

Improving performance with the Battery Backed Write Cache enabler on HP Smart Array controllers



Executive summary.....	3
Introduction.....	3
Write cache	3
Battery Backed Write Cache.....	3
Read cache.....	4
Cache size and ratio.....	4
Deploying the BBWC enabler	5
Field installation.....	7
Charging the battery	7
Performance testing.....	8
HP performance test.....	8
Tested environment.....	9
Configurations.....	10
Test tools and scripts.....	11
Tools.....	11
More information on the test tools.....	11
Heavy User script	11
Test scenarios	12
Performance test results.....	13
Write cache disabled	13
Write cache enabled	14
Results summary	15
Session logon test	16
Test environment	16
Test tools and scripts.....	16
Test scenarios	16
Write cache disabled	16
Write cache enabled	17
Test results.....	17
Logon times	17
Response times – write cache disabled.....	18
Response times – write cache enabled	19

Real-world performance.....	19
Conclusions	20
For more information.....	21

Executive summary

Enabling Battery Backed Write Cache (BBWC) can deliver significant performance improvements while providing reliable, battery-backed protection for write cache data.

This performance characterization outlines the benefits of write cache and describes how to enable this capability on the Smart Array controllers featured in HP ProLiant server platforms that are typically used in HP Server Based Computing (SBC) environments.

The results of tests run in the HP lab show that enabling BBWC can improve disk response times by up to 70%; real-world performance metrics indicate that even larger improvements are possible. This paper describes the tested components, test configurations, topology and test methodology; test results are summarized.

Intended audience: This paper is intended primarily for IT professionals planning HP SBC solution deployments or upgrades. The information and recommendations provided herein are designed to help customers optimize I/O subsystem performance and make them aware of the inherent limitations in rotational disk drive technologies.

The subject matter is fairly technical and requires an understanding of performance issues in the HP SBC space and the usage of the Microsoft® Windows® Performance Monitor (Perfmon) tool.

Introduction

HP Smart Array controllers include an allocation of RAM that can be utilized to temporarily buffer data being written to or read from disk. Since access to this RAM (cache) is significantly faster than disk access, this cache can enhance overall server performance.

Write cache

Data being written to disk are temporarily buffered in write cache before being written to disk. After receiving all the data associated with a particular write command, the controller indicates to the server that the data transfer is complete – even though data are still being written to disk (a process known as cache synchronization).

By buffering write data, disk throughput can be improved significantly. Testing carried out by HP indicated that, with the optimal number of users logged on, disk response times were as much as 70% faster with write cache enabled.

While faster access times can translate directly to enhanced system performance, this improvement comes with a potential penalty: if a system or utility failure were to occur during cache synchronization, data may be lost unless the write cache is battery-backed.

Battery Backed Write Cache

Because of the potential for data loss, HP automatically disables write cache unless an optional Battery Backed Write Cache (BBWC) enabler is installed and fully-charged. BBWC allows you to cache write data safely, knowing that data would be protected by the battery in the event of a system or utility failure.

The BBWC enabler includes the following components:

- A battery module, which includes a charger and status indicators
- A field-installable battery cable

Important:

Write cache is automatically disabled until a BBWC enabler is installed. After the BBWC enabler has been installed, write cache is enabled until battery capacity drops below 79%.

Read cache

In response to a series of sequential read commands, a larger block of data than that specified is pre-fetched and stored in read cache. If a subsequent read command finds the data it needs in cache, access is faster.

