

## Dell™ PowerEdge™ M420 blade server solution with Dell EqualLogic™ storage delivered database clustering performance



*for Microsoft® SQL Server® 2012 AlwaysOn Failover Cluster Instances*

Data center space always comes at a premium. For that reason, it is important to consider not only performance and manageability, but also the amount of physical space your servers use when you are selecting a blade server to run your database applications. The ultra-dense, quarter-height Dell PowerEdge M420 blade server lets you fit twice as many servers in a single chassis than if you were using conventional half-height blades, without sacrificing performance or other features. This essentially means you can double your data center's compute capability using your existing Dell blade chassis and network infrastructure.

We tested Dell PowerEdge M420 blade servers in our lab with a Dell EqualLogic PS6110XS array and two Dell Force10 MXL 10/40GbE switches, and found that they delivered solid performance in a dense, small package with over 400,000 orders per minute in a database application.

Using less space lets you reduce your infrastructure cost for power, cooling, racks, and floor space. These factors make the Dell PowerEdge M420 blade server an excellent choice for any company's data center.

## SPACE EFFICIENCY = SAVINGS

The power of the new Dell PowerEdge M420 delivers excellent performance in half the space of conventional half-height blades. Using existing chassis, switch, rack, and cabling infrastructure, organizations can immediately double the density of their data center by opting for the quarter-height PowerEdge M420 form factor. The ability to avoid additional chassis, switch, and cabling costs can contribute to capital savings as your data center expands and plans for the future.

## HOW WE TESTED

To measure the performance of the PowerEdge M420 blade server, we prepared two servers to run a Microsoft Windows Server 2008 R2 failover cluster. We then installed Microsoft SQL Server clustering and configured four named instances of Microsoft SQL Server 2012, each containing its own 40GB database. We assigned two instances of SQL Server to the first node of the cluster and two instances of SQL Server to the second node of the cluster.

Each node had access to shared storage, specifically a Dell EqualLogic PS6110XS 10Gb hybrid SAS/SSD array. The PS6110XS is an enterprise-class iSCSI SAN solution that provides the power and speed of SSD storage alongside the capacity advantages of SAS storage. The Dell EqualLogic PS6110XS handles tiering of data automatically between the two storage technologies, saving administrators time and effort. We connected the Dell PowerEdge M420 blade servers to the Dell EqualLogic PS6110XS via their 10Gb connections for high-performance throughput. For more on how we configured the storage, switches, and topology of our test, see [Appendix B](#).

After the server and storage were configured, we used an ecommerce workload called DVD Store to test the clustered blade server solution. To fully exercise the server in our lab setting, we enabled a longer recovery interval in our SQL Server instances. Figure 1 presents a concise configuration summary for the solution.

Solution	Dell PowerEdge M420 solution
Servers	2 x Dell PowerEdge M420
Operating system	Microsoft Windows Server 2008 R2 SP1
Database	Microsoft SQL Server 2012
Switches	2 x Dell Force10 MXL 10/40GbE
Storage	1 x Dell EqualLogic PS6110XS

Figure 1: General configuration summary for the solution we tested.

For detailed configuration information about the server, see [Appendix A](#). For step-by-step instructions on how we tested, see [Appendix B](#).

## ABOUT THE TEST ENVIRONMENT

### Features of the Dell PowerEdge M420 blade server

The new Intel® Xeon® processor E5-2470-powered Dell PowerEdge M420 blade server sets a new standard for computing density in a blade environment, and uses your existing Dell PowerEdge M1000e blade enclosure. Features include:

