



HP Integrity Superdome X Server delivers proof points for server consolidation

April 2015

Executive summary

As the need to process more and more mission critical data continues to increase, customers need a platform that can grow and scale at the pace of their business. The HP Integrity Superdome X Server is the ideal platform to support the most demanding enterprise workloads, delivering mission-critical reliability, availability, and serviceability (RAS) features in combination with blazing performance and cost-effective industry-standard x86 computing power. Superdome X delivers superior single-scale Online Transaction Processing (OLTP) performance of 1.7X as more resources are added to the system, from two to four to eight sockets. This platform also delivers the same 1.7X scaling when consolidating multiple workloads all the way out to its maximum industry-leading 16-socket size.

Revolutionary x86 architecture to optimize performance

Superdome X merges the best of Superdome and ProLiant technologies to create a revolutionary new x86 architecture designed to maximize performance and availability. With breakthrough innovations such as the fault-tolerant Crossbar Fabric and Error Analysis Engine coupled with end-to-end implementation of the Intel® Xeon® Processor E7 v2 RAS features, Integrity Superdome X sets the standard for mission-critical x86 computing. Delivered in an efficient bladed form factor from two to sixteen sockets, and with up to 12TB of memory, Superdome X grows seamlessly with the business and can address all traditional and in-memory database needs. Featuring x86 hard partitioning, HP nPars optimize license costs, lower downtime, and maximize resource utilization, with multiple nPars of different sizes supported. Each HP BladeSystem Superdome Enclosure supports 15 fans, 12 power supplies, and four HP Crossbar Fabric Modules (XFMs) to drive one to eight mission-critical, scalable, dual-socket x86 blades assuring total redundancy for even the largest configurations.

OLTP test comparison

HP conducted OLTP tests to demonstrate single-instance database scaling and performance on Superdome X. In addition, multiple mission-critical OLTP workloads were consolidated and simultaneously tested to show the growth potential of a fully configured Superdome X. The first set of tests used two-socket, four-socket, and eight-socket (2S, 4S, and 8S) configurations to compare results against the OLTP capability of an 8S HP ProLiant DL980 G7 Server. The ProLiant DL980 G7 Server was used as a benchmark due to its established leadership position with multiple 8S performance records. HP also tested multiple OLTP database workloads executing on eight-socket (8S) and then sixteen-socket (16S) configurations, providing a view of both the size and scale of mission-critical workloads that Superdome X can handle. Figures 1 and 2 show the results of the testing.

Key take aways: Scale with confidence without sacrificing performance

- **HP Integrity Superdome X Server achieved excellent scaling from 2S to 4S to 8S:**
 - 1.67X scaling from 2S to 4S
 - 1.6X scaling from 4S to 8S
 - 2.4X performance superiority delta at 8S compared to the HP ProLiant DL980 G7 Server
- **HP Integrity Superdome X Server achieved outstanding scaling from 8S to 16S, using Oracle Multitenant database configurations:**
 - 1.7X scaling from 8S to 16S
 - 4.36X performance superiority delta at compared to the HP ProLiant DL980 G7 Server

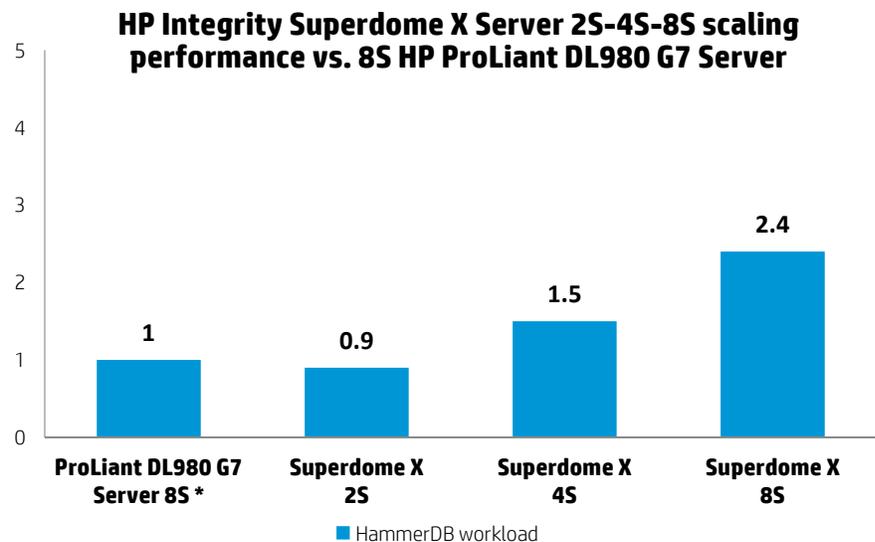
HP Integrity Superdome X Server configuration