Cache size and ratio

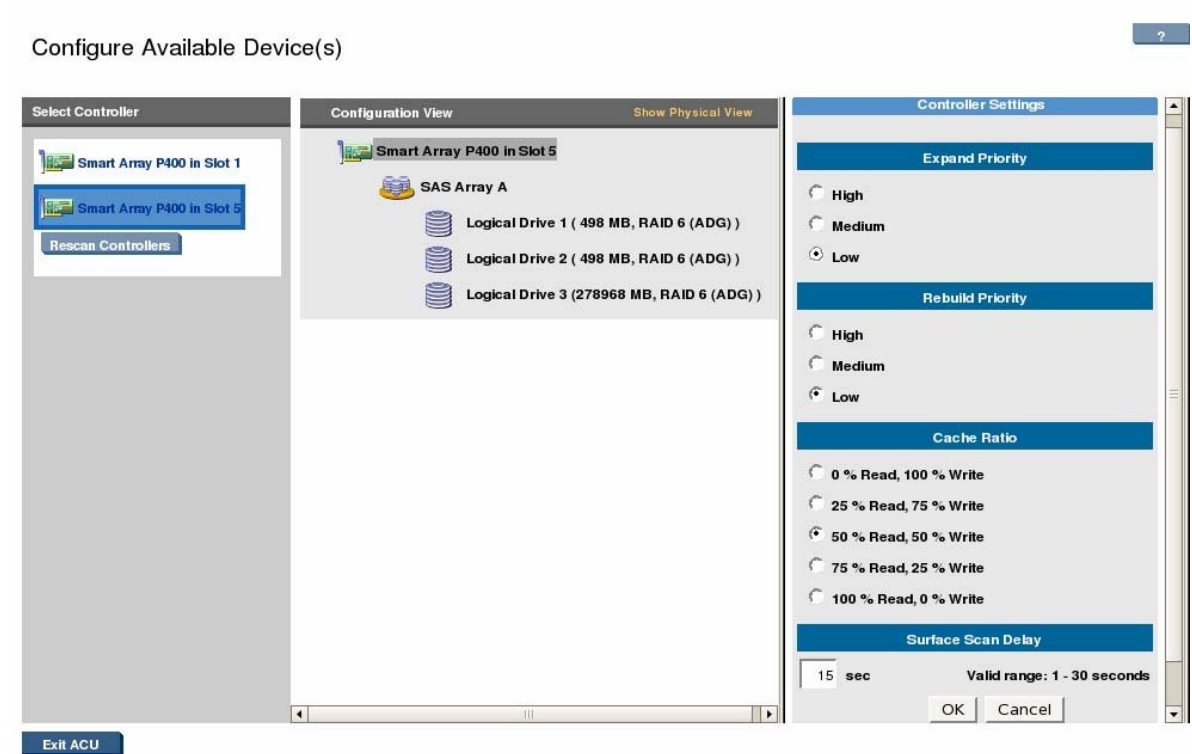
Total cache size (read cache and write cache) varies depending on the server model and options selected. Typically, cache size is between 64 MB and 512 MB.

Write cache is automatically disabled until a BBWC enabler is installed. Without this enabler, the cache ratio defaults to 100% read, 0% write. When an enabler is installed, write cache is enabled and the cache read/write ratio automatically sets to its default (e.g., 50% read, 50% write or 25% read, 75% write – depending on controller).

The size of the write cache allocation is user-selectable through the HP Array Configuration Utility (ACU) software, which allows you to select the percentages of RAM allocated to read and write cache respectively – for example, 25% read and 75% write.

Figure 1 shows a sample ACU screen.

Figure 1. Using ACU to set the percentages of RAM allocated to read and write cache



Deploying the BBWC enabler

Depending on the particular HP ProLiant server platform, there are various mechanisms for deploying a BBWC enabler:

- The enabler may be a standard feature
- The enabler may be available as an option

Table 1 indicates how the BBWC enabler can be deployed on HP ProLiant server platforms that are typically used in an HP SBC environment.

Table 1. Cache options for commonly-deployed storage controllers

HP ProLiant platform	Controller	Standard BBWC	Cache option	Part Number	
DL360 G5	Performance	P400i/256	Yes	512 MB with BBWC enabler	405148-B21
	Base	P400i/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
	Entry	E200i/64	No	128 MB with BBWC enabler	346914-B21
DL365	Performance	P400i/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
	Base	P400i/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
Entry	E200i/64	No	128 MB with BBWC enabler	346914-B21	
DL380 G5	Performance	P400/512	Yes		
	Base	P400/256	No	BBWC Enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
Entry	E200/64	No	128 MB with BBWC enabler	351580-B21	
DL385 G2	Performance	P400/512	Yes		
	Base	P400/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
Entry	E200/64	No	128 MB with BBWC enabler	351580-B21	
DL580 G4	Performance	P400/512	Yes		
	Base	P400/256	No	BBWC enabler	383280-B21
512 MB with BBWC enabler				405148-B21	
DL580 G5	Performance	P400/512	Yes		
	Base	P400/256	No	BBWC enabler	383280-B21
512 MB with BBWC enabler				405148-B21	
DL585 G2		P400/512	Yes		
		P400/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
BL460c		E200i/64	No	128 MB with BBWC enabler	351580-B21
BL465c		E200i/64	No	128 MB with BBWC enabler	351580-B21
BL480c		P400i/256	No	BBWC enabler	383280-B21
				512 MB with BBWC enabler	405148-B21
BL685c		E200i/64	No	128 MB with BBWC enabler	351580-B21
BL20p G4		E200i/64	No	128 MB with BBWC enabler	346914-B21
BL25p G2		E200i/64	No	128 MB with BBWC enabler	346914-B21
BL35p		Internal four-port SAS HBA (no cache)			
BL45p G2		E200i/64	No	128 MB with BBWC enabler	346914-B21

Field installation

Figures 2 and 3 show a typical BBWC enabler field installation.

