other documentation are also available from the same Web site.

Instructions on how to download and prepare to install your software are available at the *Electronic Software Delivery help* Web site at:

http://www.ibm.com/systems/support/software/delivery/en_US/downloadinfo.html

19.1.14 Software keys available online

You can view, print, or download the software license keys from the *Entitled software support* Web site:

<http://www.ibm.com/eserver/ess/>

Printed copies of software keys are no longer sent with your i5/OS order. You need to have your customer number and either your order number or system number that is listed on your software packing list available. For further instructions on how to access your online software license keys, refer to the Customer Notice that is included with your software order.

19.2 Upgrading from i5/OS V5R3 to IBM i 6.1

When upgrading from i5/OS V5R3 to IBM i 6.1, other special considerations apply.

19.2.1 Requirements before you install IBM i 6.1 Licensed Internal Code

The IBM i 6.1 Licensed Internal Code requires more storage space on the load-source disk unit than i5/OS V5R3M0 Licensed Internal Code. (Your load-source disk unit has enough storage if you have a system with V5R3M5 Licensed Internal Code installed.) To upgrade to the IBM i 6.1 Licensed Internal Code, you must have a load-source disk unit of 17 GB or larger. Steps for determining if you have enough storage space are included in the IBM i Information Center topic “Preparing to upgrade or replace i5/OS software.”

In addition, all systems with i5/OS V5R3 installed require up to approximately 4 gigabytes of system ASP free space before IBM 6.1 can be installed. (A system with V5R3M5 Licensed Internal Code installed already has the 4 gigabytes of free space allocated.) Your upgrade will stop during the installation of Licensed Internal Code if you do not have this space. Steps for determining storage space are included in the IBM i Information Center topic “Determining storage space required for a software upgrade.”

19.2.2 Server firmware

The level of the server firmware needs to be managed. (This does not apply to System i 8xx models that support IBM i 6.1.) Server firmware is code that resides in system flash memory and includes a number of subcomponents, including power control code and the logical partition firmware that is loaded into AIX or Linux logical partitions. The server firmware is part of your machine code. If your system is not managed by a Hardware Management Console (HMC), the level of the server firmware that is on the service partition is used. A service partition is one partition on a system that must be set up to report hardware status and software errors to IBM. Installing a new level of the server firmware will require a server IPL to be performed. A server IPL is an initial program load where all logical partitions on a system are shut down at the same time and then restarted.

If your system is managed by an HMC, the level of the server firmware on the service partition is not used. The level of the server firmware is managed by the HMC. For more information about fixes and the HMC, refer to Systems Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp>

19.2.3 Operations Console

If you have an Operations Console and plan to upgrade it along with IBM i 6.1, refer to 18.10, “Operations Console” on page 459 for additional considerations.

19.2.4 Common Information Model

If you are using the Common Information Model support, refer to 18.12, “Common Information Model” on page 462 for additional considerations.

19.3 Distributing software

The following sections describe special considerations for distributing software.

19.3.1 Using virtual tape drives

In addition to being able to use virtual optical images to distribute software using central site distribution, you can use virtual tape devices. Instructions are included in the “Optional: Preparing the central system for virtual tape images” topic in the IBM i Information Center.

19.3.2 Sharing optical resources

The virtual optical device support of i5/OS provides a method for using virtual optical images from a single image source that can be shared with other systems in a network. A client with the virtual optical device type 632B model 003 can access images located on a server that uses the Network File System (NFS). See the IBM i Information Center topic “Distributing additional licensed programs or fixes by using the Network File System” for more information.

19.3.3 Duplicating virtual images to physical media

You can use the Duplicate Optical (DUPOPT) command to duplicate virtual images to physical DVD or CD-RW media. Refer to the topic “Optional: Duplicating virtual images to physical media.”

19.3.4 Using image catalogues

A copy of the image catalog information is now stored in the image catalog directory in a file called QIMGCLG. This makes the software distribution and catalog creation after moving the image files between systems easier. The new *IMPORT parameter of the Create Image Catalog (CRTIMGCLG) command allows a catalog to be created based on existing catalog information that is located in the QIMGCLG file. Catalog entries will appear as they did in the previous catalog. Steps for this are included in “Optional: Preparing the target system for virtual images.”

19.4 Program Temporary Fix enhancements

This section describes new features of and enhancements to the Program Temporary Fix group (PTF group) and Program Temporary Fix (PTF) functions.

19.4.1 Cumulative Program Temporary Fix media

Cumulative PTF packages will now be delivered on DVD-ROM media.

19.4.2 Multiple level support for groups

IBM i now supports multiple levels for each Program Temporary Fix (PTF) group on the system. A new parameter (PTF group level) of the Display Service Attributes (DSPSRVA) and Change Service Attributes (CHGSRVA) commands allows you to specify how many levels of a PTF group to maintain on a system. The range of values for this parameter is 1-99999. The default, however, is 2, which indicates that the system keeps the previous level of a group when a new PTF group is downloaded or installed. PTF group level information is included with the Copy PTF Group (CPYPTFGRP) and Work with PTF Group (WRKPTFGRP) commands, and the PTF Group APIs.

19.4.3 Group status

Three new status values have been added for PTF groups:

Apply at next IPL	One or more PTFs in the group are set to be applied at the next IPL. All other PTFs for the installed products on the system are already temporarily or permanently applied.
Related group	The PTF group does not have any PTFs for licensed programs installed or supported on the system. However, the PTF group is identified in another PTF group as a related PTF group. If you delete a PTF group in this status, the dependent PTF group shows a status of unknown.
On order	At least one PTF in the group is on order and has not yet been installed on the system. This PTF is delivered on either physical or virtual media.

19.4.4 Target release field

Also new for PTF groups is a target release field that indicates whether a PTF group is associated with a specific i5/OS release. A PTF group is automatically deleted if it does not contain a target release or if the target release is less than the current release and the PTF group does not contain PTFs for any installed or supported licensed programs. PTF groups are automatically deleted during i5/OS install and during Operational Assist (OA) cleanup.