- Red Hat Enterprise Linux 6.6
- Oracle 12c OLTP HammerDB workload
- External storage for database
- SSD drives for log and data
- Separated storage partitions for log and data
- Balanced performance across NUMA nodes
- Up to 8S NUMA configuration
- 15c/30 CPUs per socket (HT enabled)

Results

- Superior performance
- 1.6X scaling from 2S to 4S to 8S
- 90% CPU utilization

Figure 1. HP Integrity Superdome X Server 2S-4S-8S scaling vs. 8S HP ProLiant DL980 G7 Server



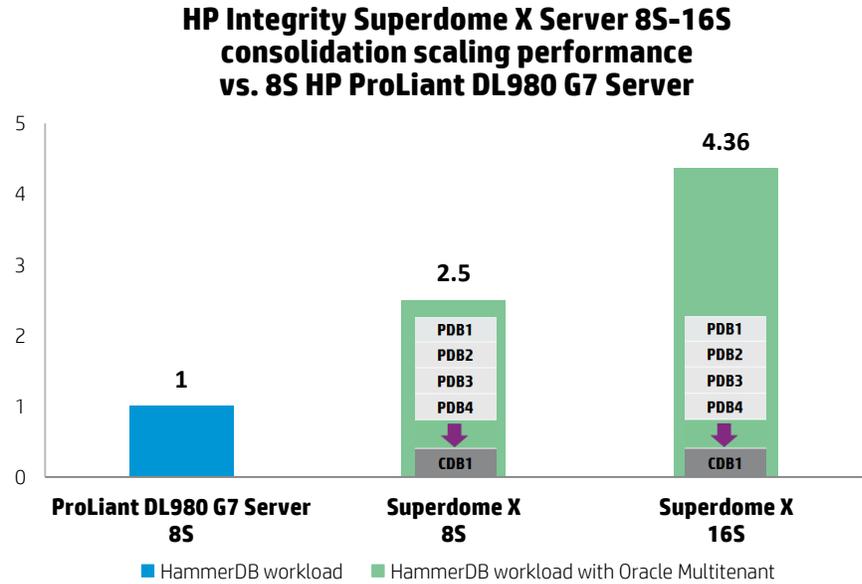
HP Integrity Superdome X Server multitenant container database configuration

- Red Hat Enterprise Linux 6.6
- Oracle 12c OLTP HammerDB workload
- Shared external storage for databases
- SSD drives for log and data
- Oracle Multitenant, with four Pluggable Databases (PDBs) in one Container Database (CDB)

Results

- 1.74X scaling of Integrity Superdome X from 8S to 16S
- Balanced performance across NUMA nodes
- Lower TCO by consolidating workloads
- 4.36X performance delta compared to the ProLiant DL980 G7 Server

Figure 2. Scaling performance of HP Integrity Superdome X Server 8S-16S using multitenant vs. HP ProLiant DL980 G7 Server (non-multitenant)