Figure 2. Mounting the BBWC enabler on the standoffs provided





Figure 3. Connecting the BBWC cable to the BBWC enabler



Charging the battery

BBWC battery charging is automatic. The battery may take up to four hours to charge when first installed; subsequently the battery receives a continuous trickle charge. With power interrupted on the server, the battery can maintain cache data for up to 72 hours.

Performance testing

This performance characterization describes a series of tests – carried out in lab and real-world environments – designed to identify the performance improvements that can be delivered by write cache.

These tests were as follows:

- **HP SBC performance test** – HP tested an HP ProLiant BL20p G1 server blade in an HP SBC test environment to determine the impact of write cache on %Disk Time, a Perfmon value that provides an indicator of user response times – lower %Disk Time values translate to a better user experience.
- **Session logon test** – HP tested a second HP SBC test environment to determine the impact of write cache on session logon times and user response times. Shorter logon times also translate to a better user experience.
- **Real-world performance** – This performance characterization concludes with an example of the performance gains possible when deploying write cache in a real-world environment – in this case, a call center application.

HP performance test

In March 2004, HP tested a BL20p G1 server blade with and without write cache to characterize performance gains in an HP SBC test environment.

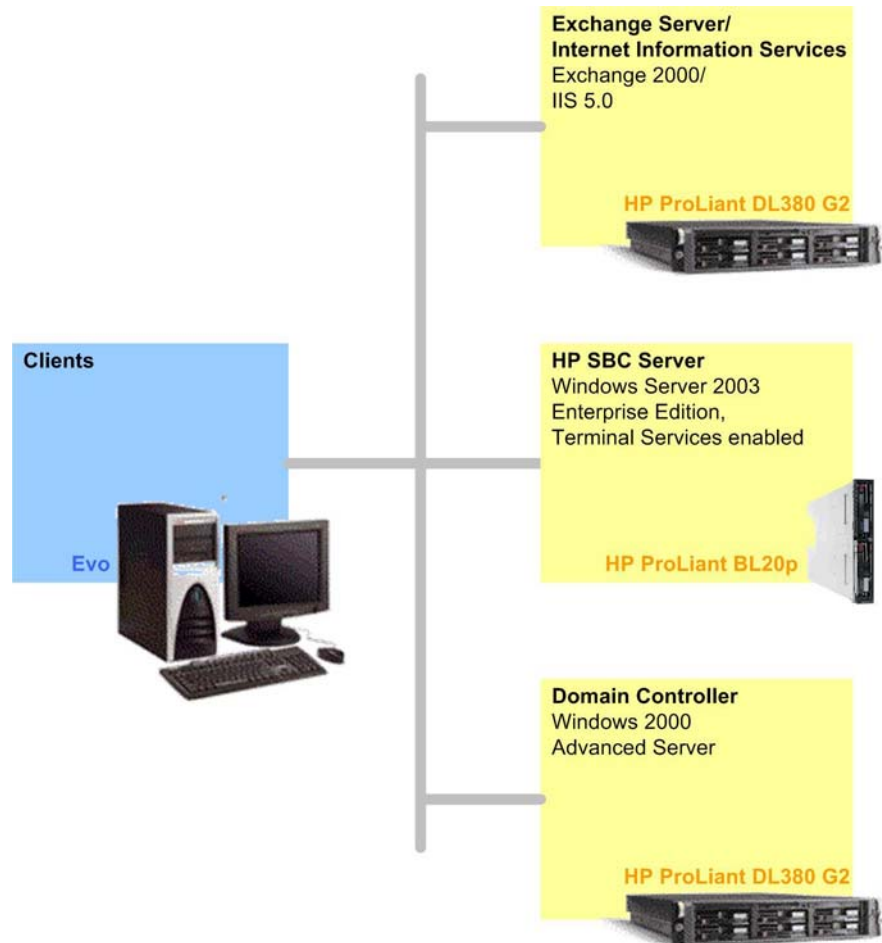
This section provides information on the following topics:

- **Tested environment** – Details of the lab environment utilized for the testing
- **Server configurations** – An overview of the servers used in the tested configuration
- **Test tools and scripts** – An overview of the tools and the test script used by HP to simulate the activities of Heavy Users
- **Test scenarios** – A description of the performance testing
- **Performance test results** – Test results for a BL20p deployed as an HP SBC server with and without write cache

Tested environment

Figure 4 illustrates the HP SBC test environment used by HP.

Figure 4. The tested environment for the HP SBC performance test



Configurations

Table 2 summarizes the configurations of systems used in the test environment.