19.4.5 Work with Program Temporary Fix Groups command

A new option on the Work with PTF Groups command is 1 (Order) that allows you to order a new level of the PTF group.

A new function key (F19=Display PTFs not installed) is provided on the Display PTF group panel of the Work with PTF groups (WRKPTFGRP) command. The function key lists only the PTFs in the PTF group that are not installed when the status of the PTF group is "Not installed."

19.4.6 Compare and Update wizard

The Compare and Update wizard of System i Navigator now updates fix group levels in addition to individual fixes.

19.4.7 Display Program Temporary Fix command

A new special value *NOTAPY is provided on the Display PTF (DSPPTF) command to determine which PTFs are not installed. Use the DSPPTF SELECT(*NOTAPY) option to display all PTFs that are not already applied or superseded for the specified product.

19.4.8 Program Temporary Fix orders

The Work with PTF Order (WRKPTFORD) command is added to help you manage program temporary fixes (PTFs).

The Work with PTF Order (WRKPTFORD) command shows program temporary fix (PTF) orders submitted by the Send PTF Order (SNDPTFORD) command or Control PTF Order (QESCPTFO) API. From the Work with PTF Order panel you can do the following:

- ▶ Download an image order submitted
- ▶ Cancel an image order submitted
- ▶ Display a PTF order

Using the Status type parameter, you can choose one of the following options to display all, or a subset of, PTF orders:

*ALL	All types of status of PTF order are shown.
*ONORDER	The PTF order has been submitted to the Fix Service Provider. The download process is the next action to do.
*DOWNLOADED	The PTF order has been submitted and downloaded.
*CANCELED	The PTF order has been canceled.
*SCHEDULED	The PTF order has been submitted to the Fix Service Provider and a job has been scheduled to download the order.
*FAILED	The PTF order has been failed.
*MAILED	The PTF order has been mailed.



IVM and HMC feature summary

This chapter discusses Integrated Virtualization Manager (IVM) and Hardware Management Console (HMC) features.

20.1 IVM and HMC comparison at a glance

Integrated Virtualization Manager (IVM) and Hardware Management Console (HMC) are both hardware management solutions.

The primary hardware management solution IBM has developed relies on an appliance server (packaged as an external tower or rack-mounted server) called Hardware Management Console (HMC). IVM performs a subset of the HMC features for a single server (IVM can manage only standalone servers), so each managed server must have its own instance of the IVM. HMC is a centralized point of hardware control and it can manage partitioned servers. A single HMC can handle multiple POWER5 and/or POWER6 systems, and two HMCs may manage the same set of servers in a dual-active configuration providing resilience. IVM can be more suitable for smaller and distributed environments (multisystem environments), where dedicated hardware and a full set of HMC features is not needed.

IVM provides a Web-based system management interface and a command-line interface that you can use to manage and configure logical partitions on Power systems (including IBM Virtual I/O Server), without the need for dedicated hardware, such as HMC. On the managed system, you can create logical partitions, manage the virtual storage and virtual Ethernet, and view service information related to the server.

The Integrated Virtualization Manager (IVM) is packaged with the Virtual I/O Server (VIOS), which services I/O and processor virtualization on Power systems.

IVM provides virtual console support for AIX and Linux partitions. HMC provides virtual console support for AIX and Linux logical partitions and virtual 5250 console emulation for IBM i logical partitions.

Using HMC it is possible to perform logical partitioning on Power systems (including IBM Virtual I/O Server) using either the Web browser-based user interface or the command line interface (CLI). HMC provides service functions on the managed system, interacting with the service processor, and is able to create, manage, modify logical partitions, modify the hardware configuration (including virtual hardware resources) of the managed system, and manage service calls.

For further information about HMC, refer to *Hardware Management Console V7 Handbook*, SG24-74911, available at:

<http://www.redbooks.ibm.com/abstracts/sg247491.html?Open>

For further information about IVM, refer to the *Integrated Virtualization Manager on IBM System p5* document, available at:

<http://www.redbooks.ibm.com/abstracts/redp4061.html?Open>

Table 20-1 on page 475 provides some comparison between IVM and HMC.

20.2 IVM and HMC feature summary

Table 20-1 IVM and HMC comparison at a glance

	Integrated Virtualization Manager (IVM)	Hardware Management Console (HMC)
Physical footprint	Integrated into the server	A desktop or rack-mounted appliance
Installation	Installed with the VIOS (optical or network). Preinstall option available on some systems.	Appliance is preinstalled. Reinstall using optical media or network is supported.
Managed operating systems supported	AIX 5L and Linux	AIX 5L, Linux, and i5/OS
Virtual console support	AIX 5L and Linux virtual console support	AIX 5L, Linux, and i5/OS virtual console support
User security	Password authentication with support for either full or ready-only authorities	Password authentication with granular control of task-based authorities and object-based authorities
Network security	<ul style="list-style-type: none"> ▶ Firewall support by command line ▶ Web server SSL support 	<ul style="list-style-type: none"> ▶ Integrated firewall ▶ SSL support for clients and for communications with managed systems
Servers supported	System p5® 505 and 505Q Express System p5 510 and 510Q Express System p5 520 and 520Q Express System p5 550 and 550Q Express System p5 560Q Express eServer p5 510 and 510 Express eServer p5 520 and 520 Express eServer p5 550 and 550 Express OpenPower® 710 and 720 BladeCenter JS21	All POWER5 and POWER5+ Processor-based servers: System p5 and System p5 Express eServer p5 and eServer p5 Express OpenPower eServer i5
Multiple system support	One IVM per server	One HMC can manage multiple servers
Redundancy	One IVM per server	Multiple HMCs can manage the same system for HMC redundancy
Maximum number of partitions supported	Firmware maximum	Firmware maximum
Uncapped partition support	Yes	Yes