- **Double the computing power per rack unit.** Quarter-height blades capable of handling eight-core Intel Xeon processor E5-2400-series units in a two-socket configuration allow you to use your rack space to its maximum potential; 32 blades per enclosure.
- **Cost-efficient upgrade.** The PowerEdge M420 utilizes the same M1000e chassis as previous Dell blades, taking advantage of the existing power, cooling, and network infrastructure already in your data center.
- **Management.** Dell's OpenManage™ suite of tools, combined with the Integrated Dell Remote Access Controller (iDRAC7), give you maximum flexibility in administering your computing infrastructure.
- **Connectivity.** Snap-in FlexIO technology provides the backbone for a flexible, high-speed interconnect fabric. With options including 1Gb and 10Gb Ethernet, Fibre Channel, InfiniBand, and the new Dell Force10 MXL 10/40GbE blade switches, connectivity options are flexible.
- **Power efficiency.** Dell PowerEdge M420 servers implement leading industry power-efficiency standards, and are designed to optimize performance while maintaining low power consumption, saving energy and money.
- **Virtualization-ready.** The M420 contains two SD card slots for hypervisor, redundant hypervisor, and/or vFlash media support, and supports various hypervisor software.

### Features of the Dell Force10 MXL 10/40GbE blade switch

The Dell Force10 MXL 10/40GbE blade switch brings new connectivity and performance options to blade computing environments with modular configuration options and standards-based 10/40GbE support. The Force10 MXL features:

- **Performance.** 10/40GbE connectivity minimizes latency and maximizes network throughput for data-intensive applications.
- **Flexibility.** Choose 1, 10, or 40GbE ports; FlexIO for plug-in QSFP+, SFP+; or 10GbE copper.
- **Extend and simplify.** The MXL is modular, and supports stacking for ease of management and higher performance. Support for converged switching allows iSCSI, NAS, Ethernet, and Fibre Channel traffic to all use the same

physical hardware, reducing additional management overhead and hardware costs.

## Features of the Dell EqualLogic PS6110XS 10GbE iSCSI array

The Dell EqualLogic PS6110XS 10GbE iSCSI array combines the low-latency and high-IOPs strength of solid-state storage with the raw capacity of traditional hard-disk drives in a single-chassis solution. Features include:

- **Agility.** Use storage that fits the demands of your server or application. Dell EqualLogic dynamically tiers data on SSD or HDD, depending on the performance requirements for a specific situation.
- **Upgradable.** Start out using your existing 10GbE SFP+ modules, and build out to cost-effective 10GBASE-T as your budget allows. Move volumes between storage pools without downtime, and change hardware without disruption.
- **Robust.** Vertical port sharing keeps your bandwidth at full speed, even if a network port fails. Automatic load-balancing keeps data moving reliably to and from your physical or virtual server farm.
- **Management.** EqualLogic Host Software, Host Integration Tools, and EqualLogic SAN Headquarters (SANHQ) provide the tools administrators need to keep a high-level view of the storage system, while providing the flexibility to make granular changes as needed.

## ABOUT DVD STORE VERSION 2.1

To create our real-world ecommerce workload, we used the DVD Store Version 2.1 (DS2) 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 (OPM) that the system could handle, to demonstrate what kind of performance you could expect for your customers. The DS2 workload also performs other actions, such as adding new customers, to utilize the wide range of database functions you would need to run your ecommerce environment. For more information about the DS2 tool, see <http://www.delltechcenter.com/page/DVD+Store>.

## WHAT WE FOUND

The two-node clustered SQL Server-powered solution running on the PowerEdge M420 blade servers achieved a total score of 436,661 orders per minute. Figure 2 shows the DS2 orders-per-minute results and CPU utilization for the two-node cluster, by SQL Server instance and node.

Dell PowerEdge M420 failover cluster DVD Store performance	
Orders per minute – Node 1, instance 1	109,662
Orders per minute – Node 1, instance 2	109,512
Orders per minute – Node 2, instance 1	108,953
Orders per minute – Node 2, instance 2	108,534
Total orders per minute for the SQL Server 2012 cluster	436,661
Percentage CPU utilization – Node 1	74.7%
Percentage CPU utilization – Node 2	77.1%

**Figure 2: Individual SQL Server instance orders per minute results, total orders per minute results, and CPU utilization for the two-node cluster running DVD Store.**

## APPENDIX A — DETAILED CONFIGURATION

Figure 3 provides detailed configuration information for the test servers.