Table 2. System configurations

Server	Configuration
Terminal Server	2P HP ProLiant BL20p G1 server blade with: <ul style="list-style-type: none"> • Intel® Pentium® III 1.4 GHz processor • 3 GB RAM • Integrated Smart Array 5i controller with optional BBWC • Two 18.2 GB hard drives • Two NC3163 Fast Ethernet NICs Windows Server 2003 Enterprise Edition, Terminal Services enabled Citrix MetaFrame XP Presentation Server FR3 Microsoft Office XP
Exchange Server/ Internet Information Services	2P HP ProLiant DL380 G2 server with: <ul style="list-style-type: none"> • 1.133 GHz Pentium III processor • 512 KB L2 cache • 2304 MB RAM • Two hot plug 36.4 GB 15,000 rpm Ultra3 SCSI hard drives • Integrated Smart Array 5i controller with RAID 1 • NC3163 10/100 Fast Ethernet NIC Windows Server 2003, Enterprise Edition Microsoft Exchange Server 2000 Microsoft Internet Information Services 5.0
Domain Controller	2P HP ProLiant DL380 G2 server with: <ul style="list-style-type: none"> • 1.4 GHz Pentium III processor • 512 KB L2 cache • 2048 MB RAM • Two hot plug 36.4 GB 15,000 rpm Ultra3 SCSI hard drives • Integrated Smart Array 5i controller with RAID 1 • NC3163 10/100 Fast Ethernet NIC Windows 2000 Advanced Server

Continued

Table 2. System configurations (continued)

Server	Configuration
Client	Variety of Pentium-based 600 MHz – 2.533 GHz Compaq Evo workstations, each with: <ul style="list-style-type: none">• At least 256 MB of memory• 640x480/256-color resolution• 10 Mbps NIC
	Windows 2000 Professional

Test tools and scripts

HP used various testing tools and scripts to simulate actual HP SBC user sessions.

Tools

Microsoft has developed a suite of tools – Terminal Services Scalability Planning Tools (TSScaling) – to help organizations test Terminal Server performance. These tools facilitate the placement and management of simulated loads on the Terminal Server.

Table 3 outlines the components of TSScaling.

Table 3: TSScaling tools

Component		Description
Automation tools	Robosrv.exe	Drives the server-side of the load simulation
	Robocli.exe	Helps drive the client-side of the load simulation
Test tools	Qidle.exe	Determines if any scripts have failed and require operator intervention
	Tbscript.exe	A script interpreter that helps drive the client-side load simulation
Help files	TBScript.doc	Terminal Server bench scripting documentation
	TSScalingSetup.doc	A scalability test environment set-up guide
	TSScalingTesting.doc	A testing guide

More information on the test tools

- Roboserver (Robosrv.exe) and Roboclient (Robocli.exe): [Terminal Server capacity planning](#)
- TSScaling: [Windows Server 2003 Terminal Server Capacity and Scaling](#)

Heavy User script

HP used a test script to simulate the activities of Heavy Users in an HP SBC environment. Heavy Users (also known as Structured Task Workers) tend to open multiple applications simultaneously and remain active for long periods of time. Heavy Users often leave applications open when not in use.

Heavy users have an average typing speed of 70 – 80 words per minute.

The Heavy User script included the following activities:

- Creating, saving and printing documents using Microsoft Office Word XP
- Browsing Web pages, and sending and receiving e-mail using Office Outlook XP
- Performing extensive Office Excel XP activity (using PivotTable dynamic views, graphing and charting data, and so on)

Test scenarios

HP used the tested configuration shown in Figure 4 to simulate an HP SBC deployment, with the Heavy User test script (described above) being utilized to simulate typical HP SBC workloads.

HP tested the configuration twice – once with write cache disabled to obtain a baseline, then with write cache enabled.

For each of these test scenarios, HP began by running the Heavy User script with a group of ten users. Start times were staggered to eliminate authentication overhead. After the sessions of the first group of users finished, HP continued to add users in groups of ten.

Performance test results

For both test scenarios (with write cache disabled or enabled), HP used Perfmon to monitor % Disk Time values as Heavy Users logged on to the HP SBC server.

Write cache disabled

Figure 5 shows test results with write cache disabled.

Figure 5. Test results with write cache disabled, showing that %Disk Time values increased in proportion to the number of Heavy User sessions

