Sun StorEdge™ Fast Write Cache Best Practices

Improving Oracle Database Performance on the Sun StorEdge™ A5x00

Overview

Fast Write Cache reduces the frequency of disk I/O access by caching the written data blocks in nonvolatile memory, and later destaging the cached data to disk asynchronously. It is implemented using nonvolatile random access memory (NVRAM) cards installed in the host computer system, through installation on the I/O boards.

Today, the Sun StorEdge™ A5000 family delivers high performance in a variety of applications. These applications include:

- Data warehousing
- Decision support
- Imaging

The disk drives used in the A5000 family line, whether 7200RPM or 10000RPM, have very high throughput. However, they do not perform optimally when doing small write I/Os that require frequent head seeks. The Fast Write Cache improves the performance of the A5000 by speeding up these small writes extending its performance leadership and enabling the A5000 family to excel at additional applications, such as:

- Online Transaction Processing (OLTP)
- NFS™

Fast Write Cache allows individual volumes to be configured by the system administrator so only volumes that would benefit from caching are cached.

Consider many factors when investigating the database Online Transaction Processing (OLTP) performance. These include:

- Transaction type (update, insert, query)
- CPU usage
- I/O performance, including performance of data tablespaces, index tablespaces, rollback segments, redo logs, and any temporary tablespaces.

May 25, 1999

Sun microsystems

I/0 Performance

It is important to understand the I/O profile for the different components and where the database bottlenecks are.

Small sequential writes: Redo logs and rollback segments

Larger random read/writes (size depends on DB block size): Data tablespaces and index tablespace

Sequential read/writes: Data tablespace (during a full table scan or a load) temporary tablespace

Because of the nature of I/O, caching of the redo logs and rollback segments can increase database throughput without any other changes. When the logs are the bottleneck, caching them increases database throughput by more than 16%.

If the database performance is CPU bound and not limited by the log I/O, caching the logs has little impact on the database performance. This is documented in Report 3.

Database Functions

The following sections summarize database functions.

Redo Logs

I/O profile: Small sequential writes

The system writes heavily to the redo logs, so they are typically put on the fastest device. They contain a sequential log of all of the changes made to the database tables, as well as information on checkpoint and other administrative information that can be used to recover a database. The log writer process (LGWR) writes redo log entries sequentially to disk. There are normally a minimum of two redo log files. If the database is running in archive log mode, normally the redo log files are spread over multiple physical disks so that current log writer activities do not compete with archive activity. Devices are either raw or file systems (synchronous writes).

Using Fast Write Cache with the A5000 family speeds up the logs, thereby improving transaction processing.

Rollback Segments

I/O profile: Small sequential writes and reads

A rollback segment is a storage space within a tablespace that holds transaction information used to guarantee data integrity during a rollback and is used to provide read consistency across multiple transactions. Most sites create a number of rollback tablespaces and use these exclusively for rollback segments.

Archive Redo Logs

I/O profile: Sequential reads

When archiving the redo logs, information is sequentially read from offline redo log files and is saved to tape or to disk. Fast Write Cache does not benefit the read operation.

Checkpoint

A checkpoint occurs periodically, normally on a log switch. Each checkpoint forces all modified database buffers to be written to the database. Each data file in the database is also updated by either the LGWR or the checkpoint process (CKPT) with the latest checkpoint information.

Areas Where Fast Write Cache May Be Beneficial

- Caching of the redo logs and rollback segments can increase database throughput without any other changes.
- Caching of the redo logs reduces log switch times. Log switches then occur at memory speeds, instead of disk speeds. It is not practical to measure this improvement, as the tools used have a one second granularity and the log switch occurs in less than one second, as opposed to two to five seconds when the logs are not cached.
- Checkpoints also benefit as a result of caching the logs. After the Database Writer (DBWR) writes out the db-buffers to the data files, it writes out the redo buffer to the redo logs. This now can occur at memory speeds.

Methodology

Several applications were run to evaluate the effect the Fast Write Cache has on the overall database performance. One application was a TPC-C like Benchmark C (TPC-C), TPC-SO. It is a server-only option that eases the running and reduces the hardware required to run a workload.

Note – Do not compare the results presented in this paper with TPC-C numbers. It is only a comparison to be used to compare cached redo logs vs. non cached redo logs.

TPC-C is an OLTP workload. It is a mixture of read-only and update intensive transactions that simulate the activities found in complex OLTP application environments. It exercises a range of system components associated with environments that are characterized by:

- The simultaneous execution of multiple transaction types that span a range of complexity
- Online and deferred transaction execution modes
- Moderate system and application execution time
- Significant disk input and output
- Transaction integrity
- Nonuniform distribution of data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of size, attributes, and relationships
- Contention of data access and update

The performance metric reported by TPC-C is a *business throughput* measuring the number of orders processed per minute. Multiple transactions are used to simulate the business activity of processing an order, and each transaction is subject to a response time constraint. The performance metric for this benchmark is expressed in transactions-per-minute-C (tpm-C).