	Integrated Virtualization Manager (IVM)	Hardware Management Console (HMC)
Dynamic Resource Movement (dynamic LPAR)	System p5 support for processing & memory BladeCenter JS21 only support for processing	Yes - full support
I/O support for AIX 5L and Linux	Virtual optical, disk, Ethernet, and console	Virtual and Direct
I/O support for IBM i	None	Virtual and Direct
Maximum # of virtual LANs	Four	4096
Fix/update process for Manager	VIOS fixes and updates	HMC e-fixes and release updates
Adapter microcode updates	Inventory scout	Inventory scout
Firmware updates	VIOS firmware update tools (not concurrent)	Service Focal Point™ with concurrent firmware updates
I/O concurrent maintenance	VIOS support for slot and device level concurrent maintenance through the diag hot plug support	Guided support in the Repair and Verify function on the HMC
Scripting and automation	VIOS command line interface (CLI) and HMC-compatible CLI	HMC command line interface
Capacity on Demand	No support	Full support
User interface	Web browser (no local graphical display)	WebSM (local or remote)
Workload Management (WLM) groups supported	One	254
LPAR configuration data backup and restore	Yes	Yes
Support for multiple profiles per partition	No	Yes
Serviceable event management	Service Focal Point Light: Consolidated management of firmware and management of partition detected errors	Service Focal Point support for consolidated management of operating system and firmware detected errors
Hypervisor and service processor dump support	Dump collection with support to do manual dump downloads	Dump collection and call home support
Remote support	No remote support connectivity	Full remote support for the HMC and connectivity for firmware remote support

Note: The IBM Systems Director 6.1 provides a systems management simplification, that is, it provides a Web console interface that integrates management consoles, supporting multiple operating systems and virtualization technologies, including IBM i. IBM Systems Director 6.1 enables capabilities for managing IBM i in the same console with the HMC, IVM, and functions for managing other Power platforms. For further information, refer to:

<http://www.ibm.com/systems/management/director/>

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 publications, see “How to get Redbooks” on page 481. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM System i5 V5R4 Technical Overview Redbook*, SG24-7271
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694
- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ *Windows-based Single Signon and the EIM Framework on the IBM eServer iSeries Server*, SG24-6975
- ▶ *IBM System i Security: Protecting i5/OS Data with Encryption*, SG24-7399
- ▶ *IBM i5/OS Program Conversion: Getting Ready for i5/OS V6R1*, REDP-4293
- ▶ *Hierarchical Storage Management*, SC41-5351
- ▶ *Implementing PowerHA for IBM i*, SG24-7405
- ▶ *PowerVM Virtualization on IBM System p: Introduction and Configuration Fourth Edition*, SG24-7940
- ▶ *VMware VI3 on BladeCenter and System x Integrated with System i*, SG24-7408
- ▶ *i5/OS Diagnostic Tools for System Administrators: An A to Z Reference for Problem Determination*, SG24-8523
- ▶ *Hardware Management Console V7 Handbook*, SG24-7491
- ▶ *End to end Performance Management on IBM i*, SG24-7808

Online resources

These Web sites are also relevant as further information sources:

- ▶ the IBM i Information Center for IBM i 6.1
<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>
- ▶ The *i5/OS V6R1 Technical Overview* presentation set, located in the Support for IBM i Web site
<http://www.ibm.com/systems/support/i/news/v6r1techoverview.html>
- ▶ The IBM Systems Software Information Center
<http://publib.boulder.ibm.com/infocenter/systems/index.jsp>
- ▶ The IBM techdoc *Stored Procedures Restore Overview*
http://www-01.ibm.com/support/docview.wss?rs=0&dc=DB520&dc=D900&dc=D800&dc=DA900&dc=DA800&q1=Restoring+External+Stored+AND+AS400KBXXYYZZRCH&uid=nas10bf6be811065714f8625707e0073a26b&loc=en_US&cs=UTF-8&lang=all

- ▶ The IBM document *i5/OS on POWER Blade Read-me First*
http://www.ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf
- ▶ The iDoctor web site
https://www-912.ibm.com/i_dir/iDoctor.nsf
- ▶ The IBM Systems Workload Estimator (a.k.a., the Estimator or WLE) web-based sizing tool
<http://www.ibm.com/systems/support/tools/estimator/>
- ▶ The IBM Systems Energy Estimator
<http://www.ibm.com/systems/support/tools/estimator/energy/>
- ▶ Searching Sizing Guides web page
<http://ibm.com/servers/sizing>
- ▶ *The IBM i Virtualization and Open Storage Read-me First* document
http://www.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf
- ▶ *Power Systems Logical partitioning* document available at:
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>
- ▶ *IBM i integration with BladeCenter and System x* at:
<http://www.ibm.com/systems/i/advantages/integratedserver/library.html>
- ▶ *IBM Power Blade servers* at:
<http://www.ibm.com/systems/power/hardware/blades/index.html>
- ▶ *Linux on integrated servers* Web site at:
<http://www.ibm.com/systems/i/advantages/integratedserver/linux/>
- ▶ *Advanced Job Scheduler Command List Changes* Web page at:
http://www.ibm.com/systems/i/software/jscheduler/command_list.html
- ▶ The Open Group
<http://www.opengroup.org>
- ▶ The *JTOpen* Web site
<http://jt400.sourceforge.net>
- ▶ You can also refer to IBM Rational Developer for IBM i web page at:
<http://www.ibm.com/software/awdtools/developer/rdi/>
- ▶ The IBM Enterprise Modernization Demos at:
http://rational.dfw.ibm.com/atdemo/atdemo_wdht_recorded.html
- ▶ The System i Access for Linux homepage
<http://www.ibm.com/systems/i/software/access/linux/>
- ▶ *The Upgrade planning IBM i, i5/OS and OS/400 level mapping* Web page
<http://www.ibm.com/systems/support/i/planning/upgrade/osmapping.html>
- ▶ *The Upgrade planning Planning statements* Web site
<http://www.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, 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

