



Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

*Kent R. Swalin
Quang Le
IBM Systems and Technology Group*

Executive Overview

File Servers are well known for providing common storage repositories for shared documents, images, and user file storage. The Server Message Block (SMB) protocol is the file sharing protocol that is used by default on Windows file servers. Now with SMB Direct in Windows Server® 2012, a Scale-Out File Server using IBM System x Servers, IBM Storage, IBM System Networking and Chelsio® T4 adapters can provide high performance file shares that are continuously available for applications such as Microsoft SQL Server® 2012. SMB Direct which is SMB over Remote Direct Memory Access (RDMA), was introduced in Windows Server 2012. SMB Direct has the typical benefits of SMB file storage: it is easy to provision, manage, and migrate.

File server administrators can provide scale-out file shares and continuously available file services to server applications and respond to increased demands quickly by simply adding more servers to the cluster or bringing more storage online. All of this can be done in a production environment and it is transparent to the server application.

Introduction

This document describes a Scale-Out solution that uses a two-node file server cluster and two servers running Microsoft SQL Server 2012 on Windows Server 2012. The application servers in this solution can access file shares using any of the cluster nodes. This is achieved by leveraging new Windows Server Failover Clustering features and new capabilities in the new version of Windows file server protocol - SMB 3.0. SMB Direct supports the use of network adapters that have remote direct memory access (RDMA) capability. SMB Direct functionality requires that the SMB client and SMB server support SMB 3.0.

The new SMB Direct storage protocol in Windows Server 2012 provides for extremely fast responses to network requests, and as a result, makes remote file storage respond similarly to directly attached block storage. SMB Direct also offloads CPU cycles when transferring data over the network, which leaves more cycles available to server applications.

Key benefits provided by Scale-Out File Server in Windows Server 2012 include:

- **Active-Active file shares:** All cluster nodes can accept and serve SMB client requests. By making the file share content accessible through all cluster nodes simultaneously, SMB 3.0 clusters and clients cooperate to provide transparent failover to alternative cluster nodes during planned maintenance and unplanned failures without service interruption.
- **Increased bandwidth:** The maximum share bandwidth is the total bandwidth of all file server cluster nodes. Unlike previous versions of Windows Server, the total bandwidth is no longer constrained to the bandwidth of a single cluster node, but rather the capability of the storage system. Bandwidth can be increased by adding nodes to the Scale-Out cluster.
- **CHKDSK with zero downtime:** CHKDSK in Windows Server 2012 is significantly enhanced to dramatically shorten the time a file system is offline for repair. Clustered shared volumes (CSVs) in Windows Server 2012 take this one step further and eliminate the offline phase. A CSV File System (CSVFS) can perform CHKDSK without impacting applications with open handles on the file system.
- **Clustered Shared Volume cache:** CSVs in Windows Server 2012 introduces support for a read cache, which can significantly improve performance in certain scenarios.
- **Simpler management:** With Scale-Out File Servers necessary CSVs and file shares can be created and added. It is no longer necessary to create multiple clustered file servers, each with separate cluster disks, and then develop placement policies to ensure activity on each cluster node.

Microsoft SQL Server 2012 offers full support for storing databases on SMB 3.0 file shares. Scale-Out file shares provide the ability to share the same folder from multiple nodes in a cluster. Utilizing a Scale-Out File Server can reduce SQL Server solution complexity. No longer does an implementation of SQL Server require large amounts of local storage or attaching a server to a Storage Area Network, when there is a Scale-Out File Server available that allows a set of shares to be accessed by Universal Naming Convention (UNC) paths.

Components

This Scale-Out File Server architecture is comprised of IBM System x servers, IBM storage, and IBM System Networking running Microsoft's Windows 2012 operating system. Each component provides a key element to the overall solution.

IBM System x3650 M4 Server

The System x3650 M4 provides the ultimate balance of performance, flexibility and cost efficiency. With technology choices like the Intel Xeon processor E5-2600 product family, eXFlash high IOPS SSD solutions and 10Gb Ethernet, the x3650 M4 meets virtualization challenges with ease. IBM System x servers and IBM System Storage help create a resilient, cost-effective and flexible infrastructure for securely storing information and mitigating risk

As business grow and change - and continue with the same tight budgets – the need efficient and flexible IT solutions increases. The new generation of x86 servers offers improved performance and efficiency designed to provide the power to transform any business. This can help improve operational efficiency optimize for business-critical workloads and unlock data for real business insights.