TPC-C throughput is defined as how many New-Order transactions per minute a system generates while the system is executing four other transactions types (Payment, Order-Status, Delivery, Stock-Level). All five TPC-C transactions have a user response time requirement, with the New-Order transaction response time set at five seconds. Therefore, for a 710tpmC number, a system is generating 710 New-Order transactions per minute while fulfilling the rest of the TPC-C transaction mix workload.

Refer to `www.tpc.org f`or more information on TPC-C.

Test Environment

The test environment consisted of the following:

- Sun Enterprise™ E6000 – 8 250MHz CPUs and 4Gbytes of main memory
- One A5200 Network Array
- Solaris™ operating environment 2.6 – HW 5/98
- Sun Enterprise Volume Manager™ 2.6
- Oracle 8.0.5
- Fast Write Cache 1.0

The Fast Write Cache resides above the volume manager and below the file system. It is not specific to any volume manager. The architecture looks like:

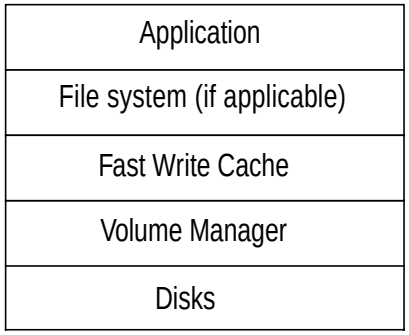

There are many benefits and reasons for the cache being above the volume manager. They include:

- Fast Write Cache can coalesce small sequential RAID 5 writes and turn them into full-stripe writes, increasing throughput and reducing response time.
- The application does one write and gets notification of the write completion, once the data is saved in NVRAM. The data is destaged through the appropriate volume manager to the appropriate RAID 1 or RAID 5 disks.

Testing Scenario

The Sun StorEdge A5200 Network array was configured to load the database. The following two configurations were tested:

Configuration	Description
RAID 5	An 8+1 RAID 5 volume was built with a 32Kbyte segment size, yielding a 256Kbyte stripe width. All tables were placed in this volume, with the exception of the rollback table and the redo logs. Those were placed on a separate RAID 5 volume (a 2+1) with a 32Kbyte segment size. Multiple redo logs were used but were placed on the same RAID 5 device.
RAID 0+1	An 8+8 RAID 0+1 volume was built with a 64Kbyte segment size. All tables were placed in this volume, with the exception of the rollback table and the redo logs. Those were placed on a separate RAID 1 volume (1 +1) with a 32Kbyte segment size. Multiple redo logs were used but were placed on the same RAID 1 device.

Various TPC-C scenarios were run on one of the above configurations. A run was done to the A5200, and then that same workload was rerun, with various volumes, within the A5200 cache. The system global area (SGA) and caching on the host were taken into account. Each run consisted of a ramp-up, steady state, and a ramp-down. Results reported are from the steady state time period only.

Test Results

Writing to memory is much faster than writing to disk, but eventually the data written to cache must be written to disk. Speeding up the writes to the redo logs generally allowed more transactions to occur in the same amount of time. That is, transactions per minute (tpm-C) went up. There was no increase in the one case of the database being CPU bound. No more transactions could be generated.

Other improvements were observed, including reduction of log switching times and faster checkpoints. They are described in more detail in the “Conclusion” section.

Note – Do not compare the results presented in this paper with TPC-C numbers. It is only a comparison to be used to compare cached redo logs vs. non cached redo logs.

RAID 5 – No Cache

Running to the RAID 5 devices without cache delivered a tpm-C of 1676. The disks, logs, and tablespaces were overused.

```
CPU state: Underused.
Network state: OK
Disk state: Significantly Overused.

CPU statistics:
 User   System  Idle
12.49%  11.22%  76.29%

Affected disks:

Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd14   92.0    167.3   100.2    67.1     2048    3361    0.00    7.32   43.7
ssd17   92.9    170.0   109.3    60.7     2048    2397    0.00    7.12   41.9
ssd19   91.6    160.2    89.6    70.6     2048    3899    0.00    6.46   40.3
ssd21   91.8    168.7    97.8    70.9     2048    3480    0.00    6.99   41.5
ssd3    91.6    168.3   100.6    67.6     2048    3773    0.00    7.09   42.1
ssd4    90.2    164.5    99.1    65.4     2048    3163    0.00    6.40   38.9
ssd5    92.9    167.1   102.3    64.9     2048    3275    0.00    7.52   45.0
ssd10   94.4    170.3   102.1    68.2     2048    3749    0.00    8.50   49.9
ssd12   90.3    150.6    86.2    64.3     2048    4790    0.00    6.46   42.9
```