System	Dell PowerEdge M420 blade server
<b>Power supplies (located in the Dell PowerEdge M1000e Blade Enclosure)</b>	
Total number	6
Vendor and model number	Dell A236P-00
Wattage of each (W)	2,360
<b>Cooling fans (located in the Dell PowerEdge M1000e Blade Enclosure)</b>	
Total number	9
Vendor and model number	Dell YK776 Rev. X50
Dimensions (h x w) of each	3.1" x 3.5"
Volts	12
Amps	7
<b>General</b>	
Number of processor packages	2
Number of cores per processor	8
Number of hardware threads per core	2
System power management policy	Balanced
<b>CPU</b>	
Vendor	Intel
Name	Xeon
Model number	E5-2470
Stepping	C2
Socket type	FCLGA1356
Core frequency (GHz)	2.3
Bus frequency	8 GT/s
L1 cache	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)
L3 cache	20 MB
<b>Platform</b>	
Vendor and model number	Dell PowerEdge M420
Motherboard model number	OMN3VC
BIOS version	1.2.4
BIOS settings	Default, DAPC power profile
<b>Memory module(s)</b>	
Total RAM in system (GB)	96
Vendor and model number	Samsung M393B2G70BH0-YH9
Type	PC3L-10600R
Speed (MHz)	1,333
Speed running in the system (MHz)	1,333
Timing/Latency (tCL-tRCD-tRP-tRASmin)	9-9-9-36
Size (GB)	16
Number of RAM module(s)	6

<b>System</b>	<b>Dell PowerEdge M420 blade server</b>
Chip organization	Double-sided
Rank	Dual
<b>Operating system</b>	
Name	Windows Server 2008 R2 SP1
Build number	7601
File system	NTFS
Kernel	ACPI x64-based PC
Language	English
<b>Graphics</b>	
Vendor and model number	Matrox G200eR
Graphics memory (MB)	16
Driver	2.4.1.0 (6/25/2012)
<b>RAID controller</b>	
Vendor and model number	Dell PERC H310 Embedded
Firmware version	20.10.1-0084
Driver version	5.1.112.64 (6/12/2011)
Cache size (MB)	0 MB
<b>Hard drive</b>	
Vendor and model number	Dell M16CSD1-50UCV-D
Number of disks in system	2
Size (GB)	50
Buffer size (MB)	N/A
RPM	N/A
Type	SSD
<b>Ethernet adapters</b>	
Vendor and model number	4x Broadcom BCM57810 NetXtreme II 10 GigE
Type	Mezzanine
Driver	7.2.8.0 (3/13/2012)
<b>USB ports</b>	
Number	2 External
Type	2.0

Figure 3: Server configuration information for the test server.

Figure 4 provides detailed information for the test storage.

Storage array	Dell EqualLogic PS-6110XS
Number of storage arrays	1
Number of storage controllers per array	2
RAID level	6 (accelerated)
Firmware version	6.0.0
Number of drives, type 1	7
Model number	LB400M
Drive size (GB)	400GB
Drive buffer size (MB)	N/A
Drive RPM	N/A
Drive type	SSD
Number of drives, type 2	17
Model number	ST9600204SS
Drive size (GB)	600GB
Drive buffer size (MB)	16MB
Drive RPM	10K
Drive type	6Gb SAS 2.5"

Figure 4: Detailed configuration information for the Dell EqualLogic PS-6110XS storage array.



