FAULT TOLERANCE PERFORMANCE AND SCALABILITY COMPARISON: NEC HARDWARE-BASED FT VS. SOFTWARE-BASED FT

Hardware-based fault tolerance that performs

with the NEC Express5800/R320d-M4 server powered by Intel®

Because no enterprise can afford downtime or data loss when a component of one of their servers fails, fault tolerance is vital. While many effective software-based fault-tolerance solutions are available, a hardware-based approach such as that employed by the NEC Express5800/R320d-M4 servers, powered by Intel Xeon® processors E5-2670 v2, can offer uninterrupted service in the event of an outage without compromising performance.

In the Principled Technologies datacenter, we set up virtual machines running database workloads using two solutions: (1) an NEC Express5800/R320d-M4 server with hardware-based fault tolerance and (2) a pair of NEC Express5800/R120d-M1 servers using VMware® vSphere® for fault tolerance.

We found that when each solution ran eight simultaneous VMs, the hardware-based solution achieved more than twice the performance of the software-based solution—processing 2.4 times the number of database orders per minute—and was able to recover from a service interruption with zero downtime or loss of performance.

This sustained strong performance across a high number of VMs in a fault-tolerant environment is an enormous asset to your business. You can get more work done with less hardware, save on datacenter space and related expenses, and be assured that you are protected.

OCTOBER 2015

A PRINCIPLED TECHNOLOGIES TEST REPORT

Commissioned by NEC Corp.

EXECUTIVE SUMMARY

Enterprises need their servers to run mission-critical applications reliably. Because any server component is subject to failure, it is essential to employ some form of fault tolerance. In a fault-tolerant computer system, the failure of a component doesn't bring the system down; rather, a backup component or procedure immediately takes over and there is no loss of service.

There are two primary approaches to fault tolerance: it can be provided with software or embedded in hardware. In the Principled Technologies datacenter, we tested two fault-tolerant server solutions:[1]

- NEC Express5800/R320d-M4 servers, powered by Intel Xeon E5-2670 v2 processors, which employ hardware fault tolerance
- NEC Express5800/R120d-M1 servers, also powered by Intel Xeon E5-2670 v2 processors, using VMware vSphere software-based fault tolerance

This report explores how well the two solutions performed and scaled with one, two, four, and eight fault-tolerant VMs.[2] To compare the performance of the two solutions, we used a benchmark that simulates an OLTP database workload and reports results in terms of orders per minute. As Figure 1 shows, when running a single VM, the hardware-based FT on NEC Express5800/ R320d-M4 outperformed the software FT solution by 28.9 percent. As we added more simultaneous VMs, this advantage increased until, with eight VMs, it delivered 2.48 times the number of OPM.

[1] On the NEC Express5800/R320d-M4, we used VMware 5.5, the latest version NEC supported at the time of testing; NEC plans to extend support to VMware vSphere 6 at a future date. On the Express5800/R120d-1M, we used VMware 6 as it was the most up-to-date implementation of software fault tolerance at the time of testing.

[2] In a companion report, available at www.principledtechnologies.com/NEC/Fault_tolerance_setup_1015.pdf, we compare the relative ease of setting up the two solutions and using them to configure eight fault-tolerant VMs.

Figure 1: At the highest VM count, the hardware-based FT on NEC Express5800/R320d-M4 delivered more than 2.4 times as many orders per minute as the software-based FT solution did.

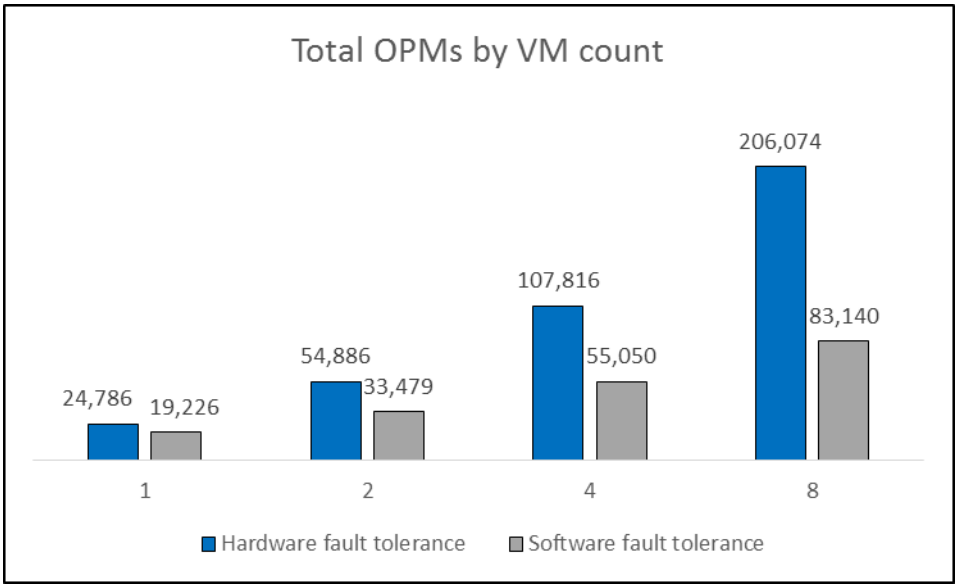

Being able to perform a greater workload while maintaining fault tolerance makes the NEC Express5800/R320d-M4 servers an attractive option. This allows the end user to obtain maximum performance while having the reliability of a fault-tolerant solution.

SOFTWARE-BASED FAULT TOLERANCE VS. HARDWARE-BASED FAULT TOLERANCE

Introduced in ESX® 4.0, the VMware vSphere fault tolerance package is designed to allow vital virtual machines to maintain greater uptime. It does so by running two virtual machines simultaneously: one VM on the primary host and a second VM on a backup host. If the primary host fails, the VM quickly and silently changes over to the backup host, preventing a loss of data.

The hardware-based fault tolerance in the NEC Express5800/R320d-M4 works differently. Its two servers operate in lockstep with each other, from their hard drives (each disk is mirrored in a RAID 1 with the disk on the other server) to their CPUs. Using a special FT appliance to achieve this, the two servers operate as one, and present themselves as one server to all other machines. In this way, any virtual machine placed on the Express5800/R320d-M4 machines is automatically fault tolerant.