RAID 5 – With Cache

Caching the redo logs and rollback table increased the tpm-C number to 1827; an increase of 5%. Analysis of the results showed that in both cases the limiting factor was access to the database tables in the RAID 5 volume. We were I/O bound. Caching the logs and rollback segments reduced the load on the A5200, thus allowing more reads and updates to the database tables.

```
CPU state: Significantly Underused.
Network state: OK
Disk state: Significantly Overused.

CPU statistics:
 User   System  Idle
10.66%  8.30%   81.04%

Affected disks:
Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd14   95.4    182.1   109.2    72.9     2048    3382    0.00    9.73   53.4
ssd17   96.6    184.9   118.4    66.5     2048    2448    0.01    9.56   51.7
ssd19   94.6    174.6    97.8    76.8     2048    4004    0.00    8.57   49.1
ssd21   95.3    182.6   105.8    76.8     2048    3532    0.00    9.08   49.8
ssd3    95.6    183.3   109.4    73.9     2048    3751    0.00    9.37   51.1
ssd4    94.1    179.3   108.0    71.4     2048    3154    0.00    8.43   47.0
ssd5    96.1    180.3   110.4    69.8     2048    3335    0.00   10.05   55.8
ssd10   97.4    183.5   110.5    73.0     2048    3831    0.01   11.36   62.0
ssd12   94.8    166.5    96.4    70.0     2048    4596    0.00    9.43   56.7
```

The caching of the logs and rollback segments enable the A5200 array to work more efficiently. Instead of the physical spindles doing a lot of small I/Os, the logs were destaged in less frequent larger I/Os. Thus, a few more I/Os were done to the data tables. See “Report 1” for the complete results.

RAID 1 – No Cache

Running to the RAID 0+1 devices without cache delivered a tpm-C of 2824.

```
CPU state: Significantly Underused.
Network state: OK
Disk state: Significantly Overused.

CPU statistics:
 User   System  Idle
14.69%  4.10%   81.21%

Affected disks:
```

```
Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd30   66.6     98.1    50.2    47.9     2048    2048    0.00    1.92   19.6
ssd31   69.5     98.0    50.1    47.9     2048    2048    0.00    2.12   21.6
ssd32   68.2    104.0    52.9    51.1     2048    2048    0.00    2.03   19.5
ssd33   70.5    103.7    52.9    50.8     2048    2048    0.00    2.21   21.3
ssd36   62.6     89.4    46.4    43.1     2048    2048    0.00    1.70   19.0
ssd37   70.3    104.3    53.3    51.1     2048    2048    0.00    2.19   20.9
ssd39   85.7    128.6    63.2    65.4     2048    2048    0.00    3.76   29.3
ssd25   70.6    102.5    52.3    50.2     2048    2048    0.00    2.22   21.6
ssd28   63.0     89.7    46.6    43.1     2048    2048    0.00    1.72   19.2
ssd29   83.1    129.7    64.3    65.4     2048    2048    0.00    3.38   26.1
ssd42   68.4    103.3    52.5    50.8     2048    2048    0.00    2.10   20.4
ssd43   70.0    102.9    52.7    50.2     2048    2048    0.00    2.13   20.7
```

RAID 1 – With Cache

Caching the redo logs and rollback table increased the tpm-C number to 3031; an increase of almost 8%.

```
CPU state: Underused.
Network state: OK
Disk state: Significantly Overused.

CPU statistics:
 User   System  Idle
15.75%  4.95%   79.30%

Affected disks:
Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd30   68.8    101.6    47.0    54.6     2048    2048    0.00    2.23   21.9
ssd31   72.6    102.3    47.7    54.6     2048    2048    0.00    2.53   24.7
ssd32   69.5    106.1    50.1    56.0     2048    2048    0.00    2.27   21.4
ssd33   71.9    105.8    49.4    56.4     2048    2048    0.00    2.46   23.2
ssd36   64.0     91.4    43.1    48.3     2048    2048    0.00    1.92   21.0
ssd37   71.1    105.5    49.5    56.0     2048    2048    0.00    2.43   23.1
ssd39   87.8    131.5    60.5    71.0     2048    2048    0.00    4.23   32.2
ssd25   72.3    104.2    48.9    55.3     2048    2048    0.00    2.48   23.8
ssd28   65.1     91.5    43.3    48.3     2048    2048    0.00    1.98   21.6
ssd29   83.9    131.2    60.2    71.0     2048    2048    0.00    3.68   28.1
ssd42   70.2    105.7    49.3    56.4     2048    2048    0.00    2.33   22.1
ssd43   71.1    103.8    48.5    55.3     2048    2048    0.00    2.39   23.0
```

Analysis of the results indicated that in both cases the limiting factor was access to the database tables in the RAID 0+1 volume. We were I/O bound. Caching the logs and rollback segments reduced the load on the A5200, thus allowing more reads and updates to the database tables. See “Report 2” for the complete results.

