

Intel and VMware Deliver



Intel® Xeon® processor E7-4870 and VMware vSphere 5.0

Reliability: Simulated RAM error isolated to single VM

Reliability is a huge factor in choosing a server to host virtualized business-critical applications. If a RAM error brings your entire server to a halt, all the applications on it stop functioning. Principled Technologies loaded an Intel® Xeon® processor E7-4870 based server running VMware® vSphere™ 5.0 with 12 virtual machines (VMs). We then injected a memory error that affected one of the VMs.

In our hands-on tests, the results backed up the claims. By taking advantage of Machine Check Architecture (MCA) recovery, a feature of the Intel® Xeon® processor E7 family, the server keeps running. Only the single VM with the simulated memory error fails; the underlying hypervisor and the remaining 11 VMs continue unharmed. In addition, the hypervisor retires the memory space that contains the fault. VMware vCenter then could gracefully restart the VM, delivering timely recovery of service without user intervention. When we repeated the test on several other VMs, the results were the same: The system continued to operate.

The Intel Xeon processor E7-4870-based server running VMware vSphere 5.0 confined the problem to a single VM and let the others keep going.



RELIABILITY, AVAILABILITY, AND SERVICEABILITY

The reliability, availability, and serviceability (RAS) features of a server are crucial to the smooth operation of your business when hosting critical business applications and services.

Reliability means the system has features that help it avoid and detect errors, availability relates to the amount of downtime a system has due to the errors, and serviceability is how the solution handles failed (or potentially failed) components—either proactively or retroactively.

MCA recovery lets the hypervisor or underlying operating system recover from memory errors that cannot be corrected at the hardware level. MCA recovery technology detects an error and reports it to the hypervisor or OS, which takes the appropriate action—fixing the error when possible or quarantining the bad memory to prevent it from continuing to corrupt data.

Enabling MCA recovery, as we did when we tested the Intel Xeon processor E7-4870-based server running VMware vSphere 5.0, allowed the system to detect the error and shut down the affected VM while keeping the other VMs on the server up and running. VMware vCenter could then restart the affected VM without issue as the hypervisor had retired the memory.

WHAT HAPPENED WHEN WE INJECTED A FAULT?

The memory errors we introduced did not cause the server to go down.

Because memory errors can cause system crashes, we tested the RAS features of our Intel Xeon processor E7-4870-based server running VMware vSphere 5.0 by injecting a fault directly into the memory of one VM and watching what happened.

Due to the advanced RAS features of the Intel Xeon processor E7-4870 and VMware vSphere 5.0, the memory error we introduced did not cause the server to go down. Only the VM with the memory fault failed; the fault did not affect the hypervisor or the other 11 VMs.

To be sure that the server and VM behavior we observed was repeatable, we injected the memory fault 12 different times into 12 different VMs. Each time, only the VM with the memory fault shut down.

SUMMARY

When we injected a fault into the memory of one virtual machine on an Intel Xeon processor E7-4870-based server running VMware vSphere 5.0, the processor's MCA recovery feature allowed the hypervisor and all the other virtual machines to keep running. In addition, the hypervisor retired the memory space that contained the fault, which let VMware vCenter gracefully restart the VM.

APPENDIX A – SERVER CONFIGURATION INFORMATION

Figure 1 provides detailed configuration information about the test server.

System	Intel® Xeon® processor E7-4870-based server
Power supplies	
Total number	4
Vendor and model number	Delta DPS-850FB
Wattage of each (W)	850
Cooling fans	
Total number	8
Vendor and model number	Delta PFB0812DHE
Dimensions (h x w) of each	3.50" x 3.50"
Volts	12
Amps	3.3
General	
Number of processor packages	4
Number of cores per processor	10
Number of hardware threads per core	2
CPU	
Vendor	Intel
Name	Xeon
Model number	E7-4870
Stepping	F2
Socket type	Socket 1567
Core frequency (GHz)	2.40
Bus frequency (GT/s)	6.40
L1 cache	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)
L3 cache (MB)	30
Platform	
Vendor and model number	Quanta QCI QSSC-S4R
Motherboard model number	31S4RMB00A0
Motherboard chipset	Intel® E7500 chipset
BIOS name and version	Intel® QSSC-S4R.QCI.01.00.T008.122020101559 (12/20/2010)
BIOS settings	Intel® Virtualization Technology Enabled, Patrol Scrub Interval set to 6
Memory module(s)	
Total RAM in system (GB)	1,024
Vendor and model number	Samsung M393B2K70CM0-CF8
Type	DDR3 PC3-8500
Speed (MHz)	1,066
Speed running in the system (MHz)	1,066
Timing/Latency (tCL-tRCD-tRP-tRASmin)	7-7-7-20
Size (GB)	16
Number of RAM module(s)	64