FEWER NETWORKS WITH NEC HARDWARE-BASED FT

Because fault tolerance is incorporated into the server itself, the NEC Express5800/R320d-M4 obviates the need for a dedicated 10Gb network. In terms of hardware, the Express5800/R320d-M4 needs only itself and a 1Gb switch to be fully FT, external storage is optional (see Figure 3). For our testing, we chose an external iSCSI array to keep both hardware-FT and software-FT environments as comparable as possible. In contrast, software-FT requires external storage and at least one dedicated 10Gb network port (see Figure 4).

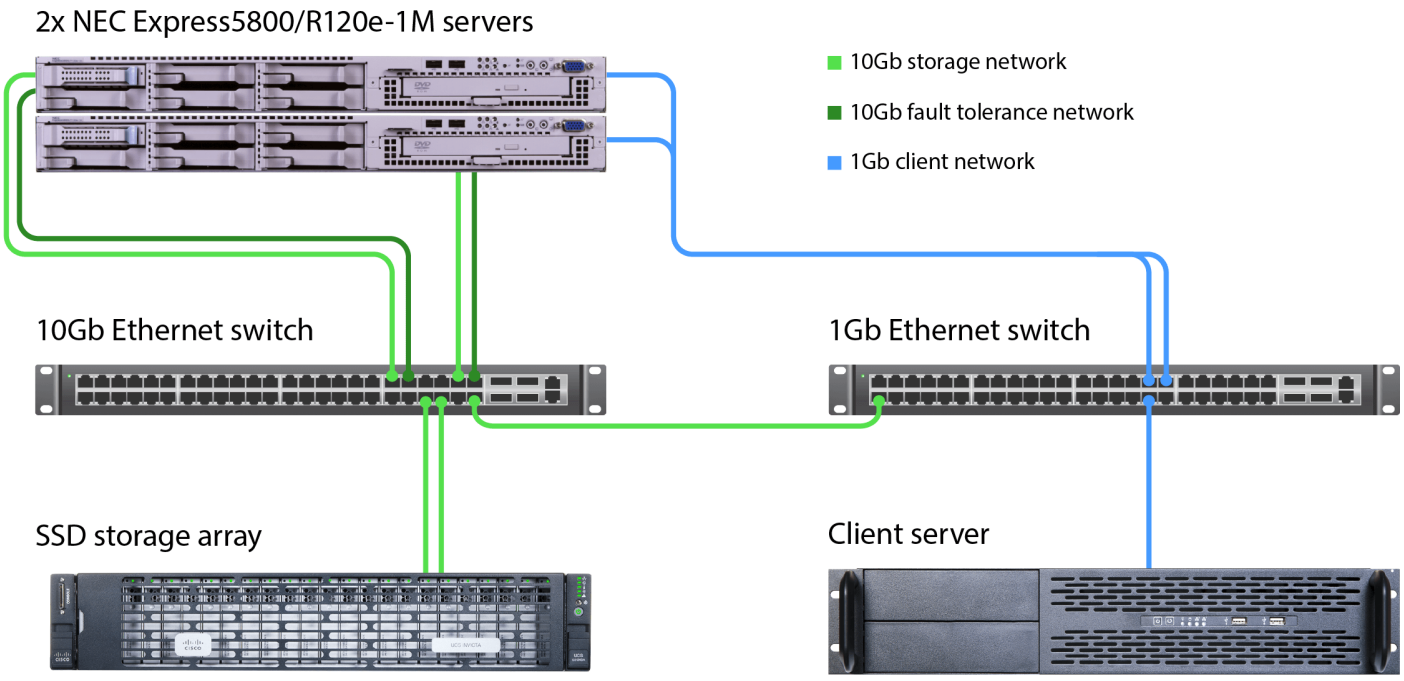

Figure 3: Testbed diagram for the NEC Express5800/R320d-M4 servers.

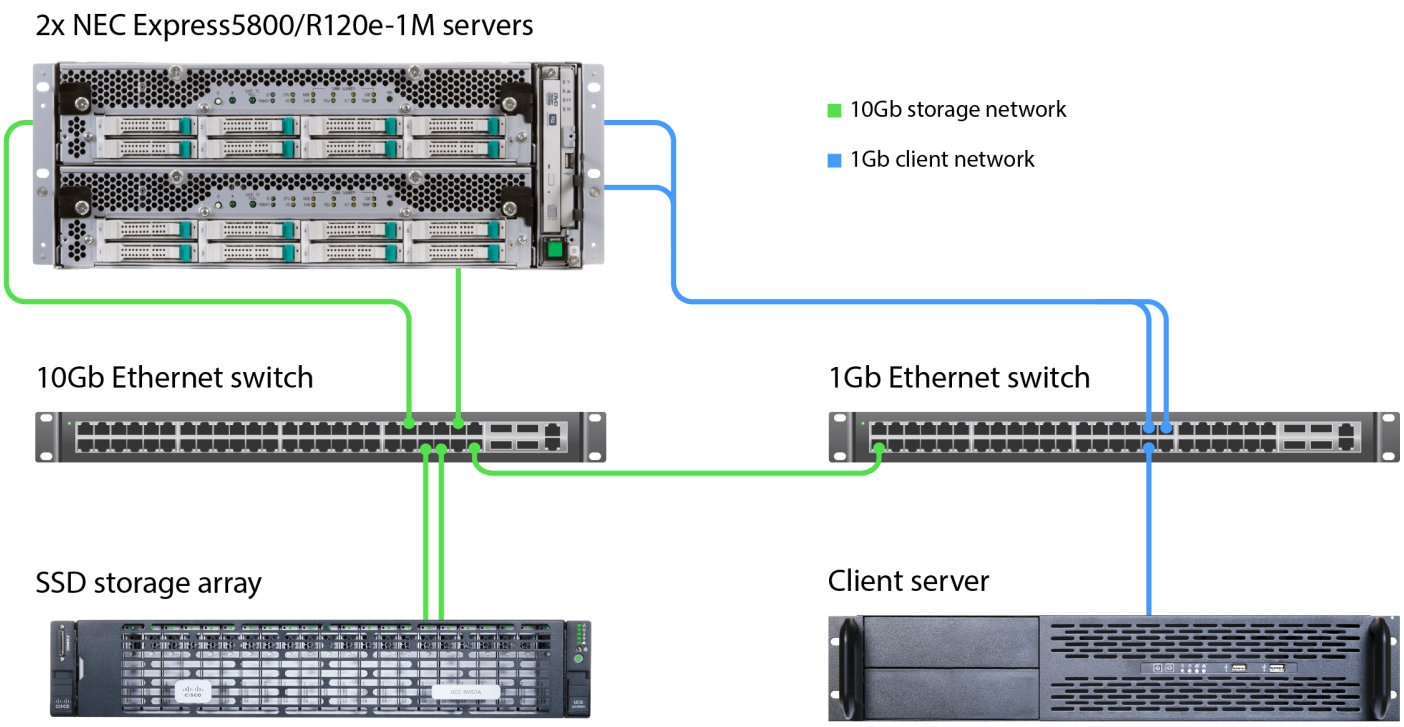

Figure 4: Testbed diagram for the NEC Express5800/R120d-M1 servers.