RAID 1 – No Cache, Out of CPU Bandwidth

Reducing the I/O load on the array was done by limiting the database size under test, forcing more of the reads to be cached on the server and mostly writes to the array. A run in this environment showed disk usage was reduced significantly. The logs disks were busy; more than 60%.

```
CPU state: Overused.
Network state: OK
Disk state: Significantly Overused.

CPU statistics:
 User   System  Idle
89.64%  7.45%   2.91%

Affected disks:
Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd24   60.3    108.8     0.0   108.8       512   29121    0.00    0.60    5.5
ssd41   64.7    108.8     0.0   108.8         0   29121    0.00    0.65    6.0
```

RAID 1 – With Cache, Out of CPU Bandwidth

Caching of the log disk reduced disk usage to less than 20%.

```
CPU state: Significantly Overused.
Network state: OK
Disk state: OK

CPU statistics:
 User   System  Idle
88.81%  9.44%   1.75%
```

The logs were working more efficiently. However, the tpm-C did not increase. Analysis of the results showed that the database was out of the CPU. The CPUs were less than 2% idle and could not generate more transactions. See “Report 3” for the complete results of these runs.

RAID 1 – No Cache, Busy Logs

Reducing the load on the CPU meant scaling back the transactions. This was done by decreasing the number of users making the logs the problem. With the logs not cached, a tpm-C of 8405 was obtained.

```
CPU state: OK
Network state: OK
Disk state: Significantly Overused.

Affected disks:
Disk    util%   xfer/s  rds/s   wrts/s  rdb/xfr wrb/xfr wtqlen  svqlen  srv-ms
ssd24   97.3    181.5   129.5    52.0     5431   45366    0.00    1.91   10.5
ssd41   98.0    165.6   113.6    52.0     5971   45366    0.00    1.94   11.7
```

In the above case, as in previous cases, the log was mirrored. The two busy disks are the primary and mirror log disks.

RAID 1 – With Cache, Busy Logs

Caching of the log volumes improved tpm-C to 9761; an increase of more than 16%.

```
CPU state: OK
Network state: OK
Disk state: OK
```

The log disks were no longer busy. See “Report 4” for the complete results of these runs.

Implementation

Determine where the bottlenecks are. The redo logs, rollback segments, or any volume can be dynamically inserted into or removed from the Fast Write Cache.

Use the `fwcadm cache -s` command to monitor cache activity. Refer to the `fwcadm(1FWC)` man page.

▼ Adding Volumes Into Fast Write Cache

1. **Edit the file** `/etc/opt/SUNWspsv/sv.cf`**, adding the volumes you want cached and the keyword** `cache`**. For example:**

```
/dev/vx/rdsk/rootdg/log1 cache
/dev/vx/rdsk/rootdg/log2 cache
/dev/vx/rdsk/rootdg/roll1 cache
/dev/vx/rdsk/rootdg/rool2 cache
```

2. **Type the following command as root:**

```
# fwcadm volume -r
```

The volumes are now in the Fast Write Cache.

▼ Removing Volumes From Fast Write Cache

1. **Edit the file** `/etc/opt/SUNWspsv/sv.cf`**, removing the appropriate volumes.**
2. **Type the following command as root:**

```
# fwcadm volume -r
```

Conclusion

Caching of the redo logs and rollback segments can increase database throughput without any changes. However, the improvement depends on the current limiting factor. When the limiting factor is the logs, the caching of them increased database throughput by more than 16%. When the logs were not the limiting

factor, but were in the same array as other database tables, making the logs operate more efficiently achieved a 5% increase. When the IO was not the limiting factor, speeding up the logs had no effect. This is the case where we were already out of CPU cycles. This is documented in “Report 3.”

Caching of the redo logs also reduces log switch times and checkpoints. Log switches occurred at memory speeds, instead of disk speeds.

Reports

The following are the complete reports of the runs described early in this document.