## APPENDIX B – HOW WE TESTED

### Storage configuration overview

Our complete storage infrastructure for the PowerEdge M420 consisted of the Dell EqualLogic PS6110XS drive array and two Force10 MXL 10/40GbE blade switches. We configured the Dell EqualLogic PS6110XS drive array in a single member group, with the RAID policy being set at RAID 6 (accelerated). The Dell EqualLogic PS6110XS array contained 7 SSD disks and 17 SAS disks. We cabled the Dell EqualLogic array to the Force10 MXL switches, which in turn were connected to the PowerEdge M1000e midplane. Fabric A was used for iSCSI SAN traffic. Fabrics B and C were used for client application network traffic.

Other topology choices are available, such as using a converged networking approach via just two Force 10 MXL switches, passing client and storage traffic over the same switch hardware. Figure 5 shows our test bed configuration.

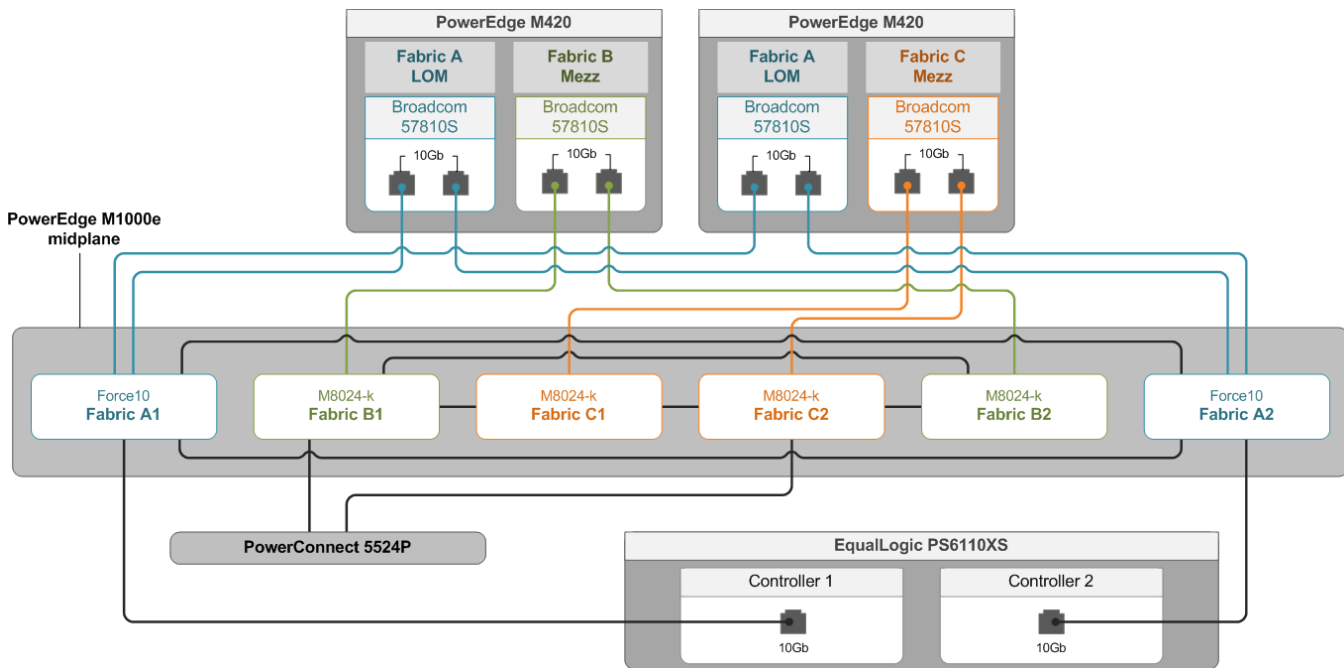


Figure 5: Test bed configuration.

Using the default storage pool, we created four LUNs per SQL Server instance: one for the SQL Server instance data, one for the SQL Server instance logs, one for the SQL Server instance tempdb files, and one for the SQL Server instance backup files. Additionally, we created LUNs presented to each node for the MSDTC and Quorum cluster resources.

For the internal storage, we configured a RAID 1 mirrored pair of disks for the operating system.