MORE VMS WITH NEC HARDWARE-BASED FT

As our test results show, the software-based solution we tested does support exceeding four VMs and eight vCPUs. However, VMware does not recommend doing so, and in fact required us to disable the following two settings:

- das.maxFtVmsPerHost
- das.maxFtVCPUsPerHost

While we needed to change these settings only once, doing so was a process that added time and steps to the initial setup. In contrast, the hardware-based NEC solution fully supports eight or more VMs.

SIGNIFICANTLY LESS NETWORK TRAFFIC WITH NEC HARDWARE-BASED FT

For the software-based FT solution to work, it must perform continual backups of the VMs between hosts. This volume of network traffic requires a dedicated 10 Gigabit network infrastructure and dedicated ports on both servers. Because we suspected this traffic was a factor in the lower performance we saw in our testing, we decided to measure it. Figure 5 shows network traffic in Mbits/sec over a 45-minute period. As it shows, the 10 Gigabit network was nearly saturated and was possibly contributing factor to the software-based FT solution not being able to scale as high as the NEC hardware-based FT solution.

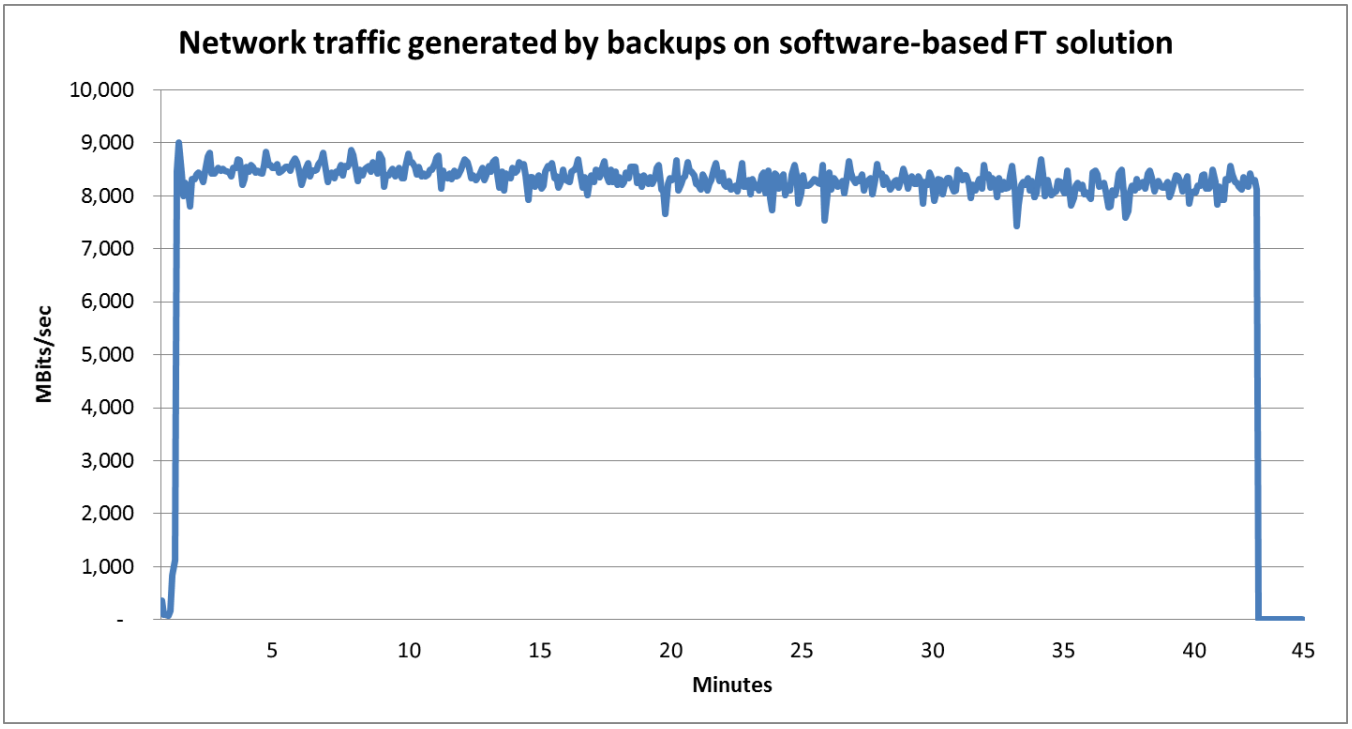

Figure 5: Network traffic in Mbits/sec between the two software-based FT hosts during an eight-VM OLTP workload.

FAULT TOLERANCE

To demonstrate the effectiveness of the hardware-based fault tolerance in the NEC solution, we simulated a system failure by removing one of the redundant servers. Before removing the server we started an eight-VM 45-minute OLTP workload run. We pulled the server 30 minutes into the run. As Figure 6 shows, recovery from the failure was instantaneous; database performance showed no interruption or decrease whatsoever, not even momentarily.