The new x3650 M4 servers offer many enhanced features:

- Industry-leading performance
- Platforms designed for flexibility and growth
- eXFlash
- Slotless 10Gb Networking (Virtual Fabric)
- High-capacity, high-performance memory



Figure 1 IBM System x3650 M4

IBM System x3550 M4 Server

Big things do come in small packages. The x3550 M4 is a 1U powerhouse with leadership high availability features, such as predictive failure analysis, redundant power and disk mirroring capabilities. The M4 generation of this business-critical platform offers up to twice the performance as previous generation servers and provides a solid foundation for cloud deployments. Hyper-V in Windows Server 2008 R2 supports 64 processors and 1TB memory on the host and 4 virtual processors and 64 GB memory in the guest.

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

Page 4

By virtualizing with Microsoft Hyper-V technology on IBM System x3550 M4 servers, businesses can make the best use of server hardware investments by consolidating multiple server roles as separate virtual machines running on a single physical machine. This results in fewer physical servers to manage. Running the same hypervisor allows IT administrators to manage more servers than without virtualization. Since applications are managed within virtual machines and isolated from the actual hardware, this again simplifies the physical server, maximizing uptime. Greater IT flexibility is also realized by the ability of virtualization to streamline the deployment and upgrade processes in both physical and virtual machines.

In addition, clients can see a lower total cost of ownership (TCO) from reduced hardware requirements by maximizing the resource utilization of each server purchased. This usually results in a substantial reduction in the number of physical servers required to support operations, thereby lowering capital acquisition, power consumption, and cooling costs. Businesses also benefit from infrastructure optimization that simplifies and standardizes IT administration, thus helping to control labor costs. Reduced downtime costs can be achieved by the ability to easily move and balance workloads across different resources, improve operational agility, and provide flexibility in managing maintenance schedules and responding to today's dynamic IT demands.



Figure 2 IBM System x3550 M4

Chelsio T4 10GbE Unified Wire Ethernet adapters

Chelsio's T420-CR adapters powered by its fourth-generation T4 ASIC provide the highest 10GbE performance available while dramatically lowering host-system processing overhead with integrated hardware offload of TCP/IP, iSCSI, FCoE and iWARP RDMA. The T420-CR frees up host CPU cycles for useful applications, resulting in increased bandwidth, lower latency, lower power consumption and improved overall system efficiency.



Figure 3 Chelsio T420-CR

IBM DS3500 Storage

The IBM DS3500 disk system is an ideal platform for Scale-Out environments. Highly expandable, clients can start with just the right amount of storage, and then scale accordingly as storage workloads increase. The DS3500 provides powerful, robust, yet easy to manage storage that serves as a platform for many Microsoft-based reference architectures.

The DS3500 can expand to support up to 192 drives, which can be an intermix of 2.5", 3.5", and SSDs. High performance SAS and NL SAS disks provide clients a wide range of deployment options. Host attachment options include 6 Gbps SAS, 8 Gbps Fibre Channel, and 1/10 Gbps iSCSI for easy guest clustering support. A full complement of advanced features such as Copy Services, Flash Copy, partitioning, full MPIO/VSS support and remote mirroring ensure the DS3500 can grow as storage and application needs increase in complexity.



Figure 4 IBM DS3524 Storage

IBM System Networking RackSwitch G8124E 10GbE Switch

The RackSwitch G8124E is a 10Gb Ethernet top of rack aggregation switch with unmatched line-rate Layer 2/3 performance. It uses a wire-speed, non-blocking switching fabric that provides simultaneous wire-speed transport of multiple packets at low latency on all ports.

The G8124E contains 24 ten Gigabit Small Form-factor, Pluggable (SFP+) slots and two 1Gb management ports. The 10Gb SFP+ slots can accept 1Gb copper transceivers, 10Gb optical transceivers, or Direct Attach Cables (DAC).



Figure 5 IBM RackSwitch G8124E 10GbE Switch

Overview of the Solution Configuration

The physical configuration of the servers in this solution is shown below in Figure 6. The two x3650 M4's are configured as a Windows 2012 cluster with the file server role enabled. They are attached to a DS3500 using 8GB Fibre Channel. In the event of a cluster node unavailability due to hardware failure or system maintenance, the other cluster node will, transparently to the client, take over serving the file shares. The two x3550 M4's are running SQL Server 2012 with all associated SQL files located on the file server as shown in Figure 9.

The DS3500 storage is configured using IBM Storage Manager Client GUI. Four data volumes were created, two for SQL data and two for Log files. These volumes were initialized and formatted using Windows Server 2012. They were added into the File Server cluster as Cluster Shared Volumes (CSV).

Each IBM System x server had one Chelsio T4 adapter installed. Latest version of the NIC-TOE driver for Windows Server 2012 was downloaded from Chelsio and installed. They were configured with a static IP address using Windows Server 2012 Network Connections.