*SYSTEM certificate store 316

Numerics

3Com telephony workload 192
5761-SS1 Option 45 - Encrypted ASP Enablement 12
62448
 Head 2
 2.4 Changed security system values 18
7463 (processor feature) 294
98182
 Head 2
 2.10 digital certificate manager 24

A

ACK storms 6
adapter microcode updates 476
Add Disk Watcher Definition (ADDDWDFN) 135
Add Job Watcher Definition (ADDJWDFN) 135
Add Kerberos Keytab Entry (ADDKRBKTE) 25
Add Kerberos Ticket (ADDKRBTKT) 25
additional fonts 468
additional settings (advisor) 303–304
address poisoning 6
Address Resolution Protocol (ARP) 222
 networking 222
Advanced Encryption Standard (AES) 11
advanced program-to-program communications (APPC)
 networking 233
advise window for Threads to process 300
announcement letter 3
Apache Xerces XML 330
appliance server 474
application development tools 1
Application Development ToolSet (ADTS) 338
application diagram 374
Application Response Measurement 287
ARM 287
associated software
 IBM i 2
availability
 cluster administrative domain 76, 82
 cluster resource groups (CRG)-level messages 81
 cluster-level messages 81
 PowerHA for i 76

B

Backup Recovery and Media Services 2
Base events (BASEVT) 181
basic operations
 IBM Systems Director Navigator for i 404
Berkeley Internet Name Domain (BIND) version 9

 networking 224
BIND 4.9.3
 networking 224
BIND 8.2.5
 networking 224
Block password change (PWDCHGBLK) 19
Block password change (QPWDCHGBLK)
 system value 16
BRMS 6.1 enhancements
 System i Navigator 438
business applications 76

C

capacity on demand 476
certificate store 316
Change Data Library using JS (CHGDLJS) 273
Change Kerberos Password (CHGKRBPWD) 25
CHRLMTAJC
 system values 13
CL Support
 application diagram 365
classic JVM 127
Close Database File (CLOSE) 342
cluster 76
cluster administrative domain 354
 availability 76, 82
cluster resource group 77
cluster resource groups (CRG)-level messages
 availability 81
cluster resource services
 IBM Systems Director Navigator for i 405
clustering 2
cluster-level messages
 availability 81
Cobol application development 362
collection services 137, 148, 312
collection services and Java 129
Command Language (CL) 273
command-line interface 474
commands
 RMVTCPTBL 11
commands moved to IBM i
 performance 137
Common Cryptographic Architecture for IBM i
(IBMJCECCAI5OS) 319
compatibility 330
Components for Unicode (ICU) 359
configuration and service
 IBM Systems Director Navigator for i 404
Configure Perf Collection (CFGPFRCOL) 183
console
 Systems Director Navigator for i 409
Control Language (CL) 342
Convert Performance Collection (CVTPFRCOL) 137

- Copy Performance Data (CPYPRDTA) 137
- copy services 76
- Create Bound C Program (CRTBNDC) 338
- Create Bound C++ Program (CRTBNDCPP) 338
- Create Bound CL Program (CRTBNDCCL) 342
- Create C Module (CRTCMOD) 338
- Create C++ Module (CRTCPPMOD) 338
- Create CL Module (CRTCLMOD) 342
- Create CL Program (CRTCLPGM) 342
- Create Command (CRTCMD) 343
- Create Java Program (CRTJVAPGM) 318–319
- Create Performance Data (CRTPFRDPA) 182
- Create Performance Summary (CRTPFRSUM) 137, 184
- Create Remote Name Daemon Control Configuration (CRTRNDCCFG) 224
- CRG 77
- cross site mirroring environment 11
- cross-site mirroring (XSM) 76
- cryptographic keys 11
- cryptographic services
 - APIs 358

D

- data source 304
- database 91
 - Materialized query table (MQT) 92
 - National Language Sort Sequences (NLSS) 90
 - plan cache 93
- database performance 90
- database support 443
- DatabaseMetaData interface 321
- databases
 - IBM Systems Director Navigator for i 405
- DB2 for i catalog views 101
- DB2 runtime engine 91
- DB2 Web Query 2
- DDM
 - Decimal floating-point data type 99
- debug 366
- Decimal floating support 337
- Decimal floating-point data type
 - DDM 99
 - DRDA 96
- Declare Processing Options (DCLPRCOPT) 343
- dedicated processors sharing 201
- Delete Kerberos Credentials Cache File (DLTKRBCCF) 25
- Delete Performance Data (DLTPFRDPA) 137
- derived key index 94
- Detection System graphical user interface (GUI) 6
- Digital Certificate Manager 24
- Direct Execution (DE) 127
- Disable User Profile (DSBUSRPRF) 262
- disk adapters 192
- disk encryption 11
- disk pools 11
- Disk Storage System sizing
 - support for DS500 and XIV 186
- Disk Watcher 136
- Display Garbage Collection Information 132

- Display Kerberos Credentials Cache File (DSPKRBBCCF) 25
- Display Kerberos Keytab Entries (DSPKRBKTE) 25
- Display Log for Job Scheduler (DSPLOGJS) 274
- Display Service Tools User ID (DSPSSTUSR) 21
- Display User Profile (DSPUSRPRF) 20
- distributing software 470
- Distributed Data Management (DDM) 99
- Distributed Relational Database Architecture (DRDA) 96
- DLPAR 201
- DNS Configuration Utility (CHKDNSCFG) 224
- DNS Zone Utility (CHKDNSZNE) 225
- Domain Information Groper (DIG)
 - networking 225
- Domain Name System (DNS)
 - networking 224
- DRDA
 - Decimal floating-point data type 96
 - Extended indicators 98
 - Heterogeneous scrollable cursors 97
 - Internet Protocol version 6 (IPv6) 98
 - Optimistic locking 98
 - Set Client Information (SQLESETI) API 98
 - SQL Limits 97
 - SQL statement interruption 97
- Dump Java Virtual Machine 128
- Dynamic Logical Partitioning 196, 201
- dynamic LPAR 201
- Dynamic Resource Movement (dynamic LPAR) 476
- Dynamic Update Utility (NSUPDATE) 226