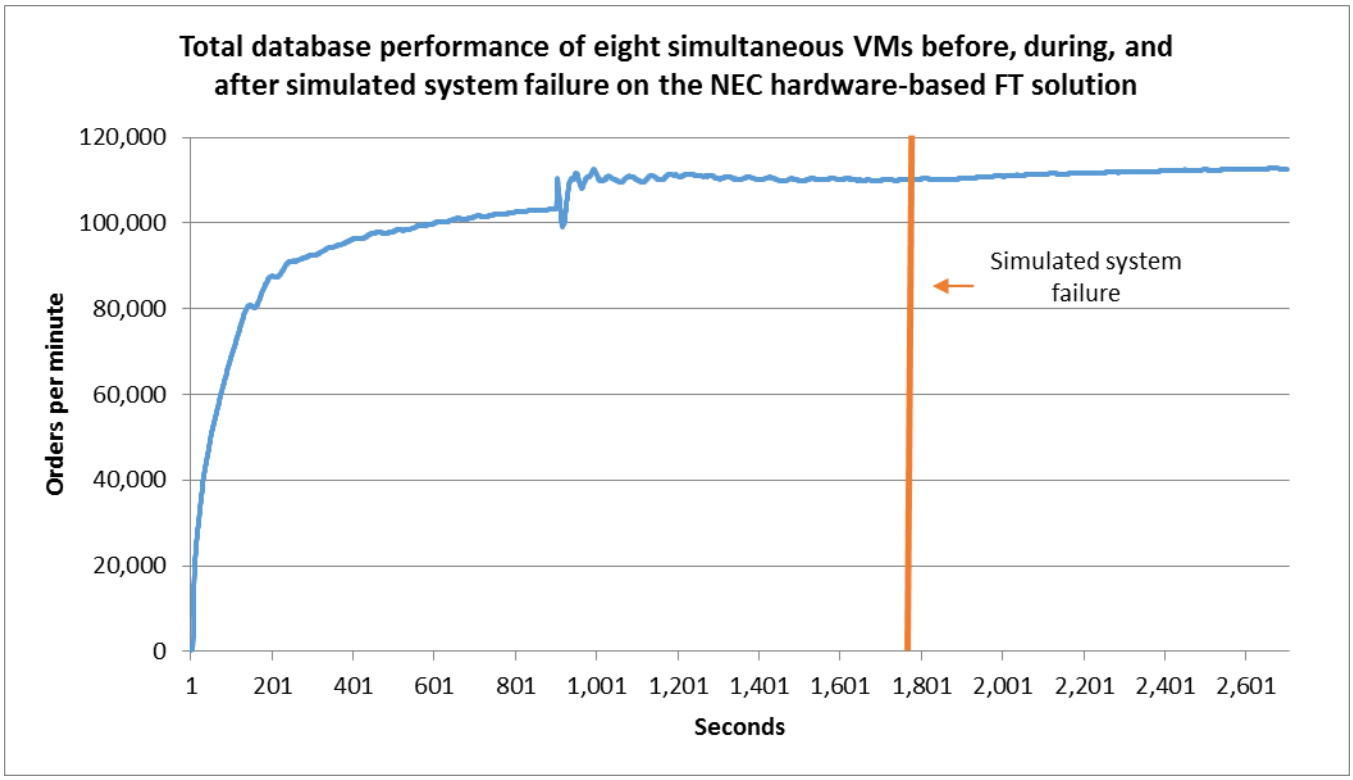

Figure 6: Performance of the eight simultaneous VMs remained constant even when we simulated a system failure on the NEC hardware-based FT solution.

CONCLUSION

Being able to rely on your server solution to deliver uncompromising levels of performance across a large number of VMs *and* to maintain these levels during an outage is a very appealing prospect. NEC Express5800/R320d-M4 servers with hardware-based fault tolerance can let you do just this.

In our datacenter, the hardware-based NEC solution with eight VMs achieved more than 2.4 times the performance of the software-based solution using VMware vSphere and recovered from a service interruption without downtime or performance loss. In addition, the hardware-based NEC solution did not require a dedicated 10-Gigabit network infrastructure to provide fault tolerance to the VMs. These advantages make the NEC Express5800/R320d-M4 server an excellent option for those businesses that don't want to choose between strong performance and fault tolerance.

APPENDIX A – SYSTEM CONFIGURATION INFORMATION

Figures 7 and 8 provide detailed configuration information for the test systems and for the NEC Storage M100 storage array, respectively.

System	NEC Express5800/R120e-1M	NEC Express5800/R320d-M4
Power supplies		
Total number	1	1
Vendor and model number	Delta Electronics® DPS-800QB A	Delta Electronics DPS-800QB A
Wattage of each (W)	800	800
Cooling fans		
Total number	8	4
Vendor and model number	Sanyo® Denki® 9CRN0412P5J003	San Ace® 9G0812P1K121
Dimensions (h × w) of each	1.5" × 2.25"	3" × 3"
Volts	12	12
Amps	1.0	1.8
General		
Number of processor packages	2	2
Number of cores per processor	10	10
Number of hardware threads per core	20	20
System power management policy	Balanced	Balanced
CPU		
Vendor	Intel	Intel
Name	Xeon	Xeon
Model number	E5-2670v2	E5-2670v2
Socket type	FCLGA2011	FCLGA2011
Core frequency (GHz)	2.50	2.50
Bus frequency	8 GT/s	8 GT/s
L1 cache	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)	256 KB (per core)
L3 cache	25 MB	25 MB
Platform		
Vendor and model number	NEC Express5800/R120e-1M	NEC Express 5800/R320d-M4
Motherboard model number	Micro-Star MS-S0821	DG7LGE
BIOS name and version	4.6.4012	7.0:25
BIOS settings	Default	Default
Memory module(s)		
Total RAM in system (GB)	192	192
Vendor and model number	Samsung® M393B2G70QH0-YK0	Samsung M393B2G70QH0-YK0
Type	PC3L-12800R	PC3L-12800R
Speed (MHz)	1,600	1,600
Speed running in the system (MHz)	1,600	1,600
Timing/Latency (tCL-tRCD-tRP-tRASmin)	11-11-11-35	11-11-11-35
Size (GB)	16	16

System	NEC Express5800/R120e-1M	NEC Express5800/R320d-M4
Number of RAM module(s)	12	12
Chip organization	Double-sided	Double-sided
Rank	2Rx4	2Rx4
Hypervisor #1		
Name	VMware ESXi 6.0.0	VMware ESXi 5.5 Update 2
Build number	2809209	1746018
File system	Ext4	Ext4
Language	English	English
RAID controller		
Vendor and model number	Emulex® Light Pulse LPe12002-M8-N	LSI® SAS2008
Firmware version	UB2.02a2	V7.23.01.00
Cache size (MB)	N/A	N/A
Hard drives #1		
Vendor and model number	Seagate® ST9500620NS	Seagate ST300MP0005
Number of drives	1	8
Size (GB)	500	300
RPM	7,200	15,000
Type	SATA 6.0 Gb/s	SAS
Ethernet adapter #1		
Vendor and model number	Broadcom® BCM5718B0KFBG	Intel 82576
Type	Integrated	Integrated
Driver	Tg3	ftSys_igb
Ethernet adapter #2		
Vendor and model number	NEC N8190-154	NEC D4G7LDR Gigabit X540-AT2
Type	PCIe®	PCIe
Driver	Bnx2x	ftSys_ixgbe

Figure 7: System configuration information for the test systems.