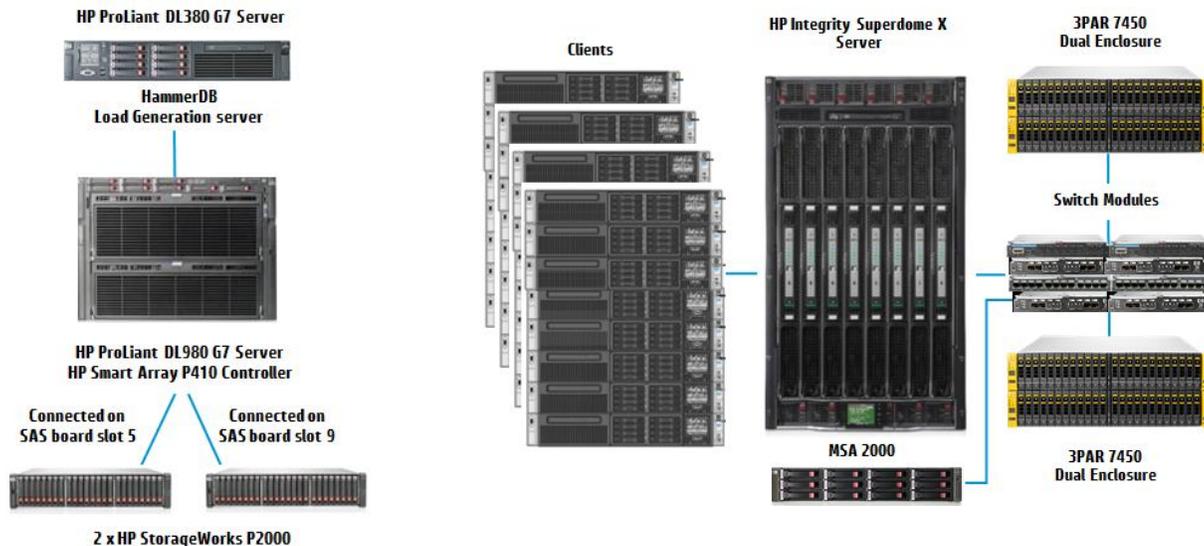


HP Integrity Superdome X Server and HP ProLiant DL980 G7 Server benchmark configurations

HP Integrity Superdome X Server, HP 3PAR StoreServ 7450 All-flash array, Oracle Database 12c, and the Red Hat Enterprise Linux 6.6 operating system were configured to run an OLTP benchmark as shown in Figures 1 and 2. The test benefitted from the low latency capabilities of the HP 3PAR StoreServ 7450 storage with Solid State Disks (SSDs). For all the tests, the workload was HammerDB, an I/O intensive OLTP benchmark tool that could stress the server as well as the I/O subsystem. All transactional numbers were normalized, and only a trend of transactions is shown.

Database performance testing with the HP ProLiant DL980 G7 Server was based on an industry-standard OLTP type of workload mixing read-only and update intensive transactions to simulate the activities found in complex OLTP application environments. These tests simulated the high transaction rate Order entry system of a wholesale supplier company having a number of geographically distributed sales districts and associated warehouses. The chosen database loading tool was HammerDB, a leading load generation tool for many databases, including Oracle and MySQL. HammerDB applies an OLTP workload-based upon industry-standard benchmark specifications. It comes with pre-built simulations and can provide multi-user workloads out of the box. HammerDB is available under a public license and its use is free of charge. Figure 3 shows the setup for each system tested.

Figure 3. Setups for measuring performance of the HP ProLiant DL980 G7 Server and HP Integrity Superdome X Server.



Summary

HP Integrity Superdome X Server

The HP Integrity Superdome X Server offers scalability that surpasses the market, flexibility through HP nPars, and mission-critical reliability, availability, and serviceability. Blending industry-standard efficiencies with proven innovations to deliver groundbreaking performance and a superior uptime experience, Superdome X sets the standard for mission-critical x86 computing.

HP 3PAR StoreServ 7450 All-flash storage array

HP 3PAR StoreServ 7450 is an all-flash-based storage array that combines accelerated performance with all the enterprise-class, tier-1 features and functionality expected for a mission-critical application environment. This flash-optimized architecture is founded on several unique HP 3PAR StoreServ innovations.

Red Hat Enterprise Linux 6.6 operating system

Built for the modern datacenter environment, Red Hat Enterprise Linux anticipates shifts in the IT landscape that blur the lines between physical, virtual, and cloud computing. With long-standing market leadership and a 14-year partnership, HP and Red Hat actively collaborate to engineer and deliver solutions, reference architectures, and services supported by joint investment in business development and marketing. HP delivers Linux solutions that combine the security, dependability, value, and flexibility of Open Source with enterprise-quality management applications, services, communication tools, and hardware. As a leader in Linux revenue and units worldwide, HP is a proven and trusted choice for Linux requirements.¹

Bottom line

As demonstrated by the HP benchmark testing results addressed in this paper, the HP Integrity Superdome X Server delivers blazing performance, scaling evenly at about 1.7X from two to four to eight sockets. It also proves itself as an ideal server consolidation platform with 1.7X scaling from eight to sixteen sockets. Superdome X offers new possibilities for customers who want to maximize performance or consolidate critical, demanding workloads. HP delivers the right compute for the right workload at the right economics, every time.