Report 1

RAID 5 TPCCS Report No Cache

```
Test Data

Ramp-up       : 100
Ramp-down     : 60
Run-time      : 200
Trigger-time  : 25
Database scale: 50
Users         : 30
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 4 4 4 43 45
Count of servers : 12 8 2 8 28


MQTh (Maximum Qualified Throughput): 1748.00 tpmS


TRANSACTION MIX

Total number of transactions = 11820
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      05244           44.37%          -
```

Payment:	05040	42.64%	43
Order-Status:	00508	4.30%	4
Delivery:	00510	4.31%	4
Stock-level:	00518	4.38%	4

TPC-C Requirement: Passed

RESPONSE TIMES	AVG.	MAX.	90TH	REQD. 90TH
New-Order	0.607	2.630	1.000	5
Payment	0.335	1.345	0.800	5
Order-Status	0.399	1.431	0.800	5
Delivery(interactive)	0.106	0.680	0.240	5
Delivery(deferred)	0.858	4.000	3.000	80
Stock-level	1.652	5.952	4.000	20

TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response : Passed

THINK TIMES	AVG.	MIN.	MAX.	REQD. AVG.
New-Order	0.000	0.000	0.000	12
Payment	0.000	0.000	0.000	12
Order-Status	0.000	0.000	0.000	10
Delivery	0.000	0.000	0.000	5
Stock-level	0.000	0.000	0.000	5

MISC. TPC-C REQUIREMENTS

Average items per order	10.020	Passed
Remote order-lines	1.043	Passed
New-orders rolled back	1.144	Passed
Remote payments	14.821	Passed
Payments through C_LAST	59.821	Passed
Order-Status thru C_LAST	60.039	Passed
Delivery wih no New-order	0.000	Passed

LITTLE’s LAW VERIFICATION

Number of users = 30
Sum of Avg. RT * TPS for all Tx. Types = 29.908550

RAID 5 TPCCS Report Cache

```
Test Data

Ramp-up       : 100
Ramp-down     : 60
Run-time      : 200
Trigger-time  : 25
Database scale: 50
Users         : 30
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 4 4 4 43 45
Count of servers : 12 8 2 8 28


MQTh (Maximum Qualified Throughput): 1827.33 tpmS


TRANSACTION MIX

Total number of transactions = 12474
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      05482           43.95%          -
Payment:        05355           42.93%          43
Order-Status:   00532           4.26%           4
Delivery:       00560           4.49%           4
Stock-level:    00545           4.37%           4
TPC-C Requirement: Passed


RESPONSE TIMES          AVG.            MAX.            90TH    REQD. 90TH

New-Order               0.585           2.363           1.000           5
Payment                 0.219           1.533           0.600           5
Order-Status            0.520           1.865           1.000           5
Delivery(interactive)   0.100           0.666           0.230           5
Delivery(deferred)      0.973           5.000           3.000           80
Stock-level             2.269           8.615           5.000           20
TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response  : Passed
```

```
THINK TIMES     AVG.            MIN.            MAX.            REQD. AVG.

New-Order        0.000          0.000            0.000          12
Payment          0.000          0.000            0.000          12
Order-Status     0.000          0.000            0.000          10
Delivery         0.000          0.000            0.000          5
Stock-level      0.000          0.000            0.000          5

MISC. TPC-C REQUIREMENTS

Average items per order                9.960         Passed
Remote order-lines                     0.932         Passed
New-orders rolled back                 0.985         Passed
Remote payments                       15.761         Passed
Payments through C_LAST               60.448         Passed
Order-Status thru C_LAST              57.519         Passed
Delivery wih no New-order              0.000         Passed


LITTLE’s LAW VERIFICATION

Number of users = 30
Sum of Avg. RT * TPS for all Tx. Types = 29.752840



SAMPLE 14           ******      Storage Cache     ******             20:26:16

                          disk_io       cache          write_blocks
  cd cached_partition  reads writes  reads writes  dirty todisk failed
   0 .../raid5/roll3       0    164      0    140      1     51      0
   1 ...k/raid5/log1       0    595      0    570      2      2      0
   2 ...k/raid5/log2       0      0      0      0      0      0      0
                      ------ ------ ------ ------
       Kbytes/s total:     0    759      0    710

 accesses/s    read/s    write/s   %readh   %writeh
             (misses/s) (misses/s)
    144.12       0.00     144.12      0.0     100.0
             (   0.00 ) (   0.00 )

$
```

Report 2

RAID 1 TPCCS Report - 1 No cache

```
Test Data

Ramp-up       : 75
Ramp-down     : 60
Run-time      : 200
Trigger-time  : 25
Database scale: 50
Users         : 30
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 4 4 4 43 45
Count of servers : 12 8 2 8 28


MQTh (Maximum Qualified Throughput): 2824.00 tpmS


TRANSACTION MIX

Total number of transactions = 19323
$ cat summary
                TPCCS Report RAID 1 no cache

Test Data

Ramp-up       : 75
Ramp-down     : 60
Run-time      : 200
Trigger-time  : 25
Database scale: 50
Users         : 30
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 4 4 4 43 45
Count of servers : 12 8 2 8 28
```