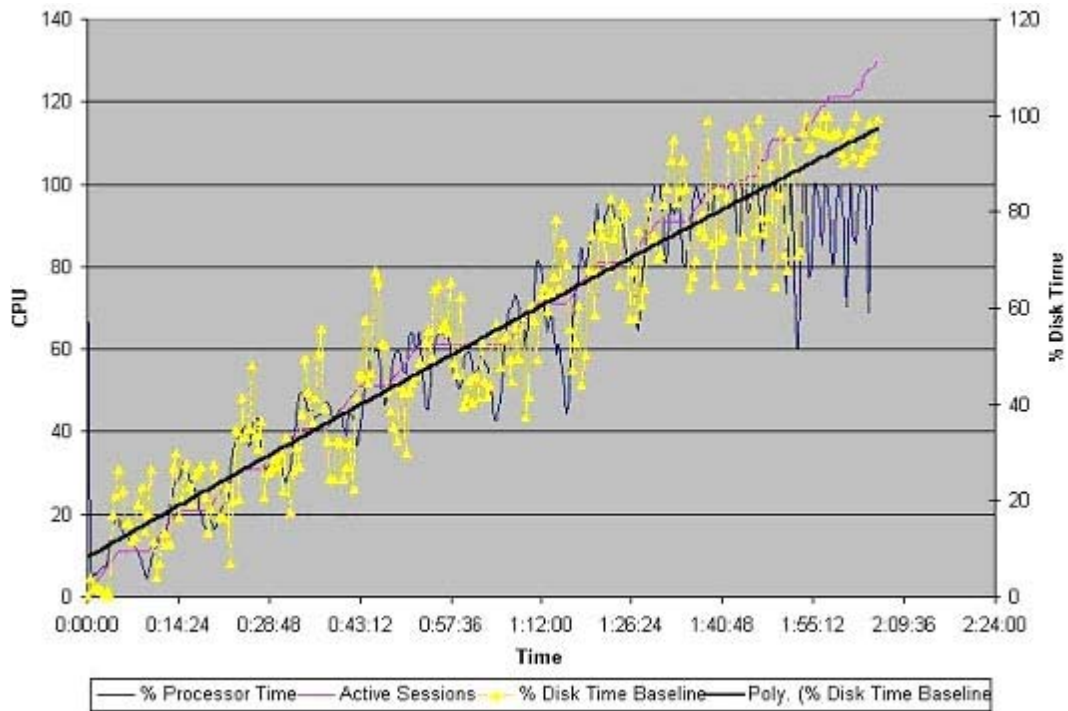


Figure 5 shows that, as more Heavy Users began sessions, % Disk Time values increased proportionally, indicating that user response times would also increase.

Write cache enabled

Figure 6 shows test results with write cache enabled.

Figure 6. Test results with write cache enabled

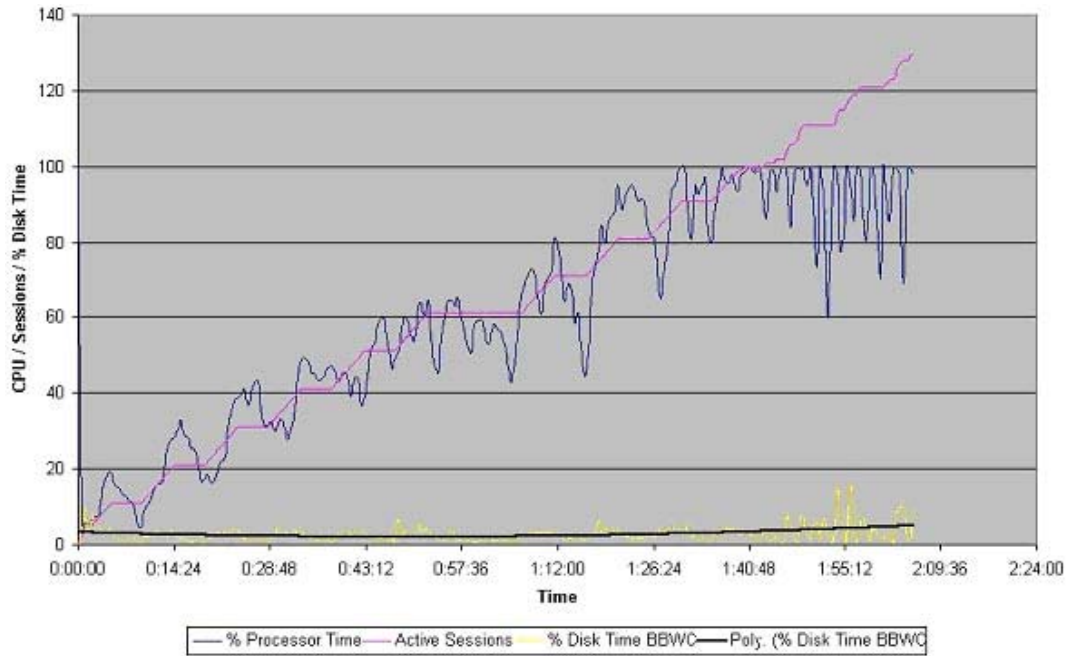
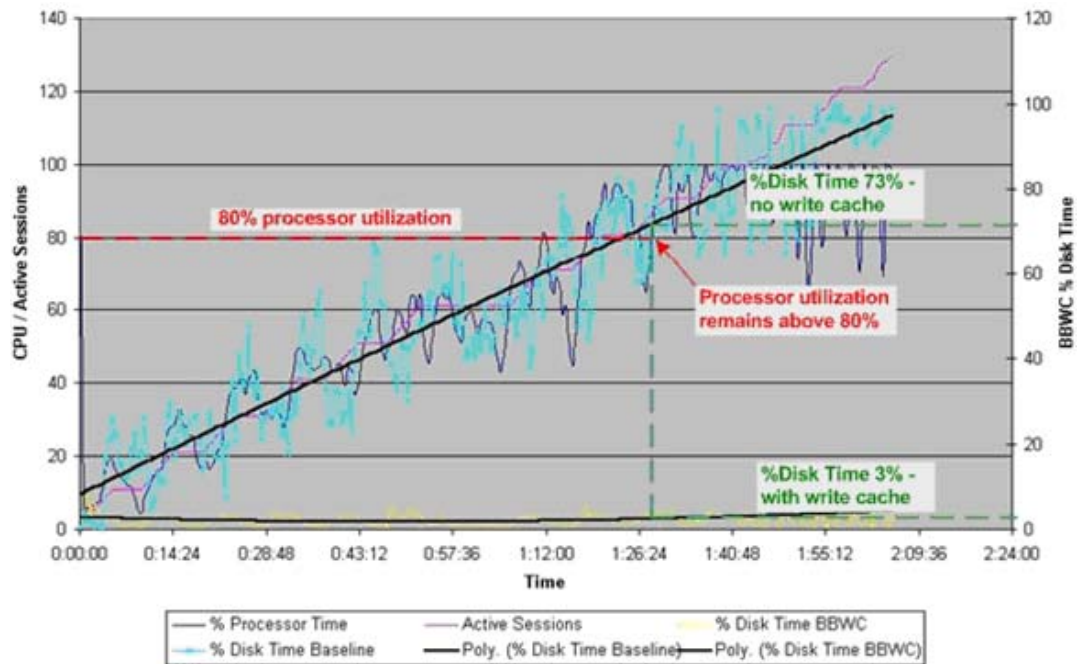