Storage array	SSD storage array
Number of storage controllers per array	1
RAID level	5
Number of drives per array	24
Drive vendor and model number	Micron MTFDDAK240MAV-1AE12ABYY
Drive size (GB)	240
Drive type	SSD, 6 Gbps SAS

Figure 8: Detailed configuration information for the SSD storage array.

APPENDIX B – HOW WE TESTED

Implementing fault tolerance using the two solutions

Figure 9 presents the steps we performed to implement fault tolerance using the two solutions.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
Preparing the system for fault tolerance	
Performing pre-install tasks on the NEC Express5800/R320d-M4 1. Pull all drives from the NEC Express5800/R320d-M4 storage except the first drive on node 0. 2. Disconnect all network cables, and make sure there's nothing except power connected to either server. 3. Press F2 to enter the BIOS of the server during POST. 4. In the BIOS, change the following settings: • Advanced → PCI Configuration → SAS Option ROM Scan: Disabled • Advanced → PCI Configuration → LAN1-4 Option ROM Scan: Disabled • Advanced → PCI Configuration → PCI Slot 1-4 Option ROM: Disabled • Server → OS Boot Monitoring: Disabled 5. Navigate to Save & Exit, and select Save Changes and Exit. 6. When asked to confirm your changes, select Yes.	Note: This step is not necessary on the NEC Express5800/120e-1M.
Installing ESXi 1. Insert NEC's build of ESXi 5.5 Update 2 into the server, and boot into the ESXi installation. 2. At the Welcome screen, press Enter to start. 3. At the confirmation screen, press F11 to begin the ESXi installation. 4. At Select a Disk to Install or Upgrade, select your installation drive, and press Enter. 5. At Please select a keyboard layout, select your language, and press Enter. 6. At Enter a root password, enter your password, and press Enter. 7. At Confirm Install, press F11 to start the installation. 8. When the installation is complete, plug a network cable into the first 1Gb slot on each server, and press Enter to reboot. 9. When the machine reboots, press F2 to log in. 10. Enter your username and password, and press Enter. 11. Navigate to Configure Management Network, and press	**Installing ESXi** 1. Insert the installation disk into the first server, and boot into the ESXi installation. 2. At the Welcome screen, press Enter to start. 3. At the confirmation screen, press F11 to begin the ESXi installation. 4. At Select a Disk to Install or Upgrade, select your installation drive, and press Enter. 5. At Please select a keyboard layout, select your language, and press Enter. 6. At Enter a root password, enter your password, and press Enter. 7. At Confirm Install, press F11 to start the installation. 8. When the installation is complete, press Enter to reboot. 9. When the machine reboots, press F2 to log in. 10. Enter your username and password, and press Enter. 11. Navigate to Configure Management Network, and press Enter. 12. Select Network Adapters, and press Enter.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
Enter. 12. Select Network Adapters, and press Enter. 13. In Network Adapters, select the NIC you wish to use, and press Enter. 14. Navigate to IPv4 Configuration, and press Enter. 15. In IPv4 Configuration, enter your IPv4 address and subnet mask, and press Enter. 16. When prompted to restart your management network, press Y to restart your management network. 17. Return to the main menu, the select Troubleshooting Options, and press Enter. 18. Enable SSH, and press Escape.	13. In Network Adapters, select the NIC you wish to use, and press Enter. 14. Navigate to IPv4 Configuration, and press Enter. 15. In IPv4 Configuration, put in your IPv4 address and subnet mask, and press Enter. 16. When prompted to restart your management network, press Y to restart your management network. 17. Insert the installation disk into the second server, and boot into the ESXi installation. 18. At the Welcome screen, press Enter to start. 19. At the confirmation screen, press F11 to begin the ESXi installation. 20. At Select a Disk to Install or Upgrade, select your installation drive, and press Enter. 21. At Please select a keyboard layout, select your language, and press Enter. 22. At Enter a root password, enter your password, and press Enter. 23. At Confirm Install, press F11 to start the installation. 24. When the installation is complete, press Enter to reboot. 25. When the machine reboots, press F2 to log in. 26. Enter your username and password, and press Enter. 27. Navigate to Configure Management Network, and press Enter. 28. Select Network Adapters, and press Enter. 29. In Network Adapters, select the NIC you wish to use, and press Enter. 30. Navigate to IPv4 Configuration, and press Enter. 31. In IPv4 Configuration, enter your IPv4 address and subnet mask, and press Enter. 32. When prompted to restart your management network, press Y to restart your management network.
Configuring ESXi and installing the ftSys Management Appliance 1. With the vSphere client, log into your ESXi server. 2. Click the Configuration tab. 3. Click Security Profile. 4. In Security Profile, scroll to Firewall, and click Properties. 5. In Firewall Properties, scroll to syslog, check its checkbox, and click OK.	**Configuring a high-availability cluster** 1. Log into vCenter. 2. Right-click your datacenter and select New Cluster. 3. Name your cluster, check Turn On vSphere HA, and click Next. 4. In vSphere HA, leave settings on defaults, and click Next. 5. In Virtual Machine Options, leave settings on defaults, and click Next.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
6. Insert NEC's FT control software install DVD into the host running vSphere client. Then, in the vSphere client, select File → Deploy OVF template. 7. In Deploy OVF Template, navigate to the ftSysMgt appliance OVA (if you have mounted the DVD on your D: drive, it is located at D:\appliance\ftSysMgt-5.1.1-233_OVF10.ova), and click Next. 8. In OVF Template Details, click Next. 9. Accept the EULAs, and click Next. 10. In Name and Location, enter the name of your appliance, and click Next. 11. In Storage, select the local storage for the ESXi server, and click Next. 12. In Disk Format, select Thick Provision Lazy Zeroed, and click Next. 13. In Ready to Complete, check the Power on after deployment checkbox, and click Finish. 14. After the appliance has been deployed, right-click it, and select Open Console. 15. When the VM has finished booting, navigate to Configure Network, and press Enter. 16. In the network configuration Main Menu, type 6, and press Enter. 17. Type your IP address (it must be in the same subnet as the host), and press Enter. 18. Type your subnet, and press Enter. 19. Type 1, then press Enter to exit network configuration. 20. Navigate to Login, and press Enter. 21. Log in with username `root` and password `ftServer`. 22. Change the password to your desired password with the command `passwd root`. 23. Insert a hard drive into slot 0 on node 1. 24. Mount NEC's FT control software install DVD to the appliance. 25. After the DVD is mounted, run the following command: `/opt/ft/sbin/ft-install /dev/cdrom` 26. Enter the IP address of the ESXi host, and press Enter. 27. Enter the root username of the ESXi host, and press Enter.	6. In VM Monitoring, leave settings on defaults, and click Next. 7. In VMware EVC, check Enable EVC for Intel Hosts, select Intel "Ivy Bridge" Generation, and click Next. 8. In VM Swapfile Location, leave the settings on the default recommended option, and click Next. 9. In Ready to Complete, click Finish.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
28. Enter the root password of the ESXi host, and press Enter. 29. When asked to review your system documentation, press Y to continue. 30. If you get any more prompts, press Y to continue. 31. Finally, a prompt to reboot the host will appear. Press Y to reboot the host. 32. After several reboots (and roughly 90 minutes after the host has finished rebooting), node 0 will synchronize with node 1 and the host will become fault tolerant.	