MQTh (Maximum Qualified Throughput): 2824.00 tpmS

TRANSACTION MIX

Total number of transactions = 19323

TYPE	TX. COUNT	MIX	REQD. MIX.
New-Order:	08472	43.84%	-
Payment:	08346	43.19%	43
Order-Status:	00836	4.33%	4
Delivery:	00847	4.38%	4
Stock-level:	00822	4.25%	4

TPC-C Requirement: Passed

RESPONSE TIMES	AVG.	MAX.	90TH	REQD. 90TH
New-Order	0.399	2.179	0.800	5
Payment	0.156	1.236	0.400	5
Order-Status	0.227	0.808	0.400	5
Delivery(interactive)	0.099	0.662	0.230	5
Delivery(deferred)	0.341	2.000	2.000	80
Stock-level	1.251	5.397	3.000	20

TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response : Passed

THINK TIMES	AVG.	MIN.	MAX.	REQD. AVG.
New-Order	0.000	0.000	0.000	12
Payment	0.000	0.000	0.000	12
Order-Status	0.000	0.000	0.000	10
Delivery	0.000	0.000	0.000	5
Stock-level	0.000	0.000	0.000	5

MISC. TPC-C REQUIREMENTS

Average items per order	10.053	Passed
Remote order-lines	0.935	Passed
New-orders rolled back	0.944	Passed
Remote payments	14.630	Passed
Payments through C_LAST	60.388	Passed

```
Order-Status thru C_LAST              62.201         Passed
Delivery wih no New-order              0.000         Passed


LITTLE’s LAW VERIFICATION

Number of users = 30
Sum of Avg. RT * TPS for all Tx. Types = 29.923990


/s  w/s   kr/s   kw/s wait actv wsvc_t asvc_t  %w  %b device
  0.0  0.4    0.0    2.1  0.0  0.0    0.0   12.2   0   0 c0t1d0
  0.0  0.6    0.0    4.8  0.0  0.0    0.0   18.1   0   1 c0t2d0
  0.0  1.2    0.0    7.7  0.0  0.0    0.0   14.9   0   1 c0t3d0
  0.0 74.0    0.0  148.0  0.8  2.9   10.6   39.7   3  28 c2t23d0
  0.0 75.7    0.0  639.1  0.0  0.4    0.0    4.9   0  37 c2t22d0
 62.3 19.8  124.6   39.7  0.0  1.3    0.0   15.4   0  55 c2t18d0
 61.5 13.7  123.1   27.4  0.0  1.1    0.0   14.1   0  51 c2t16d0
 58.3 12.0  116.7   24.0  0.0  0.9    0.0   12.4   0  49 c2t2d0
 61.9 17.9  123.8   35.9  0.0  1.3    0.0   15.9   0  55 c2t19d0
 74.1 24.0  148.3   48.0  0.0  1.7    0.0   17.4   0  64 c2t9d0
 64.0 17.4  128.1   34.9  0.0  1.2    0.0   15.1   0  54 c2t7d0
 62.7 17.5  125.4   35.0  0.0  1.3    0.0   16.8   0  56 c2t20d0
 65.0 18.0  130.1   36.0  0.0  1.2    0.0   14.3   0  56 c2t4d0
 64.5 20.8  129.1   41.7  0.0  1.4    0.0   16.8   0  57 c2t8d0
 60.0 13.7  120.1   27.4  0.0  1.1    0.0   14.3   0  50 c2t3d0
  0.0 74.0    0.0  148.1  0.7  2.9    9.4   39.3   3  28 c2t24d0
 63.5 17.9  127.1   35.9  0.0  1.3    0.0   15.6   0  56 c2t6d0
 68.6 18.0  137.2   36.0  0.0  1.3    0.0   14.9   0  58 c2t17d0
 58.6 12.0  117.2   24.0  0.0  0.8    0.0   11.5   0  48 c2t10d0
 72.9 24.0  145.8   48.0  0.0  1.8    0.0   19.0   0  65 c2t1d0
  0.0 75.7    0.0  639.1  0.0  0.3    0.0    4.5   0  34 c2t21d0
 62.3 20.9  124.7   41.8  0.0  1.4    0.0   16.8   0  55 c2t0d0
 69.8 19.8  139.6   39.7  0.0  1.5    0.0   16.3   0  60 c2t5d0
```