E

- enhancements, IBM i 2
- encryption
 - enhanced functions 2
- End Disk Watcher (ENDDW) 135
- End Job Watcher (ENDJW) 135
- End Performance Explorer (ENDPEX) 137
- End Program (ENDPGM) 342
- energy estimation 186
- Enterprise Edition Connector (J2C) 364
- Enterprise Generation Language (EGL) 362, 382
- Enterprise Identity Mapping (EIM) 24
- EWLM 288
- expanded Web enablement 3
- export or import the performance profile (Web Performance Advisor) 293
- Extended indicators
 - DRDA 98
- Extended Passive (EPSV) 227
- Extended Port (EPRT) 227
- extrusions
 - security 6

F

- features 406
- file systems
 - IBM Systems Director Navigator for i 405
- firmware updates 476

- fix/update process for Manager 476
- flash copy 77
- FlashCopy 76
- Forwarding mailhub server (FWDHUBSVR) 229
- Fraggle attack 6
- Full outer join 95
- full outer join 95
- Functional Zoom
 - application diagram 365

G

- Generate JVM Dump 133
- Generate JVM Dump (GENJVMDMP) 321
- GENJVMDMP 133
- geographic mirroring 77
- Global Mirror 76
- global mirror 78
- Go Licensed Program (GO LICPGM) 324
- Graphical user interface (GUI) 139

H

- hardware 2
- hardware encryption
 - security 12
- hardware management 440
- Hardware Management Console 297
- Hardware Management Console (HMC) 473
- hardware management solutions 474
- HASM 2
- HATS
 - WebFacing 395
- HATS interoperability 397
- HATS V7.5 enhancements 398
- health center 436
- Heap Analyzer 141
- heap dump 133
- heap size 128
- Heterogeneous scrollable cursors
 - DRDA 97
- Hierarchical Storage Management
 - Redbooks 66
- high availability 2
- high availability product 76
- HMC 297
- Host Server Variable Length compression 92
- HTTP server 288, 291
- Hypervisor 196
- Hypervisor and service processor dump support 476

I

- I/O concurrent maintenance 476
- I/O queuing 159
- I/O Support for AIX 5L and Linux 476
- I/O Support for IBM i 476
- i5/OS
 - nomenclature 2
- i5/OS Diagnostic Tools for System Administrators
 - An A to Z Reference for Problem Determination

- Redbook 455
- i5/OS projects
 - development 372
- i5/OS Tasks 286
- i5/OS V6R1 1
- i5/OS V6R1 Technical Overview 1, 479
- IBM Advanced Job Scheduler for i5/OS (5761-JS1) 266
- IBM High Availability Solutions Manager (HASM) 2
- IBM i
 - associated software 2
 - installing or upgrading 466
- IBM i 6.1
 - IBM i 6.1 191
 - new functions 2
 - overview 2
 - POWER5 and POWER6 technology systems 2
- IBM i 6.1 Licensed Internal Code
 - requirements 469
- IBM i applications 3
- IBM i database
 - frequently asked questions 352
- IBM i Debugger 338
- IBM i Doctor
 - Preferences 141
- IBM i Information Center 1, 479
- IBM i Information Center for IBM i 6.1 Web 1
- IBM i performance data 3
- IBM i program call Java Platform 2 364
- IBM i RPG 362
- IBM i Run Priority (RUNPTY) 340
- IBM i storage
 - virtualization 2
- IBM i support on the POWER6 Blade 191
- IBM i system builder 2
- IBM i terminology 2
- IBM i users and groups
 - IBM Systems Director Navigator for i 435
- IBM i workloads 192
- IBM Java Cryptography Extension 319
- IBM Rational Developer for IBM i (RDi) 364
- IBM Rational Team Concert for i (RTCi) 400
- IBM Sizing Guides 190
- IBM System i5 V5R4 Technical Overview
 - Redbooks 1
- IBM Systems Director for Linux on System z, 6.1 404
- IBM Systems Director for Power Systems V6.1 3
- IBM Systems Director for Power Systems, V6.1 (AIX, IBM i, Windows, Linux) 404
- IBM Systems Director for x86, V6.1 404
- IBM Systems Director Navigator for i 6, 404–405
 - IBM Systems Director 6.1 3
 - network interfaces 220
 - Web based systems management interface to IBM i 3
- IBM technology for Java 126
- IBM Technology for Java (32-bit) 127
- IBM Technology for Java (64-bit) 127
- IBM technology for Java 32-bit 128
- IBM technology for Java 64-bit 128
- IBM Tivoli Directory Server 231
- IBM Tivoli Directory Server for i5/OS (LDAP) 230

- IBM Toolbox for Java 330
 - compatibility 330
- IBM WebSphere application development tools 2
- ICU 2.6.1 Sort Sequences 90
- identity mapping (EIM) 364
- ILE C
 - compile 337
- ILE C compiler 337
- ILE C++
 - compiler 337
- ILE COBOL
 - compiler 336
- image catalogues 470
- Implementing IBM Systems Director 6.1
 - Redbooks 3
- Implementing PowerHA for IBM i
 - Redbooks 80
- import customer performance data 187
- Include CL Source (INCLUDE) 342
- Incoming Remote Command (IRC) 273
- independent disk pool 11
- Information Center 4
- Install Window Server (INSWNTSVR) 263
- installation and set up information
 - System i Access for Windows 441
- installing or upgrading
 - IBM i 466
 - planning statements 466
- Integrated File System 436
- Integrated IBM i Debugger 377
- integrated server administration
 - IBM Systems Director Navigator for i 405
- Integrated Virtualization Manager 196
- Integrated Virtualization Manager (IVM) 473
- Integrated Virtualization Manager on IBM System p5 document 474
- Interior Gateway Protocol (IGP)
 - networking 223
- internet configurations
 - IBM Systems Director Navigator for i 405
- Internet Protocol version 6 (IPv6) 327, 430
 - DRDA 98
- interval details 143
- intrusion detection
 - security 6
- Intrusion Detection Management panel 7
- intrusion events 6
- intrusion monitor 6
- ioctl function 340
- IPv4
 - networking 218
- IPv6
 - networking 218
- IVM 196, 204