Adding the host to the vCenter 1. Log into your vCenter. 2. Right-click the datacenter for your FT server, and select Add Host. 3. Enter the hostname or IP address, then the authentication credentials, and click Next. 4. If prompted, click Yes to accept the security credentials for your new host. 5. At Host Information, click Next. 6. At Assign License, apply the relevant license to your host, and click Next. 7. At Lockdown Mode, leave at defaults, and click Next. 8. At Ready to Complete, verify your settings, and click Finish.	**Adding the hosts to the vCenter and high-availability cluster** 1. Right-click the HA cluster, and select Add Host. 2. Enter the hostname or IP address of the first host, then the authentication credentials, and click Next. 3. If prompted, click Yes to accept the security credentials for your new host. 4. At Host Information, click Next. 5. At Assign License, apply the relevant license to your host, and click Next. 6. At Lockdown Mode, leave the defaults, and click Next. 7. At Ready to Complete, verify your settings, and click Finish. 8. Click your new host, and click the Manage tab. 9. Click the Networking tab, and click Edit settings on the host's VMkernel port. 10. In Port properties, check vMotion traffic and Fault Tolerance logging, and click OK. 11. Right-click the HA cluster, and select Add Host. 12. Enter the hostname or IP address of the second host, then the authentication credentials, and click Next. 13. If prompted, click Yes to accept the security credentials for your new host. 14. At Host Information, click Next. 15. At Assign License, apply the relevant license to your host, and click Next. 16. At Lockdown Mode, leave the defaults, and click Next. 17. At Ready to Complete, verify your settings, and click Finish. 18. Click your new host, and click the Manage tab. 19. Click the Networking tab, and click Edit settings on the host's VMkernel port.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
	20. In Port properties, check vMotion traffic and Fault Tolerance logging, and click OK.
Adding iSCSI storage to the server 1. Select your host in the vCenter. 2. Click Manage, then Networking. 3. Click Add Host Networking. 4. In Select connection type, select VMKernel Network Adapter, then click Next. 5. In Select target device, select New standard switch, then click Next. 6. In Create a Standard Switch, click Add adapters. 7. Select the network adapter you want and click OK. 8. Click Next. 9. In Port properties, label your VMkernel adapter, then click Next. 10. In IPv4 settings, select Use static IPv4 settings, type in your IP address and netmask, then click Next. 11. In Ready to complete, verify the settings are correct, then click Finish. 12. Click Manage, then Storage. 13. In Storage Adapters, click Add New Storage Adapter. 14. Select iSCSI software adapter, then click OK. 15. Select your iSCSI software adapter, then select Network Port Binding. 16. Select Add. 17. Choose the VMKernel port group you created previously for iSCSI traffic, then click OK. 18. Select Targets, then click Add. 19. In the Add Send Target Server window, type the IP address of your iSCSI Server, then click OK. 20. When prompted, rescan your host's storage information. It should detect your iSCSI storage and attach it.	**Adding iSCSI storage to the servers** 1. Select your first host in the vCenter. 2. Click Manage, then Networking. 3. Click Add Host Networking. 4. In Select connection type, select VMKernel Network Adapter, then click Next. 5. In Select target device, select New standard switch, then click Next. 6. In Create a Standard Switch, click Add adapters. 7. Select the network adapter you want and click OK. 8. Click Next. 9. In Port properties, label your VMkernel adapter, then click Next. 10. In IPv4 settings, select Use static IPv4 settings, type in your IP address and netmask, then click Next. 11. In Ready to complete, verify the settings are correct, then click Finish. 12. Click Manage, then Storage. 13. In Storage Adapters, click Add New Storage Adapter. 14. Select iSCSI software adapter, then click OK. 15. Select your iSCSI software adapter, then select Network Port Binding. 16. Select Add. 17. Choose the VMKernel port group you created previously for iSCSI traffic, then click OK. 18. Select Targets, then click Add. 19. In the Add Send Target Server window, type the IP address of your iSCSI Server, then click OK. 20. When prompted, rescan your host's storage information. It should detect your iSCSI storage and attach it. 21. Select your second host in the vCenter. 22. Click Manage, then Networking. 23. Click Add Host Networking. 24. In Select connection type, select VMKernel Network Adapter, then click Next. 25. In Select target device, select New standard switch, then click Next. 26. In Create a Standard Switch, click Add adapters.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
	27. Select the network adapter you want and click OK. 28. Click Next. 29. In Port properties, label your VMkernel adapter, then click Next. 30. In IPv4 settings, select Use static IPv4 settings, type in your IP address and netmask, then click Next. 31. In Ready to complete, verify the settings are correct, then click Finish. 32. Click Manage, then Storage. 33. In Storage Adapters, click Add New Storage Adapter. 34. Select iSCSI software adapter, then click OK. 35. Select your iSCSI software adapter, then select Network Port Binding. 36. Select Add. 37. Choose the VMKernel port group you created previously for iSCSI traffic, then click OK. 38. Select Targets, then click Add. 39. In the Add Send Target Server window, type the IP address of your iSCSI Server, then click OK. 40. When prompted, rescan your host's storage information. It should detect your iSCSI storage and attach it.
Preparing the VMs for fault tolerance	
Note: This step is not necessary on the NEC Express5800/R320d-M4 because once the system is prepared for fault tolerance, every VM is automatically fault tolerant.	**Configuring a VM to be fault tolerant** 1. Right-click the first VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 2. If a verification popup window appears, click Yes. 3. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 4. In Select host, select the second host in your cluster, and click Next. 5. In Ready to complete, verify the details of your VM, and click Finish. 6. Right-click the second VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 7. If a verification popup window appears, click Yes. 8. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
	split the VMs across multiple datastores), and click Next. 9. In Select host, select the second host in your cluster, and click Next. 10. In Ready to complete, verify the details of your VM, and click Finish. 11. Right-click the third VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 12. If a verification popup window appears, click Yes. 13. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 14. In Select host, select the second host in your cluster, and click Next. 15. In Ready to complete, verify the details of your VM, and click Finish. 16. Right-click the fourth VM you want to become FT and select Fault Tolerance → Turn On Fault Tolerance. 17. If a verification popup window appears, click Yes. 18. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 19. In Select host, select the second host in your cluster, and click Next. 20. In Ready to complete, verify the details of your VM, and click Finish. 21. Right-click the fifth VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 22. If a verification popup window appears, click Yes. 23. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 24. In Select host, select the second host in your cluster, and click Next. 25. In Ready to complete, verify the details of your VM, and click Finish. 26. Right-click the sixth VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance.