Best practices for performance

Analysis and recommendations

Default firmware and storage settings were used in this platform configuration, including default firmware settings. For best I/O performance on the HP Integrity Superdome X Server, HP recommends using multiple paths to maintain high availability while also maximizing performance and minimizing latencies. A way to maximize performance in extremely demanding environments is to minimize inter-communication between NUMA nodes. This can be achieved using tightly zoned hardware configurations and operating system-to-hardware configurations, such as setting CPU affinity to minimize latencies across the NUMA nodes.

The approach taken in this effort to achieve maximum I/Os and throughput was to connect and zone the server to storage in such a way that cross-node activity is minimized both from the server and the storage. By dedicating virtual volumes to specific HBAs and NUMA nodes, all of the I/O for a specific volume stays local to specific storage nodes and server nodes. For applications that do a good job with NUMA awareness, this can deliver extremely good performance. For those applications that are not as good with NUMA awareness, more manual tuning may be required, but the flexibility to tune the environment exists.

Note: For zoning best practices, see the technical white paper [HP Universal Database Solution: Oracle and HP 3PAR StoreServ 7450 All-flash array](#).

Table 1. Best practices recommendations.

Server	OS and Kernel configuration: Disable cstates at kernel boot.
Storage	UDEV settings for performance: Set udev parameters per values in appendix B. <ul style="list-style-type: none"> • SSD loading – Load SSDs in groups of four per enclosure at a minimum • Volume size – Virtual volumes should all be the same size and SSD type for each Oracle ASM group. • Use Full Provisioning for the storage of the database and logs
RHEL 6.6 OS	See appendix C for details.
Oracle database	<ul style="list-style-type: none"> • ASM – Use separate DATA and LOG ASM groups • Logs – Assign at least three volumes to the log group • Oracle parameters <ul style="list-style-type: none"> – Set HUGE pages only – Disable automatic memory management if applicable – Set buffer cache memory size large enough per implementation to avoid physical reads – Enable NUMA support

¹ Source: IDC Worldwide Quarterly Server Tracker for 4Q14, March 2015. See hp.com/go/therealstory for details.

Appendixes

Appendix A. Server configurations

Table 2. HP ProLiant DL980 G7 Server and HP Integrity Superdome X Server benchmark configurations.

	ProLiant DL980 G7 Server	Superdome X straight-scaling configuration	Superdome X multitenant configuration
Processors	8 x Intel® Xeon® Processors X7560 (8c/2.26 GHz/24M/130 W)	2, 4, or 8 x Intel Xeon Processors E7-2890 v2 (15c/2.8 GHz/37.5M/155 W)	8 or 16 Intel Xeon Processors E7-2890 v2 (15c/2.8 GHz/37.5M/155 W)
Memory	512GB RAM	16x32GB 2Rx4 PC3-12800R DIMMs per blade 2S: 512GB 4S: 1TB 8S: 2TB; DIMMs were installed across all memory sockets of 8 CPUs	16x32GB 2Rx4 PC3-12800R DIMMs per blade 8S: 2TB; DIMMs were installed across all memory sockets of 8 CPUs 16S: 4TB; DIMMs were installed across all memory sockets of 16 CPUs
OS storage	Embedded controller	2x400GB SAS SSD 1 HP QMH2672 16Gb FC HBA mezzanine	2x400GB SAS SSD 1 HP QMH2672 16Gb FC HBA mezzanine
Data and log storage	2 HP Smart Array P410 RAID controllers, 1 for data and 1 for log External storage: 2 x P2000, 1 for data and 1 for log 6x120GB SSD RAID 0	External storage: 2 HP 3PAR 7450 StoreServ all-flash array four-enclosure units RAID 1: 16x256GB partitions 18x4 SAS cMLC SSD per 3PAR Total 480GB All configurations: 1 QLogic 16Gb 2672 for OS 2S: 2 QLogic 16Gb 2672 4S: 4 QLogic 16Gb 2672 8S: 8 QLogic 16Gb 2672 2 FC mezzanine HBAs in x16 slots (slot numbers 2 and 3) (Type B) per blade	External storage: 2 HP 3PAR 7450 StoreServ all-flash array four-enclosure units RAID 1: 16x256GB partitions 18x4 SAS cMLC SSD per 3PAR Total 480GB Both configurations: 1 QLogic 16Gb 2672 for OS 8S: 8 QLogic 16Gb 2672 16S: 16 QLogic 16Gb 2672 2 FC mezzanine HBAs in x16 slots (slot numbers 2 and 3) (Type B) per blade
Database data and log storage	Oracle redolog files : HP Smart Array Controller P410 raw device 6x120GB SDD RAID 0 or HP StorageWorks P2000 SAS card upper half of HP ProLiant DL980 G7 Server Oracle database files: P2000 SAS card in ProLiant DL980 G7 lower half	Database data and log stored on the two 3PAR StoreServ 7540 All-flash enclosures 1TB available redolog capacity 4TB user database space 128GB available database temp space	Database data and log stored on the two 3PAR StoreServ 7540 All-flash enclosures 1TB available redolog capacity 4TB user database space 128GB available database temp space
Storage zoning	No zoning	Storage zoning with switches 5, 6, 7, 8 in the server interconnect module 1 and 2 for network, 3 and 4 for OS drives 5-8 for FC connections to 3PAR 7450	Storage zoning with switches 5, 6, 7, 8 in the server interconnect module 1 and 2 for network, 3 and 4 for OS drives 5-8 for FC connections to 3PAR 7450
RHEL version	Red Hat Enterprise Linux 5.6 x64	Red Hat Enterprise Linux 6.6 x64	Red Hat Enterprise Linux 6.6 x64
Database	Oracle 11g Release 2	Oracle 12c	Oracle 12c
Clients	1 HP ProLiant DL380 G7 Server	8 ProLiant DL360 Gen8 servers	8 and 16 ProLiant DL360 Gen8 servers