The detailed steps to create the solution are documented in the appendix.

**Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection
– Solution Utilizing IBM System x Servers**

Page 7

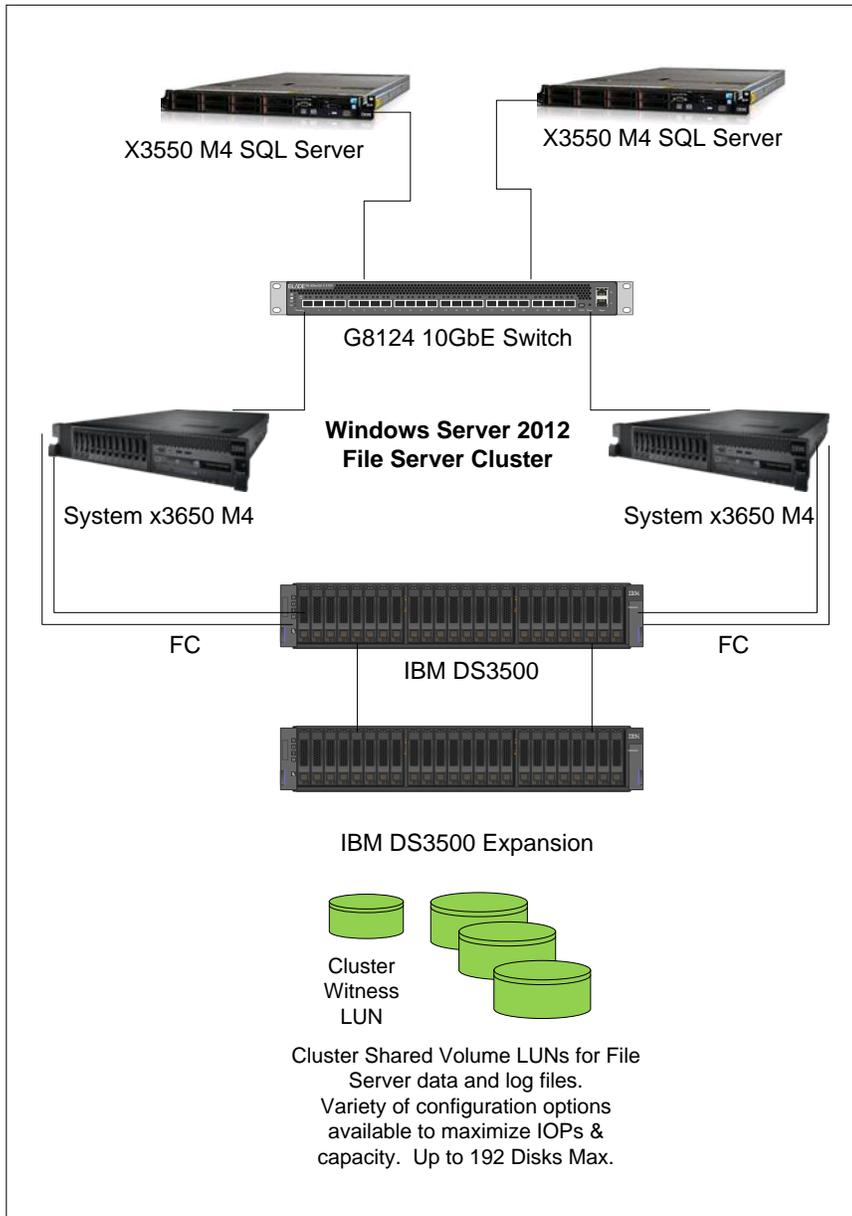


Figure 6 Physical view of the Cluster Configuration

Figure 7 shows a view of the interaction between the SMB Direct client and the SMB Direct server. In this scenario the SMB Direct clients are the two SQL Servers while the SMB Direct server is the x3650 M4 cluster. The SMB Direct client and server roles are included in Windows Server 2012.