### Configuring the external storage

1. Using the command-line console, via serial cable, reset the Dell EqualLogic PS6110XS by using the `reset` command.
1. Supply a group name, group IP address, and IP address for eth0.

2. After group creation, using a computer connected to the same subnet as the storage, use the Dell EqualLogic Web interface to do the following:
  - a. Click the array, and choose Yes when prompted to configure the member. Choose RAID 6 (accelerated).
  - b. Create LUNs as specified above.
  - c. Enable shared access to the iSCSI target from multiple initiators on the volume.
  - d. Create an access control record for the volume, specifying access via the iSCSI initiator name for each server.

## Configuring the server operating system – Windows Server 2008 R2

### Adjusting BIOS settings

We used the latest released BIOS updates on the systems, and used the default BIOS settings with Dell Performance Per Watt (DAPC).

### Installing Windows Server 2008 R2

1. Mount the install media for Windows Server 2008 R2 to the server via the iDRAC virtual media wizard.
2. Power on the server.
3. Choose the language, time and currency, and keyboard input. Click Next.
4. Click Install Now.
5. Choose Windows Server 2008 R2 Enterprise (Full Installation), and click Next.
6. Accept the license terms, and click Next.
7. Click Custom.
8. Click the Disk, and click Drive options (advanced).
9. Click New→Apply→Format, and click Next.
10. After the installation completes, click OK to set the Administrator password.
11. Enter the administrator password twice, and click OK.
12. Connect the machine to the Internet and install all available Windows updates. Restart as necessary.
13. Enable remote desktop access.
14. Change the hostname and reboot when prompted.
15. Set up advanced networking (iSCSI offload):
  - a. Open the Broadcom Advanced Control Suite 4 application.
  - b. Select the Broadcom 10Gb NIC port 1 from the list.
  - c. In the Configuration tab, enable iSCSI Offload Engine from the Resource Reservation sub-section.
  - d. Select the newly created VBD interface under the same port name.
  - e. In the configuration tab, set the MTU size to 9000, and assign a static IP address. Click Apply.
  - f. Repeat steps b-e for Broadcom 10Gb NIC port 2 and apply the appropriate IP address.
16. We teamed two NICs for SQL Server application traffic, and two NICs for cluster heartbeat traffic (shared with iSCSI offload).

Set up advanced networking (NIC Teaming):

- a. Open the Broadcom Advanced Control Suite 4 application.
- b. From the Team Menu, select Create a Team.
- c. At the Broadcom Teaming Wizard welcome screen, click Next.
- d. Enter a name for the team (ex: Team 1), and click Next.
- e. Select the Team Type Smart Load Balancing(TM) and Failover (SLB), and click Next.
- f. Select both Broadcom 10Gb adaptors from the list, and click Add. Click Next.
- g. Select Do not configure a standby member, and click Next.
- h. From the Configure LiveLink list, select No, and click Next.
- i. Select Skip Manage VLAN, and click Next.
- j. Select Commit changes to system and Exit the wizard, and click Finish.
- k. Repeat for the second NIC team.

## Connecting to the volumes with Microsoft iSCSI Initiator

1. In the Dell EqualLogic console, ensure the volumes are online.
2. Click Start→Administrative Tools→iSCSI Initiator.
3. Select the Discovery Tab, and click Discover Portal.
4. Enter the IP address for the Storage Group, and click OK.
5. Select the Targets tab, and click Refresh.
6. Select the first Inactive Target listed, and click Connect.
7. Ensure that Add this connection to the list of Favorite Targets is selected, check the Enable multi-path check box, and click OK.
8. Repeat until you have connected to all volumes, and click OK.
9. Repeat steps on the remaining node in your cluster.