For more information about Superdome X, see the [Quick Specs](#).

Appendix B. UDEV configurations

Below is an example of the udev configurations used in the testing.

```
/etc/udev/rules.d/10-3par.rules
ACTION=="add|change", KERNEL=="dm-*", PROGRAM="/bin/bash -c
'cat /sys/block/$name/slaves/*/device/vendor | grep 3PARdata'", ATTR{queue/rotational}="0",
ATTR{queue/scheduler}="noop", ATTR{queue/rq_affinity}="2", ATTR{queue/nomerges}="1"
/etc/udev/rules.d/12-dm-permissions.rules
ENV{DM_NAME}=="mpathch", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcg", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcf", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathce", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathag", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcd", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathaf", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcc", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathae", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcb", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathad", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathac", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcl", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathck", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathcj", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
ENV{DM_NAME}=="mpathci", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
```

Appendix C. Red Hat Enterprise Linux 6.6 operating system settings

Recommended Linux boot-time settings

The Linux boot parameter keeps the processing cores in C1 state when used in conjunction with "intel_idle.max_cstate=1." For this OLTP workload, and environment, the C1 state setting provides a small performance gain.

For RHEL systems, edit /boot/grub/grub.conf and add intel_idle.max_cstate=1 " parameter as mentioned previously.

Preparing Linux-based servers for low-latency tuning

Before configuring the server for low latency, do the following:

1. Make the following edits:
 - a. For Red Hat Enterprise Linux systems:
 - b. Edit /boot/grub/grub.conf and add "intel_idle.max_cstate=1" to the kernel line
2. Reboot the server.
3. After reboot, run the services-off.sh script to stop extraneous services. The following example stops the services shown and prevents them from starting on subsequent boots:

```
for SERVICE in \
NetworkManager abrt-ccpp abrt-d acpid auditd autofs certmonger cups kdump nfslock openct pcsd
portreserve ostfix rhnsd rhsmcertd
chkconfig --level 123456 $SERVICE off
service $SERVICE stop
done
```

4. Use the irqbalancer to preclude some cores from servicing software IRQs:
 - a. Enter the following command:


```
# service irqbalance stop
```
 - b. Do a one-time run of the irq balancer:


```
# irqbalance --oneshot -h exact
```