Figure 6 shows that, as more Heavy Users began sessions, % Disk Time values still tended to remain below 5%, indicating a uniform user experience with minimal impact from disk I/O.

Results summary

Figure 7 directly compares the test results.

Figure 7. Test result comparison, showing the dramatic improvement with write cache enabled



Results for the test carried out with write cache enabled are shown in yellow; results for the test carried out with write cache disabled are shown in light blue.

Figure 7 shows that, with the optimal number¹ of users logged on to a BL20p server blade, response times are 70% faster with write cache enabled.

Important:

The 70% performance improvement described above is based on a specific test environment. Actual values achieved in a production environment will depend on server platform, configuration, and application.

¹ Traditionally, HP has regarded the number of users active when processor utilization reaches 80% to be the optimal number of users.

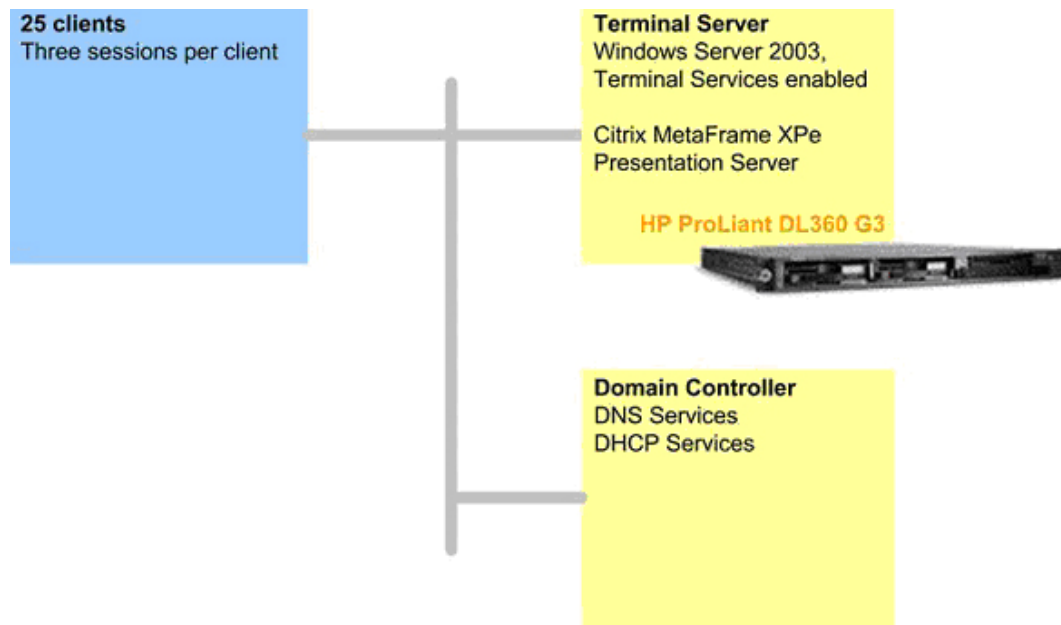
Session logon test

HP measured the time taken for a sample session logon in various HP SBC scenarios to determine the impact of write cache on logon times and user response times.