J

- J9 126
- Java 2 Platform 320
- Java 32-bit 2
- Java Development Kits (JDKs) 322

- Java dump 133
- Java JIT (Just In Time) compiler 318
- Java pPerformance 129
- Java Secure Socket Extension (JSSE) 320
- Java Secure Socket Extension 6 (JSSE 6) 320
- Java virtual machine 126
- Java Virtual Machine Debugger Interface (JVMDI) 319
- Java Virtual Machine Profiler Interface (JVMPPI) 319
- Java Virtual Machine Tool Interface (JVMTI) 319
- Java Virtual Machines (JVMs) 320
- JDBC provider 304
- JDBC Resources tab (advisor) 303
- Job name (JOB) 181
- Job watcher 136
- Job Watcher Wizard 143
- journal management
 - IBM Systems Director Navigator for i 405
- journaling 2
- JVM
 - choosing which to use 128
- JVM settings Advise examples 302

K

- Key messages 4

L

- LCA password 316
- ldapmodify
 - networking 230
- ldapmodrdrn 231
- License and Web standard 186
- license and Web standard 186
- Licensed Internal Code (LIC)
 - security 23
- licensed programs 1
- Limit device sessions (LMTDEVSSN) 20
 - system value 18
- link-state advertisements (LSAs)
 - networking 223
- Live Partition Mobility 196
 - logical partitioning 196
- Local Certificate Authority (LCA) 315
- Lock events (LCKEVT) 181
- lock holders 117
- logical partitioning 196
 - Live Partition Mobility 196
 - Multiple Shared Processor Pools 196
 - POWER Hypervisor 197
- logical partitions 474
- low-cost solution
 - software encryption 13
- LPAR configuration data backup and restore 476

M

- Machine Interface (MI) 357
- major release
 - overview 2
- manage application server and HTTP server attributes

- (Web Performance Advisor) 293
- Manage attributes 299
- manage system attributes (Web Performance Advisor) 293
- managed operating systems supported 475
- Management Central 439
- Management Central or the IBM i 266
- Materialized query table (MQT)
 - database 92
- materialized query table refresh 92
- maximum # of virtual LANs 476
- maximum number of partitions supported 475
- Maximum storage to use (MAXSTG) 181
- memory 294
- Metro Mirror 76, 78
- micro-partitioning 196
- miscellaneous
 - APIs 356
- multi-path disk units 159
- multiplatform suite
 - IBM Systems Director 3
- multiple scheduling environment 273
- multiple scheduling environments 273
- Multiple Shared Processor Pools
 - logical partitioning 196
- multiple shared processor pools 196
- multiple system support 475
- multiple threads 230

N

- N_Port ID Virtualization 196
- name changes 139
- National Language Sort Sequences (NLSS) 90
 - database 90
- Native Abstract Windowing Toolkit (NAWT) 327
- Netserver 226
- network
 - IBM Systems Director Navigator for i 404
- network security 475
- networking 222
 - advanced program-to-program communications (AP-PC) 233
 - Berkeley Internet Name Domain (BIND) version 9
 - networking 224
 - BIND 4.9.3 224
 - BIND 8.2.5 224
 - Domain Information Groper (DIG) 225
 - Domain Name System (DNS) 224
 - Interior Gateway Protocol (IGP) 223
 - IPv4 218
 - IPv6 218
 - ldapmodify 230
 - link-state advertisements (LSAs) 223
 - Open Shortest Path First (OSPF) 223
 - Point-to-Point Protocol (PPP) 222
 - Post Office Protocol (POP) 229
 - Request for Comments (RFC) 2136 226
 - Secure e-mail (S/MIME) 229
 - Secure Sockets Layer (SSL) 230
 - Shortest Path First (SPF) algorithm 223

- Simple Mail Transfer Protocol (SMTP) 228
- Virtual IP 222
- Virtual Private Network (VPN) 227
- networking functions 432
- new and enhanced
 - IBM Power System models 186
- new components 139
- new database monitor format 96
- new disk models 191
- new IBM Power Systems 189
- new IBM system p POWER6 models 191
- new IBM system x models 191
- new SAS storage 189
- nomenclature
 - i5/OS 2
- non-SQL Interfaces 90
- number of disk units 294

O

- object conversions 466
- OLE DB enhancements 430
- On-Line Analytical Processing (OLAP) 349
- Open Shortest Path First (OSPF)
 - networking 223
- Operating System events (OSEVT) 181
- optical media 468
- Optimistic locking
 - DRDA 98
- optimization time 91
 - database 91
- OS/400 V5R2 23
- output queue monitor
 - scheduled job 271
- overview
 - IBM i 6.1 2
- overview information 1

P

- password encryption 230
- Password expiration warning (QPWDEXPWRN)
 - system value 16
- Password rules (QPWDRULES)
 - system value 13
- performance
 - commands moved to IBM i 137
 - IBM Systems Director Navigator for i 405
- Performance Data 137
- Performance Explorer 137
- Performance Monitoring Infrastructure 305
- performance setting Web page 298
- Permission class 330
- Personal Communications 5250 430
- PEX-Analyzer changes 144
- physical footprint 475
- Ping-Of-Death 6
- PIV 196
- plan cache
 - database 93
- planning statements