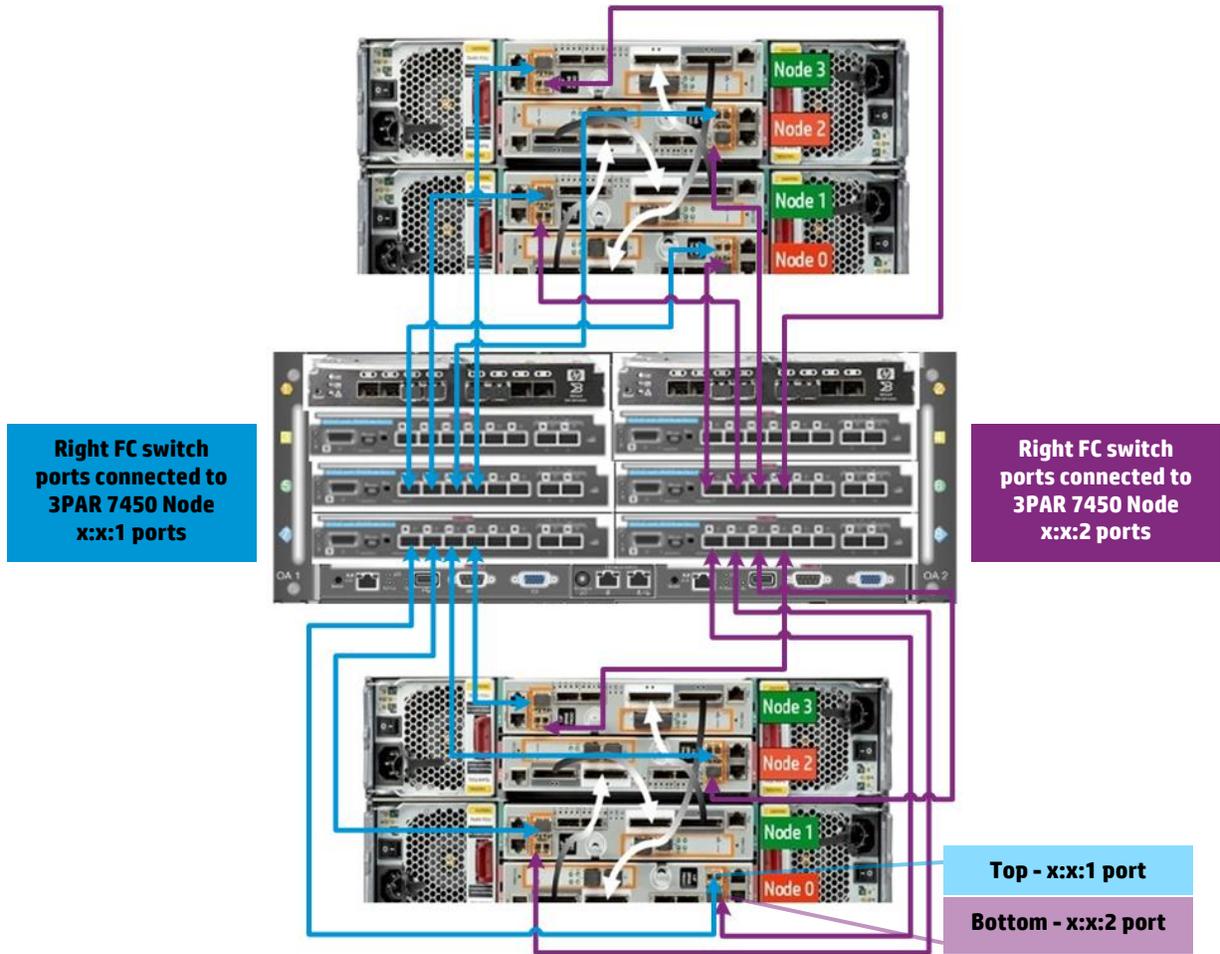
Wait until the command service irqbalance status returns "irqbalance is stopped."
5. For low-latency performance, we also set the latency-performance as shown:


```
# tuned-adm profile latency-performance
```

Appendix D. HP Integrity Superdome X Server test external storage to switch module connections

The switch module is located within the HP Integrity Superdome X Server enclosure. The connections between the HP 3PAR StoreServ 7450 All-flash array and the Fibre Channel switch modules are shown in Figure 4. For more information about storage zoning, refer to the [HP 3PAR Command Line Interface Administrator's Manual: HP 3PAR OS 3.1.2](#).

Figure 4. HP 3PAR StoreServ 7450 All-flash array rear view Fibre Channel connection diagram.



For more information

- HP ProLiant performance: hp.com/servers/benchmarks
- HP Integrity Superdome X: hp.com/servers/superdomex
- HP Integrity Servers: hp.com/go/integrity

Sign up for updates
hp.com/go/getupdated



Share with colleagues

© Copyright 2015 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.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc in the United States and other countries. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries.