Hardware-based FT on the NEC Express5800/R320d-M4	Software-based FT on the NEC Express5800/120e-1M
	27. If a verification popup window appears, click Yes. 28. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 29. In Select host, select the second host in your cluster, and click Next. 30. In Ready to complete, verify the details of your VM, and click Finish. 31. Right-click the seventh VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 32. If a verification popup window appears, click Yes. 33. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 34. In Select host, select the second host in your cluster, and click Next. 35. In Ready to complete, verify the details of your VM, and click Finish. 36. Right-click the eighth VM you want to become FT, and select Fault Tolerance → Turn On Fault Tolerance. 37. If a verification popup window appears, click Yes. 38. In Select datastores, select the appropriate backup datastores for your secondary VM (we chose the same datastore as the VM, but normally it is recommended to split the VMs across multiple datastores), and click Next. 39. In Select host, select the second host in your cluster, and click Next. 40. In Ready to complete, verify the details of your VM, and click Finish.

Figure 9: The steps required to implement fault tolerance using the two solutions.

Conducting our performance testing

About our test tool, DVD Store Version 2.1

To create our real-world ecommerce workload, we used the DVD Store Version 2.1 benchmarking tool. DS2 models an online DVD store, where customers log in, search for movies, and make purchases. DS2 reports these actions in orders per minute that the system could handle, to show what kind of performance you could expect for your customers. The DS2 workload also performs other actions, such as adding new customers, to exercise the wide range of database functions you would need to run your ecommerce environment.

For more details about the DS2 tool, see www.delltechcenter.com/page/DVD+Store.

Installing SQL Server 2014 on SQL virtual machine

1. Power on the server.
2. Insert the SQL Server 2014 installation media into the DVD drive.
3. Click Run SETUP.EXE. If Autoplay does not begin the installation, navigate to the SQL Server 2014 DVD, and double-click it.
4. In the left pane, click Installation.
5. Click New SQL Server stand-alone installation or add features to an existing installation.
6. Select the Enter the product key radio button, and enter the product key. Click Next.
7. Click the checkbox to accept the license terms, and click Next.
8. Click Use Microsoft Update to check for updates, and click Next.
9. Click Install to install the setup support files.
10. If no failures are displayed, click Next.
11. At the Setup Role screen, choose SQL Server Feature Installation, and click Next.
12. At the Feature Selection screen, select Database Engine Services, Full-Text and Semantic Extractions for Search, Client Tools Connectivity, Client Tools Backwards Compatibility, Management Tools –Basic, and Management Tools – Complete. Click Next.
13. At the Installation Rules screen, after the check completes, click Next.
14. At the Instance configuration screen, leave the default selection of default instance, and click Next.
15. At the Server Configuration screen, choose NT Service\SQLSERVERAGENT for SQL Server Agent, and choose NT Service\MSSQLSERVER for SQL Server Database Engine. Change the Startup Type to Automatic. Click Next.
16. At the Database Engine Configuration screen, select the authentication method you prefer. For our testing purposes, we selected Mixed Mode.
17. Enter and confirm a password for the system administrator account.
18. Click Add Current user. This may take several seconds.
19. Click Next.
20. At the Error and usage reporting screen, click Next.
21. At the Installation Configuration Rules screen, check that there are no failures or relevant warnings, and click Next.
22. At the Ready to Install screen, click Install.
23. After installation completes, click Close.
24. Close the installation window.
25. Shutdown the virtual machine.

Configuring the database

We generated the data using the Install.pl script included with DVD Store version 2.1 (DS2), providing the parameters for our 10GB database size and the database platform we used. We ran the Install.pl script on a utility system running Linux, to generated the database schema.

After processing the data generation, we transferred the data files and schema creation files to a Windows-based system running SQL Server. We built the 10GB database in SQL Server, and then performed a full backup, storing the backup file remotely for quick access. We used that backup file to restore the database when necessary.