Test environment

Figure 8 illustrates the HP SBC environment tested by HP.

Figure 8. Test environment for the logon session testing



Test tools and scripts

HP used the following testing tools and scripts to launch simultaneous SBC user sessions on each client machine:

- An internal Citrix test tool launched three simultaneous user sessions on each client
- A logon script automatically logged the sessions off after 90 seconds

Test scenarios

Two test scenarios were utilized – write cache disabled and enabled.

Write cache disabled

1. Using the Citrix test tool, configure 25 client machines to launch three user sessions each.
2. Run the logon test with groups of five, 10, 15, 20, and then all 25 clients.
3. Using a stopwatch, measure the time taken for a sample logon session on one client machine.
4. Use Perfmon to monitor disk, memory, and processor utilization on the HP SBC server.

Write cache enabled

1. Install a BBWC enabler on the HP SBC server.
2. Wait three hours for the battery to fully charge.
3. Repeat the test scenarios.

HP performed three runs for each scenario, obtaining metrics for logon times and disk utilization (an indicator of user response times).

Test results

HP obtained results for logon times and user response times.

Logon times

Table 4 shows single session logon times with write cache disabled and enabled. These results reflect average values for the three runs.

Table 4. Single session logon times

Number of clients	Number of client logons	Session logon time (seconds)	
		Write cache disabled	Write cache enabled
5	15	16	30
10	30	25	16
15	45	43	38
20	60	57	51
25	75	80	60

Table 4 shows logon times improving by up to 33% with write cache enabled.

Response times – write cache disabled

Figure 9 shows the performance of the HP SBC server with write cache disabled over two iterations of the logon test.

Figure 9. Disk utilization averages approximately 20% without write cache

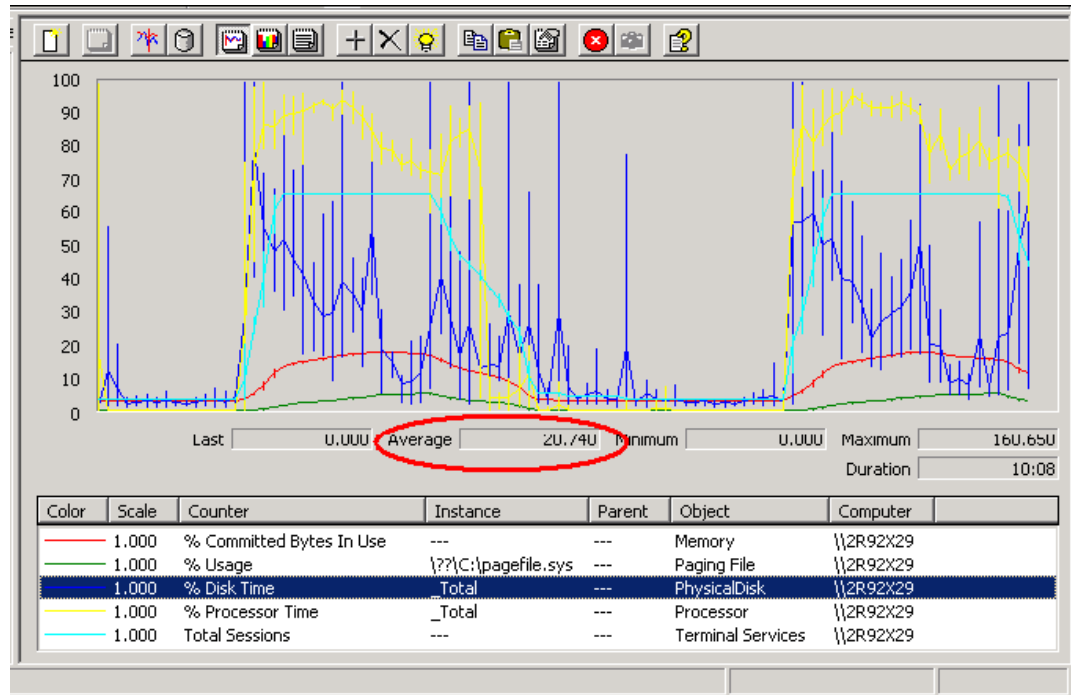


Figure 9 shows disk utilization (%Disk Time) averaging 20.74%.