System	Intel® Xeon® processor E7-4870-based server
Chip organization	Double-sided
Rank	Quad
Hard disk	
Vendor and model number	Seagate ST9146803SS
Number of disks in system	2
Size (GB)	146
Buffer size (MB)	16
RPM	10,000
Type	SAS
Operating system	
Name	VMware vSphere 5.0 (ESXi 5.0 Build 384862)
File system	EXT3
Kernel	5.0.0
Build	384862
Language	English
Ethernet	
Vendor and model number	Intel® 82576NS Gigabit Ethernet Adapter
Type	Integrated
Optical drive(s)	
Vendor and model number	Sony Optiarc AD-7580S
Type	DVD-RW
USB ports	
Number	5
Type	USB 2.0

Figure 1: Server configuration information.

APPENDIX B – DETAILED TEST METHODOLOGY

In this appendix, we present the steps we followed as we conducted our testing.

Setting up and configuring the EMC storage

We used an EMC® CLARiiON® CX3-40 Fibre Channel-connected SAN to store all VMs for this testing. We used a QLogic® QLE2462 dual-port host bus adapters (HBA) in our server to connect to the external storage.

Installing VMware vSphere 5.0

We installed VMware vSphere 5.0 (Build 384862) using the steps below. For testing, we kept a clean version of VMware vSphere 5.0 by booting and running from CD after each reboot.

1. Turn on the server, and insert the vSphere CD.
2. Once the server boots from the CD, the setup process begins automatically. This should take about 45 minutes.
3. Once the setup completes, VMware vSphere boots and is ready for use.

Installing vSphere client on a management workstation

1. Configure your management workstation to be on the same subnet as your vSphere server. In Internet Explorer, navigate to the IP address of your vSphere server.
2. Accept certificate warnings by clicking Continue to this website (not recommended).
3. Click Download vSphere Client.
4. Choose Run to Install.
5. Choose your language, and click OK.
6. At the Welcome screen, click Next.
7. Read the End-User Patent Agreement, and click Next.
8. Accept the license terms, and click Next.
9. Enter user and organization details, and click Next.
10. At the Destination folder screen, click Change if you wish to change your install destination; otherwise click Next.
11. Click Install.
12. To complete the installation, click Finish.

Creating a virtual machine with vSphere

1. Click Start→All Programs→VMware→VMware vSphere client.
2. Enter the IP address or hostname, user name, and password, and click Login.
3. Right-click on the host, and choose New Virtual Machine.
4. Choose Custom, and click Next.
5. Assign a name to the virtual machine, and click Next.
6. Select a datastore for the virtual machine files, and click Next.
7. Choose Virtual Machine Version: 8, and click Next.
8. Choose Microsoft Windows, then Microsoft Windows Server 2008 R2 (64-bit), and click Next.
9. Select the number of virtual processors and cores per virtual processor you want to allocate (in our case, six virtual processors and one core per virtual processor), and click Next.
10. Select the amount of RAM you want to allocate (in our case, 82 GB), and click Next.
11. Select 1 as the number of NICs, select your network, and select your desired network adapter. Click Next.
12. Choose LSI Logic SAS as the SCSI controller, and click Next.
13. Select Create a new virtual disk, and click Next.
14. Choose disk size, store location, select Thin Provision and click Next.
15. Select the appropriate virtual device node, and click Next.
16. Click Finish.