- installing or upgrading 466
- Point-to-Point Protocol (PPP)
 - networking 222
- policies 7
 - security 7
- Post Office Protocol (POP)
 - networking 229
- Power 595 (9119-FHA) 187
- POWER Hypervisor
 - logical partitioning 197
- Power systems 474
- Power Systems virtualization
 - PowerVM 196
- POWER5 and POWER6 technology systems
 - IBM i 5.4 2
 - IBM i 6.1 2
- POWER6 virtualization 190
- PowerHA for i 2
 - availability 76
- PowerVM
 - Power Systems virtualization 196
- PowerVM Lx86 196
- PowerVM Virtualization on IBM System p
 - Introduction and Configuration Fourth Edition
 - Redbooks 216
- Preferences
 - IBM i Doctor 141
- pre-upgrade verification tool 466
- prevention system (IDS)
 - security 6
- Print JVM Job 132
- Print JVM Job command (PRTJVMJOB) 321
- Print PEX Report (PRTPEXRPT) 137
- Print Schedule using JS (PRTSCDJS) 276
- Print System Report (PRTSYSRPT) 138
- privileges
 - security 23
- processor feature 294
- Program (PGM) 342
- Program bracketing events (PGMBKTEVT) 181
- Program Calls
 - application diagram 365
- Program events (PGMEVT) 181
- property sheet views 140
- PRTJVMJOB 132
- PTF groups recommendation 285
- PTF tabs - Web Performance Advisor 296
- PWDSYSVAL
 - system values 13

Q

- QAPMWASAPP file 312
- QAPMWASCFG file 312
- QAPMWASEJB file 312
- QAPMWASRSC file 312
- QAPMWASSVR file 312
- QAQQINI
 - updates 93
- QoS server 6
- Qshell support 341

- QSQRVR job 304
- Quiesce processing 100
- QZDASOINIT job 304

R

- Rational brand name 2
- Rational Business Developer 362
- Rational Business Developer (RBD) V7.1 379
- Rational Developer for i 2
- Rational Developer for i for SOA Construction (RDi SOA) 3
- Rational Developer for IBM i 362
- Rational Development for i (RDi) 3
- Rational Host Access Transformation Services (HATS) Toolkit 362
- real time notification 6
- Redbook
 - i5/OS Diagnostic Tools for System Administrators
 - An A to Z Reference for Problem Determination 455
- Redbooks
 - Hierarchical Storage Management 66
 - IBM System i5 V5R4 Technical Overview 1
 - Implementing IBM Systems Director 6.1 3
 - Implementing PowerHA for IBM i 80
 - PowerVM Virtualization on IBM System p
 - Introduction and Configuration Fourth Edition 216
 - Security Guide for IBM i V6.1 10
 - VMware VI3 on BladeCenter and System x Integrated with System i 261
- Redbooks Web site 481
 - Contact us xx
- redundancy 475
- Remote Name Daemon Control (RNDC) 226
- remote recovery center 76
- remote SQL status view 140
- remote support 476
- Remote System Explorer (RSE) 364
- Remove Disk Watcher Definition (RMVDWDFN) 135
- Remove Job Watcher Definition (RMVJWDFN) 135
- Remove Kerberos Keytab Entry (RMVKRBKTE) 25
- Remove TCP/IP Table 11
- reporting setup 142
- Request for Comments (RFC) 2136
 - networking 226
- requirements
 - IBM i 6.1 Licensed Internal Code 469
- reserved pool capacity 198
- resources tab (advisor) 303
- RMVTCPTBL 11
- RMVTCPTBL command 11
- Run Remote Command (RUNRMTCMD) 273
- Run SQL Statements (RUNSQLSTM) 100

S

- SAN environment
 - security 11
- save while active 2
- scheduled job 268

- output queue monitor 271
- Scheduled Jobs Properties 275
- Screen Designer 376
- Screen Designer capabilities 365
- scripting and automation 476
- SEA 202
- Secure e-mail (S/MIME)
 - networking 229
- Secure Hash Algorithm (SHA) 23
- Secure Sockets Layer (SSL) 320
 - networking 230
- Secure Sockets Layer cipher control (QSSLCSLCTL)
 - system value 18
- Secure Sockets Layer protocols (QSSLPCL)
 - system value 18
- security
 - common attacks 6
 - enhanced functions 2
 - IBM Systems Director Navigator for i 405
 - intrusion detection 6
 - policies 7
- security audit journal 6
- Security Guide for IBM i V6.1
 - Redbooks 10
- server firmware 469
- servers supported 475
- serviceable event management 476
- session timeout
 - WLE 190
- Set Client Information (SQLESETI) API
 - DRDA 98
- Set Data Library using JS (SETDLJS) 273
- set password to expired 20
- setting up the HTTP administration server to use SSL 314
- shared Ethernet adapters 202
- shared processor pools 197
 - multiple 197
- shared-processor pools 197
- sharing optical resources 470
- Shortest Path First (SPF) algorithm
 - networking 223
- sign-on attempts not valid 20
- Simple Mail Transfer Protocol (SMTP)
 - networking 228
- sizing capability
 - disk storage systems 187
- sizing enablement 189
- sizing guides
 - workload 187
- SMTP server attributes (CHGSMTPA) 229
- Smurf attack 6
- SOA Construction (RDi SOA) 362
- software encryption
 - low-cost solution 13
 - security 12
- software license keys 469
- Source file (SRCFILE) 342
- spooled file conversions 467
- SPT 196
- SQE optimizer 90
- SQL CALL cache 92
- SQL limits
 - DRDA 97
- SQL performance monitors 436
- SQL plan cache 436
- SQL plan cache event monitor 121
- SQL plan cache size threshold 118
- SQL plan cache statements 119
- SQL Query Engine (SQE) 90
- SQL Runtime Performance Improvements 91
- SQL statement interruption
 - DRDA 97
- SQL_STMT_REUSE INI option 92
- SQLBindCol 346
- SQLBindParam 346
- SQLBindParameter 346
- SQLColumns 346
- SQLForeignKeys 346
- SQLGetCol 346
- SQLGetTypeInfo 346
- SQLPrimaryKeys 346
- SQLProcedureColumns 346
- SQLProcedures 346
- SQLSetConnectAttr 346
- SQLSetCursorName 346
- SQLSpecialColumns 346
- SQLStatistics 346
- SQLTables 346
- Standard data retention (days) (STDDTARET) 184
- Standard Edition (J2SE) 320
- Start Disk Watcher (STRDW) 135
- Start Group using JS (STRGRPJS) 276
- Start HOST Query (HOST) 225
- Start Job Watcher (STRJW) 135
- Start Journal Library (STRJRNLIB) 99
- Start Performance Explorer (STRPEX) 137
- Structured Query Language (SQL) 345, 435
- support for multiple profiles per partition 476
- switchable devices 81
- Synchronization events (SYNCEVT) 181
- System base (*SYSBAS) 86
- System base (*SYSBAS) commitment control transactions 99
- system CPW 294
- system dump 133
- System i Access for Linux 442
- System i Access for Web 442
- System i Access for Windows 267, 430
- System i Navigator 6
 - network enhancements 219
- System i Navigator for Wireless 446
- system model 294
- System Planning Tool 188, 196
- system resources 297
- system value
 - Block password change (QPWDCHGBLK) 16
 - Limit device sessions (LMTDEVSSN) 18
 - Password expiration warning (QPWDEXPWRN) 16
 - Password rules (QPWDRULES) 13