RAID 1 TPCCS Report Cache

```
Test Data

Ramp-up       : 75
Ramp-down     : 60
Run-time      : 200
Trigger-time  : 25
Database scale: 50
Users         : 30
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 4 4 4 43 45
Count of servers : 12 8 2 8 28


MQTh (Maximum Qualified Throughput): 3031.33 tpmS


TRANSACTION MIX

Total number of transactions = 20649
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      09094           44.04%          -
Payment:        08877           42.99%          43
Order-Status:   00900           4.36%           4
Delivery:       00894           4.33%           4
Stock-level:    00884           4.28%           4
TPC-C Requirement: Passed


RESPONSE TIMES          AVG.            MAX.            90TH    REQD. 90TH

New-Order               0.380           2.095           0.800           5
Payment                 0.151           1.230           0.400           5
Order-Status            0.231           0.835           0.400           5
Delivery(interactive)   0.095           0.607           0.220           5
Delivery(deferred)      0.338           2.000           2.000           80
Stock-level             1.017           4.416           2.000           20
TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response  : Passed
```

THINK TIMES	AVG.	MIN.	MAX.	REQD. AVG.
New-Order	0.000	0.000	0.000	12
Payment	0.000	0.000	0.000	12
Order-Status	0.000	0.000	0.000	10
Delivery	0.000	0.000	0.000	5
Stock-level	0.000	0.000	0.000	5

MISC. TPC-C REQUIREMENTS

Average items per order	9.999	Passed
Remote order-lines	0.944	Passed
New-orders rolled back	0.891	Passed
Remote payments	15.129	Passed
Payments through C_LAST	60.471	Passed
Order-Status thru C_LAST	59.444	Passed
Delivery wih no New-order	0.000	Passed

LITTLE’s LAW VERIFICATION

Number of users = 30
Sum of Avg. RT * TPS for all Tx. Types = 29.975800

Report 3

RAID 1 TPCCS Report NO cache
Out of CPU Bandwitdh

```
Test Data

Ramp-up       : 200
Ramp-down     : 100
Run-time      : 300
Trigger-time  : 15
Database scale: 10
Users         : 50
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 5 10 5 45 35
Count of servers : 16 20 4 14 33


MQTh (Maximum Qualified Throughput): 5644.20 tpmS


TRANSACTION MIX

Total number of transactions = 82231
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      28221           34.32%          -
Payment:        37199           45.24%          43
Order-Status:   04298           5.23%           4
Delivery:       08240           10.02%          4
Stock-level:    04273           5.20%           4
TPC-C Requirement: Passed


RESPONSE TIMES          AVG.            MAX.            90TH    REQD. 90TH

New-Order               0.275           2.355           0.600           5
Payment                 0.142           1.230           0.400           5
Order-Status            0.134           1.104           0.400           5
Delivery(interactive)   0.100           0.809           0.240           5
```

```
Delivery(deferred)      0.227           5.000           2.000           80
Stock-level             0.124           0.856           1.000           20
TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response  : Passed


THINK TIMES     AVG.            MIN.            MAX.            REQD. AVG.

New-Order        0.000          0.000            0.000          12
Payment          0.000          0.000            0.000          12
Order-Status     0.000          0.000            0.000          10
Delivery         0.000          0.000            0.000          5
Stock-level      0.000          0.000            0.000          5

MISC. TPC-C REQUIREMENTS

Average items per order               9.993         Passed
Remote order-lines                    0.980         Passed
New-orders rolled back                0.992         Passed
Remote payments                      14.710         Passed
Payments through C_LAST              60.195         Passed
Order-Status thru C_LAST             59.121         Passed
Delivery wih no New-order             0.000         Passed


LITTLE's LAW VERIFICATION

Number of users = 50
Sum of Avg. RT * TPS for all Tx. Types = 49.930647
```

RAID 1 TPCCS Report Cache
Out of CPU Bandwidth

```
Test Data

Ramp-up       : 200
Ramp-down     : 100
Run-time      : 300
Trigger-time  : 15
Database scale: 10
Users         : 50
Slave machines: saratoga
```

```
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 5 10 5 45 35
Count of servers : 16 20 4 14 33


MQTh (Maximum Qualified Throughput): 6060.40 tpmS


TRANSACTION MIX

Total number of transactions = 88193
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      30302           34.36%          -
Payment:        39878           45.22%          43
Order-Status:   04621           5.24%           4
Delivery:       08809           9.99%           4
Stock-level:    04583           5.20%           4
TPC-C Requirement: Passed


RESPONSE TIMES          AVG.            MAX.            90TH    REQD. 90TH

New-Order               0.259           2.386           0.600           5
Payment                 0.129           1.206           0.400           5
Order-Status            0.124           0.799           0.400           5
Delivery(interactive)   0.100           0.904           0.240           5
Delivery(deferred)      0.042           1.000           1.000           80
Stock-level             0.122           0.841           1.000           20
TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response  : Passed


THINK TIMES     AVG.            MIN.            MAX.            REQD. AVG.

New-Order        0.000          0.000            0.000          12
Payment          0.000          0.000            0.000          12
Order-Status     0.000          0.000            0.000          10
Delivery         0.000          0.000            0.000          5
Stock-level      0.000          0.000            0.000          5

MISC. TPC-C REQUIREMENTS
```