## Configuring the external volumes

1. Click the Server Manager icon in the taskbar.
2. In the left pane, expand Storage and click Disk Management.
3. Right-click the first external volume and choose Initialize Disk.
4. In the right pane, right-click the volume and choose New Simple Volume...
5. At the Welcome window, click Next.
6. At the Specify Volume Size window, leave the default selection, and click Next.
7. At the Assign Drive Letter or Path window, choose a drive letter, and click Next.
8. At the Format Partition window, choose NTFS and 64K allocation unit size, and click Next.
9. At the Completing the New Simple Volume Wizard window, click Finish.
10. Repeat steps 3 through 9 for the remaining external volumes.

## Configuring the Windows cluster using the Failover Cluster Manager

### Adding the clustering feature and validating the cluster

1. Open Server Manager and choose Features.
2. Click Add Features.
3. Choose Failover Clustering, and proceed through the wizard to complete the addition of the feature.
4. Repeat on the second node.
5. On the first node, open the Windows Failover Cluster Manager.
6. Click Validate a Configuration.
7. Click Next.
8. Add server names, and click next.
9. On the Testing Options screen, click Next.
10. On the Confirmation screen, click Next.
11. After the cluster validation completes, click Finish.

### Creating and configuring the cluster

1. On the first node, open the Windows Failover Cluster Manager.
2. Click Create a Cluster.
3. Click Next.
4. Add the server name, and click Next.
5. Assign a cluster name, specify a cluster IP address, and click Next.
6. On the confirmation screen, click Next.
7. After the cluster creation completes, click Finish.
8. In the Failover Cluster Manager, right-click Storage and choose Add a Disk.
9. Select the LUN reserved for MSDTC use and click OK.

10. In the Failover Cluster Manager, right-click Services and Applications and choose Configure a Service or Application.
11. Select the Distributed Transaction Coordinator (DTC) application, and choose Next.
12. Assign a name for the MSDTC service and a network address, and click Next.
13. Choose a disk resource, the one configured in step 9 above.
14. On the Confirmation screen, click Next.
15. After the addition of DTC to the cluster, click Finish.
16. In the Failover Cluster Manager, right-click Nodes, and choose Add Node.
17. Supply the server name of node 2 and complete the Add Node wizard.

## Installing the SQL Server 2012 cluster instances

To install SQL Server named instances on the cluster and make them highly available, you must use the cluster-specific installation option. We followed the process below four times, one time per SQL Server named instance.

### Installing SQL Server 2012 on the first Failover Cluster node

1. Mount the SQL Server install media and choose Installation, then New SQL Server Failover Cluster installation.
2. Click Next to proceed to the Setup Support Rules screen.
3. On the Setup Support Rules screen, click Next.
4. Enter a product key and click Next.
5. Accept the License Agreement and click Next.
6. Choose SQL Server Feature Installation and click Next.
7. Choose Database Engine Services, SQL Server Replication, Full-Text and Semantic Extractions for Search, Client Tools Connectivity, Client Tools Backwards Compatibility, Management Tools – Basic, Management Tools – Complete, then click Next.
8. On the Feature Rules screen, ensure all the states are passed, and click Next.
9. On the Instance configuration screen, enter a name for the SQL cluster under SQL Server Network Name, choose Named Instance, provide an instance name, and click Next.
10. On the Disk Space Requirements screen, click Next.
11. On the Cluster Resource Group screen, click Next.
12. On the Cluster Disk Selection screen, select the applicable disks for this names instance. The disk resources must already be online, formatted, and added to the Failover Cluster Manager.
13. On the Cluster Network configuration screen, configure a static IP address for the SQL Server named instance, and click Next.
14. On the Server Configuration screen, provide Active Directory authentication details for the SQL Server and SQL Server Agent accounts, and click Next.
15. On the Database Engine Configuration screen, click the Server Configuration tab, click Mixed Mode, provide an SA password, and click the Add current user button.
16. On the Database Engine Configuration screen, click the Data Directories tab and provide the appropriate paths to the data, logs, and tempdb locations.
17. On the Error Reporting screen, click next.
18. On the Cluster Installation Rules screen, click Next.
19. On the Ready to Install screen, click Install.
20. After installation, connect to the instance via SQL Server Management Studio, right-click the instance, choose database settings, and set Recovery Interval to at least 1000.