- QALWBJRST 140
- Secure Sockets Layer cipher control (QSSLCSLCTL) 18
- Secure Sockets Layer protocols (QSSLPCL) 18
- Verify object on restore (QVFYOBJRST) 19
- Systems Director Navigator for i 406
 - console 409
 - features 406

T

- technical overview 1
- technology 9406 models 2
- Temporary distinct sorted list 92
- Temporary values list 92
- The i5/OS V6R1 Technical Overview 1, 479
- time management and zones 440
- time synchronization 262
- Trace TCP/IP Application (TRCTCPAPP) 231
- transaction counts 291
- Transport Layer Security (TLS) 320

U

- uncapped partition support 475
- un-encrypted disk pool 11
- user interface 476
- user interfaces (UI) 362
- user security 475
- User-defined file systems (UDFS) 37
- users and groups
 - IBM Systems Director Navigator for i 405

V

- Verify object on restore (QVFYOBJRST)
 - system value 19
- Version Availability 185
- VIOS 196
- virtual console support 474–475
- Virtual I/O events (VRTIOEVT) 181
- Virtual I/O Server 196, 202
- Virtual I/O Server client support 191
- Virtual IP
 - networking 222
- virtual optical device 468
- Virtual Private Network (VPN)
 - networking 227
- virtual tape device 468
- virtualization
 - IBM i storage 2
- VMware VI3 on BladeCenter and System x Integrated with System i
 - Redbooks 261

W

- Wait buckets 138
- Web based systems management interface to IBM i
 - IBM Systems Director Navigator for i 3
- Web console interface 404
- Web Environment area (advisor) 301

- Web Environment JVM settings 301
- Web Environment of Web Performance Advisor 299
- Web Performance Advisor 292–293
 - performance attributes 293
- Web Performance Monitor 287
- Web-based system management interface 474
- WebFacing
 - HATS 395
- WebFacing applications 397
- WebSphere Development Studio (WDS) 331
- WebSphere Development Studio Client (WDS) Advanced Edition 364
- WebSphere Development Studio for iSeries client (WDS) development tool 3
- WebSphere Development Studio for iSeries client development tool
 - WDS 3
- WebSphere Development Tools for iSeries Client (WDS) 2
- WebSphere Enterprise Service Bus (WESB) 188
- WebSphere Enterprise Workload Manager 288
- WebSphere Process Server (WPS) 188
- well-known attacks 6
- WLE
 - session timeout 190
 - WLE 188–190
 - Workload Developer Toolkit 188
- WLE Workload Developer Toolkit enhancements 187
- work management 431
 - IBM Systems Director Navigator for i 404
- Work with History using JS (WRKHSTJS) 274
- Work with Jobs using JS (WRKJOBJS) 274
- Work with JVM Jobs (WRKJVMJOB) 129–130, 320
- Work with System Activity (WRKSYSACT) 137
- workload
 - sizing guides 187
- workload developer enhancements 190
- Workload Developer Toolkit 188
 - WLE 188
- Workload Management (WLM) 476
- workload updates 190
- workloads and sizing guides 187
- WRKACTJOB 136
- WRKSYSACT 136

X

- X-An 364
- X-Analysis tool from Databorough
 - vendor tools 364
- XML Toolkit for IBM System i 341



Redbooks

IBM i 6.1 Technical Overview

(1.0" spine)

0.875" x 1.498"

460 x 788 pages



IBM i 6.1 Technical Overview



Manage complexity and risk with new Web-based management and security functions

Boost efficiency with virtualization and dynamic workload management

Realize innovation with open application design choices

This IBM Redbooks publication introduces a technical overview of the main new features, functions and enhancements available in IBM i 6.1 (formerly called i5/OS V6R1). It gives a summary and brief explanation of new capabilities and what has changed in the operating system, and also discusses many of the licensed programs and application development tools associated with IBM i.

Many other new and enhanced functions are described, such as virtualization of storage, security, Java performance, improved performance with IBM System Storage devices, backup and recovery, including base IBM i, Backup, Recovery and Media Services (BRMS). The book introduces the PowerHA product, IBM Systems Director-based system management and an easier Web enablement.

The information provided in this book will be useful for customers, Business Partners, and IBM service professionals involved with planning, supporting, upgrading, and implementing IBM i 6.1 solutions.

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

SG24-7713-00

ISBN 0738433608