Average items per order	10.010	Passed
Remote order-lines	0.997	Passed
New-orders rolled back	0.977	Passed
Remote payments	14.803	Passed
Payments through C_LAST	59.943	Passed
Order-Status thru C_LAST	60.398	Passed
Delivery wih no New-order	0.000	Passed

LITTLE’s LAW VERIFICATION

Number of users = 50
Sum of Avg. RT * TPS for all Tx. Types = 49.995730

Report 4

RAID 1 TPCCS Report NO Cache. Busy Logs.

```
Test Data

Ramp-up       : 250
Ramp-down     : 100
Run-time      : 300
Trigger-time  : 15
Database scale: 10
Users         : 90
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 5 10 5 45 35
Count of servers : 16 20 4 14 33


MQTh (Maximum Qualified Throughput): 8405.60 tpmS


TRANSACTION MIX

Total number of transactions = 122089
TYPE            TX. COUNT       MIX             REQD. MIX.
----            ---------       ---             ----------
New-Order:      42028           34.42%          -
Payment:        55110           45.14%          43
Order-Status:   06366           5.21%           4
Delivery:       12200           9.99%           4
Stock-level:    06385           5.23%           4
TPC-C Requirement: Passed


RESPONSE TIMES          AVG.            MAX.            90TH    REQD. 90TH

New-Order               0.301           2.409           0.800           5
Payment                 0.207           1.447           0.600           5
Order-Status            0.133           1.121           0.400           5
Delivery(interactive)   0.099           1.000           0.230           5
Delivery(deferred)      0.096           1.000           1.000           80
Stock-level             0.126           0.996           1.000           20
```

```
TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response  : Passed


THINK TIMES     AVG.            MIN.            MAX.            REQD. AVG.

New-Order        0.000          0.000            0.000          12
Payment          0.000          0.000            0.000          12
Order-Status     0.000          0.000            0.000          10
Delivery         0.000          0.000            0.000          5
Stock-level      0.000          0.000            0.000          5

MISC. TPC-C REQUIREMENTS

Average items per order               9.975         Passed
Remote order-lines                    1.002         Passed
New-orders rolled back                1.009         Passed
Remote payments                      15.139         Passed
Payments through C_LAST              60.054         Passed
Order-Status thru C_LAST             60.289         Passed
Delivery wih no New-order             0.000         Passed


LITTLE’s LAW VERIFICATION

Number of users = 90
Sum of Avg. RT * TPS for all Tx. Types = 89.735247
```

RAID 1 TPCCS Report Cache Busy Log

```
Test Data

Ramp-up       : 250
Ramp-down     : 100
Run-time      : 300
Trigger-time  : 15
Database scale: 10
Users         : 90
Slave machines: saratoga
Think times     : 0 0 0 0 0
Term delays     : 100 100 110 120 230
Tx. mix : 5 10 5 45 35
Count of servers : 16 20 4 14 33
```

MQTh (Maximum Qualified Throughput): 9761.20 tpmS

TRANSACTION MIX

Total number of transactions = 142138

TYPE	TX. COUNT	MIX	REQD. MIX.
New-Order:	48806	34.34%	-
Payment:	64364	45.28%	43
Order-Status:	07388	5.20%	4
Delivery:	14150	9.96%	4
Stock-level:	07430	5.23%	4

TPC-C Requirement: Passed

RESPONSE TIMES	AVG.	MAX.	90TH	REQD. 90TH
New-Order	0.285	5.324	0.600	5
Payment	0.145	17.521	0.400	5
Order-Status	0.143	1.130	0.400	5
Delivery(interactive)	0.102	8.747	0.240	5
Delivery(deferred)	0.090	14.000	1.000	80
Stock-level	0.132	0.963	1.000	20

TPC-C Requirement for 90th percentile: Passed
TPC-C Requirement for Avg. response : Passed

THINK TIMES	AVG.	MIN.	MAX.	REQD. AVG.
New-Order	0.000	0.000	0.000	12
Payment	0.000	0.000	0.000	12
Order-Status	0.000	0.000	0.000	10
Delivery	0.000	0.000	0.000	5
Stock-level	0.000	0.000	0.000	5

MISC. TPC-C REQUIREMENTS

Average items per order	10.003	Passed
Remote order-lines	0.979	Passed
New-orders rolld back	0.992	Passed
Remote payments	15.398	Passed

```
Payments through C_LAST               60.071         Passed
Order-Status thru C_LAST              59.854         Passed
Delivery wih no New-order              0.000         Passed


LITTLE’s LAW VERIFICATION

Number of users = 90
Sum of Avg. RT * TPS for all Tx. Types = 89.067080
```