### Adding a node to the SQL Server 2012 Failover Cluster instance

1. On the second node, mount the SQL Server media and start the installation wizard.
2. Choose Add node to a SQL Server failover cluster.

3. Proceed through the first introductory screens to specify product key, license agreement, and setup support rules.
4. On the Cluster Node Configuration screen, choose the SQL Server instance that you wish to add the node to, and click Next.
5. On the Cluster Network Configuration screen, select the network settings you wish to use, and click Next.
6. On the Service Accounts screen, fill in Active Directory authentication details, and click Next.
7. On the Error Reporting screen, click next.
8. On the Add Node Rules screen, click Next.
9. On the Ready to Add Node screen, click Install.

## Configuring the database (DVD Store)

### Data generation overview

We generated data using the Install.pl script included with DVD Store version 2.1 (DS2), providing the parameters for our 40GB database size and the database platform on which we ran, Microsoft SQL Server. We ran the Install.pl script on a utility system running Linux. The database schema was also generated by the Install.pl script.

After processing the data generation, we transferred the data files and schema creation files to a Windows-based system running SQL Server 2008 R2 SP1. We built the 40GB database in SQL Server 2012, and then performed a full backup, storing the backup file on shared storage LUN we designated for backup storage. We used that backup file to restore on all instances between test runs. We performed this procedure once, and used the same backup file for all both nodes in all three test runs.

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, allowing for approximately 40 percent free space in the database files. Besides this file size modification, the database schema was created and loaded according to the DVD Store documentation. Specifically, we followed the steps below:

1. Generate the data and create the database and file structure using database creation scripts in the DS2 download. We made size modifications specific to our 40GB database and the appropriate changes to drive letters.
2. Transfer the files from our Linux data generation system to a Windows system running SQL Server.
3. Create database tables, stored procedures, and objects using the provided DVD Store scripts.
4. Set the database recovery model to bulk-logged to prevent excess logging.
5. Load 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. Create indices, full-text catalogs, primary keys, and foreign keys using the database-creation scripts.
7. Update statistics on each table according to database-creation scripts, which sample 18 percent of the table data.
8. On the SQL Server instance, create 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. Set the database recovery model back to full.
10. Create the necessary full text index using SQL Server Management Studio.
11. Create a database user, and map this user to the SQL Server login.
12. Perform a full backup of the database. This backup allowed us to restore the databases to a pristine state relatively quickly between tests.

### 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 track the last OPM reported by each client/instance pair. We used four client machines, each client targeting one SQL Server clustered instance.

Each complete test cycle consisted of the general steps listed below. For each scenario, we ran three test cycles, and chose the median outcome.

1. Clean up prior outputs from the host system and all client driver systems.
1. Drop all databases from all target SQL Server cluster instances.
2. Restore all databases on all target SQL Server cluster instances.
3. Move all SQL Server cluster resources to node A.
4. Reboot node B.
5. Move all SQL Server cluster resources to node B.
6. Reboot node A.
7. Take all SQL Server cluster resources offline.
8. Bring all SQL Server cluster resources online.
9. Move SQL Server instances 1 and 2 to node A.
10. Move SQL Server instances 3 and 4 to node B.
11. Let the test server idle for ten minutes.
12. Start the DVD Store driver on all respective clients.

We used the following DVD Store parameters for testing the virtual machines in this study:

```
ds2sqlserverdriver.exe --target=<target_IP>\<target_instance_name> --  
run_time=30 --n_threads=30 --db_size=40GB --think_time=0 --warmup_time=5
```

## ABOUT PRINCIPLED TECHNOLOGIES



Principled Technologies, Inc.  
1007 Slater Road, Suite 300  
Durham, NC, 27703  
[www.principledtechnologies.com](http://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.

---