The only modification we made to the schema creation scripts were the specified file sizes for our database. We explicitly set the file sizes higher than necessary to ensure that no file-growth activity would affect the outputs of the test. Other than this file size modification, we created and loaded the database in accordance to the DVD Store documentation. Specifically, we followed these steps:

1. We generated the data, and created the database and file structure using database creation scripts in the DS2 download. We made size modifications specific to our 10GB database, and made the appropriate changes to drive letters.
2. We transferred the files from our Linux data generation system to a Windows system running SQL Server.
3. We created database tables, stored procedures, and objects using the provided DVD Store scripts.
4. We set the database recovery model to bulk-logged to prevent excess logging.
5. We loaded the data we generated into the database. For data loading, we used the import wizard in SQL Server Management Studio. Where necessary, we retained options from the original scripts, such as Enable Identity Insert.
6. We created indices, full-text catalogs, primary keys, and foreign keys using the database-creation scripts.
7. We updated statistics on each table according to database-creation scripts, which sample 18 percent of the table data.
8. On the SQL Server instance, we created a ds2user SQL Server login using the following Transact SQL (TSQL) script:

```
USE [master]
GO
CREATE LOGIN [ds2user] WITH PASSWORD=N'',
        DEFAULT_DATABASE=[master],
        DEFAULT_LANGUAGE=[us_english],
        CHECK_EXPIRATION=OFF,
        CHECK_POLICY=OFF
GO
```

9. We set the database recovery model back to full.
10. We created the necessary full text index using SQL Server Management Studio.
11. We created a database user, and mapped this user to the SQL Server login.
12. We then performed a full backup of the database. This backup allowed us to restore the databases to a pristine state.

Running the DVD Store tests

We created a series of batch files, SQL scripts, and shell scripts to automate the complete test cycle. DVD Store outputs an orders-per-minute metric, which is a running average calculated through the test. In this report, we report the last OPM reported by each client/target pair.

We used the following DVD Store parameters for testing:

```
ds2sqlserverdriver.exe --target=<target_IP> --ramp_rate=10 --run_time=30 --
n_threads=32 --db_size=10GB --think_time=0 --detailed_view=Y --
warmup_time=15 --csv_output=<drive path>
```

APPENDIX C – DETAILED PERFORMANCE TEST RESULTS

Figure 10 shows database performance results for the hardware-based FT solution and the software-based FT solution.

	1 simultaneous VM		2 simultaneous VMs		4 simultaneous VMs		8 simultaneous VMs	
	Hardware FT	Software FT	Hardware FT	Software FT	Hardware FT	Software FT	Hardware FT	Software FT
VM 1	24,786	19,226	27,162	16,787	27,185	13,699	26,048	10,489
VM 2			27,724	16,692	26,985	13,751	25,928	10,857
VM 3					26,736	14,045	25,551	10,912
VM 4					26,910	13,555	25,955	10,940
VM 5							25,799	9,525
VM 6							25,585	10,193
VM 7							25,641	10,019
VM 8							25,567	10,205
TOTAL	**24,786**	**19,226**	**54,886**	**33,479**	**107,816**	**55,050**	**206,074**	**83,140**
AVERAGE	**24,786**	**19,226**	**27,443**	**16,740**	**26,954**	**13,763**	**25,759**	**10,393**

Figure 10: Database orders per minute for the two FT solutions.

APPENDIX D – DISK LATENCY RESULTS

We used the same all-flash storage array for all testing, and to ensure that this array was not the bottleneck during our testing, we measured the guest latency for the VMs using esxtop. Figure 11 shows the read and write guest latency during our testing. As it shows, both hardware and software fault tolerance have latencies well below the recommended 20-millisecond latency during the test run, and both hardware and software averaged below 2ms latency in both read and write during the test, which indicates that the bottleneck is not the storage.

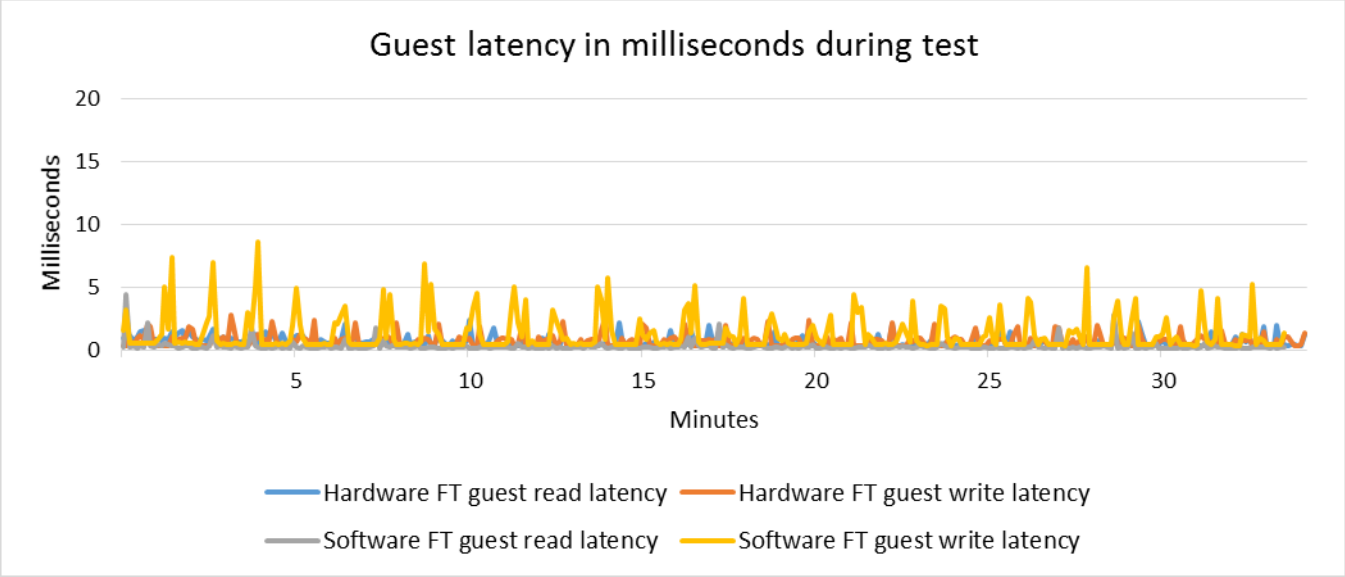

Figure 11: Average guest latency of the hardware and software fault tolerance during maximum load (eight simultaneous VMs).

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.