Response times – write cache enabled

Figure 10 shows the performance of the HP SBC Server with write cache enabled over two iterations of the logon test.

Figure 10. Disk utilization averages approximately 3% with write cache enabled

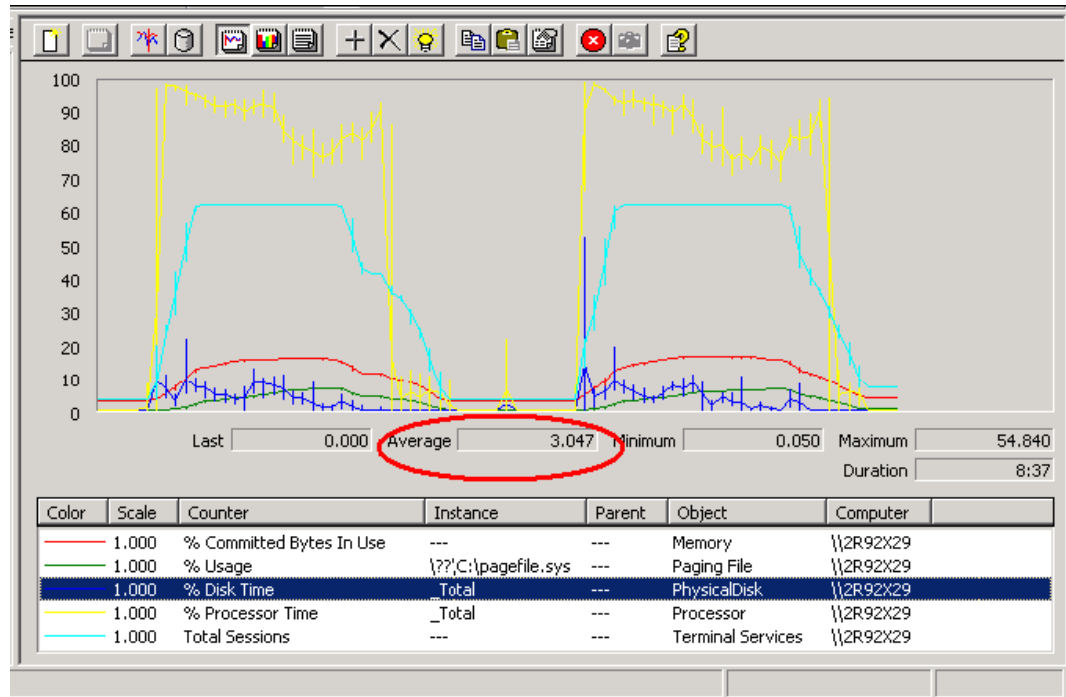


Figure 10 shows that enabling write cache reduces the average %Disk Time value from 20.74% to 3.047% (and correspondingly improves user response times). The values of other counters are essentially unchanged.

Real-world performance

This performance characterization also demonstrates the value of enabling write cache in a real-world application – a call center.

In this example, the call center deployed an HP ProLiant server with three different write array configurations:

- Smart Array 5304 controller with 256 MB write cache
- Smart Array 5i Plus controller with 64 MB write cache
- Smart Array 5i controller with write cache disabled

Call center IT staff measured the maximum numbers of users supported by each configuration, as shown in Figure 11.

Figure 11. Real-world call center performance, showing the value of write cache



Figure 11 shows that, in this call center environment, the maximum numbers of users supported were as follows:

- HP Smart Array 5i controller with no write cache – 39 users
- HP Smart Array 5i Plus controller with 64 MB write cache – 51 users
- HP Smart Array 5304 controller with 256 MB write cache – 96 users

The test results indicate that 64 MB write cache translated to support for up to 31% more users. Deploying additional write cache (256 MB) translated to an even more significant improvement – up to 246% more users.

Conclusions

The write cache installed on HP Smart Array controllers can deliver significant performance enhancements by temporarily buffering data written to disk.

Performance tests carried out in a lab environment showed various performance metrics improving by as much as 70%. Real-world metrics obtained in a call center environment showed that even larger improvements are possible – as much as 246% in this example.

To obtain these improvements, however, customers must first deploy the BBWC enabler (which, depending on the HP ProLiant server model, may be standard, optional, or bundled with a Smart Array Controller).

For more information

HP ActiveAnswers for Server Based Computing, <http://www.hp.com/solutions/activeanswers/citrix>

HP ProLiant servers, <http://www.hp.com/go/proliant>

HP Smart Array controllers,

<http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/index.html>

HP Array Configuration Utility,

<http://h18004.www1.hp.com/products/servers/proliantstorage/software-management/acumatrix/index.html>

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback.

© 2004, 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, and Pentium are trademarks of Intel Corporation in the U.S. and other countries.

4AA0-8805ENW, Revision 3, September 2007