Installing Windows Server 2008 R2 Datacenter Edition on the hosts

1. Insert the installation DVD for Windows Server 2008 R2 x64 into the drive.
2. In vSphere, right-click the virtual machine, and choose Edit Settings.
3. Click CD/DVD drive 1, click the Host Device radio button, and choose the appropriate drive.
4. Click OK.
5. Right-click the machine, and choose Power→Power On.
6. Right-click the machine, and choose Open console.
7. At the Language Selection Screen, click Next.
8. Click Install Now.
9. Select Windows Server 2008 R2 Datacenter (Full Installation), and click Next.
10. Click the I accept the license terms check box, and click Next.
11. Click Custom.
12. Click Drive options (advanced).
13. Ensure you select the proper drive, and click New.
14. Click Apply.
15. Click Next.
16. At the User's password must be changed before logging on warning screen, click OK.
17. Set the Administrator password, and click the arrow to continue.
18. At the Your password has been changed screen, click OK.

Installing VMware tools on the Virtual machine

1. Right-click the virtual machine, select Guest→Install/Upgrade VMware tools.
2. Click OK to close the VMware Tools warning.
3. Right-click the VM, and choose Open Console.
4. Double-click Run setup.exe.
5. At the Welcome screen, click Next.
6. Select Typical Installation, and click Next.
7. Click Install.
8. After the installation completes, click Finish, and restart the virtual machine.

Installing system updates in Windows Server 2008 R2

We installed all the following updates available on April 13, 2011 on each virtual machine:

- Security Update for Windows Server 2008 R2 x64 Edition (KB2503658)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2506223)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2425227)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2511455)
- Update for Windows Server 2008 R2 x64 Edition (KB2505438)
- Update for Windows Server 2008 R2 x64 Edition (KB2511250)
- Update for Windows Server 2008 R2 x64 Edition (KB2506014)
- Cumulative Security Update for Internet Explorer 8 for Windows Server 2008 R2 x64 Edition (KB2497640)
- Windows Internet Explorer 9 for Windows Server 2008 R2 for x64-based Systems
- Cumulative Security Update for ActiveX Killbits for Windows Server 2008 R2 x64 Edition (KB2508272)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2507618)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2510531)
- Update for Windows Server 2008 R2 x64 Edition (KB2484033)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2509553)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2506212)

- Security Update for .NET Framework 3.5.1 on Windows 7 and Windows Server 2008 R2 SP1 for x64-based Systems (KB2446710)
- Update for Windows Server 2008 R2 x64 Edition (KB2524375)
- Security Update for Windows Server 2008 R2 x64 Edition (KB2508429)

Executing the test

This shows the steps we used for testing. For testing, we used a boot cd containing VMware ESXi. Because it is a boot CD, we had to add our VMs to the inventory after each reboot.

1. Power on the server.
2. Add the server to the vCenter datacenter.
3. Add VMs to inventory.
4. Start VMs 1-12 in order (with a 20-second delay between each start).
5. Let the VMs sit idle for 10 minutes.
6. Check each VM's Heartbeat status (all should report a status of "green").
7. Load the Intel machine check exception (MCE) injector driver.
8. Start the MCE injector Python script:
 - a. Check the number of VMs running.
 - b. Do the following for each VM serialized:
 - i. Select a random virtual page number (VPN) for the current VM.
 - ii. Lookup the machine page number (MPN) associated with that random VPN.
 - iii. Check MPN "Selected" and "Retired" status; both should report a status of '0'.
 - iv. Start a new MCE Injector instance.
 - v. Using the select MPN, the script prepares, injects, and triggers the MCE.
 - vi. Record the time it takes for the VM to shut down.
 - vii. Record the time for the MPN "Selected" and "Retired" status to change from '0' (false) to '1' (true).
 - viii. Wait 60 seconds
 - ix. Check each VM's Heartbeat status. All VMs which have had an error injected should report a status of "gray" since they have been shut down. All other VMs should report a status of "green" to indicate they are still responding.
 - x. Cleanup MCE Injector instance.
 - xi. Repeat these steps for the next VM until all VMs have successfully simulated a memory error and shutdown down accordingly.
 - c. Do a final heartbeat check on all VMs (all should report a status of "gray").
9. Verify that vCenter properly reported the memory errors.
10. Reboot the server.

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.

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.