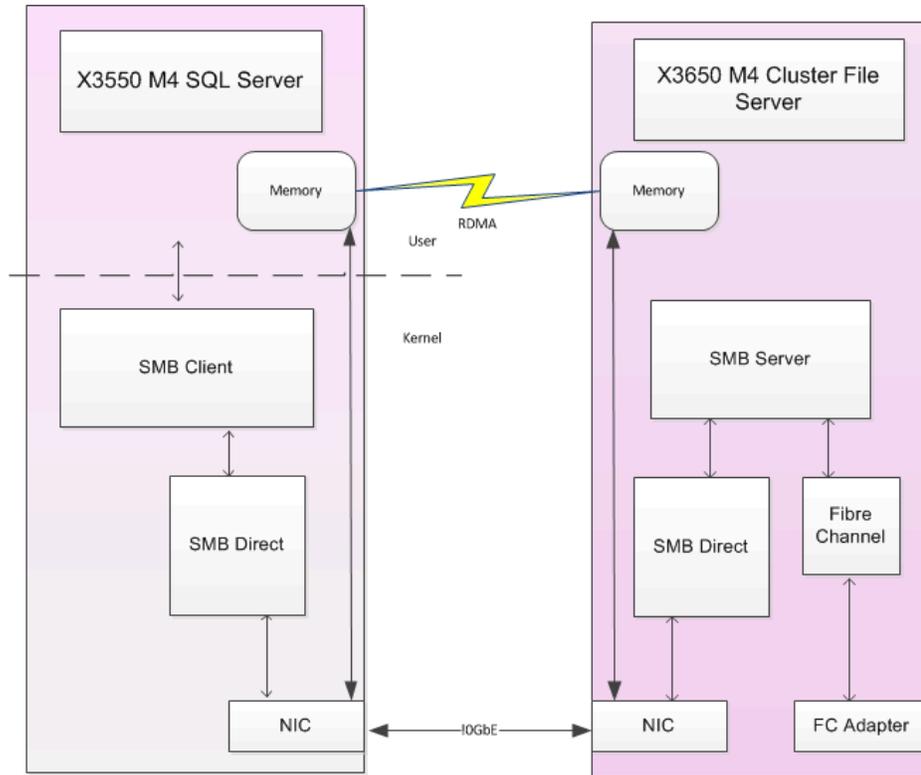


Figure 7 SMB Direct Client and Server

Figure 8 shows the two options when adding the File Server Role to the cluster. For this solution the option for File Server for scale-out application data was chosen.

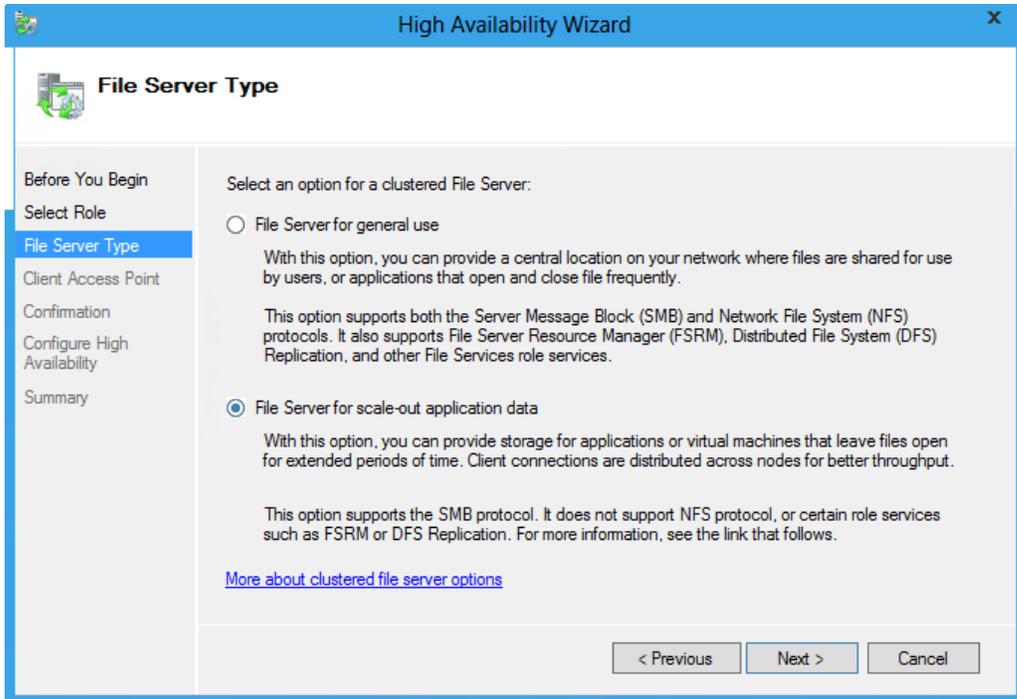


Figure 8 Creating the High Available File Server

Table 1 and table 2 detail the configuration of the cluster nodes and SQL Servers.

X3650 M4 File Server Node Configuration	
Processors	Two Intel Xeon Processor E5-2670 with eight 2.6GHz cores
Total Physical Memory	64GB using DDR-3 RDIMMs
Operating System	Microsoft Windows Server 2012 Datacenter Edition
Network	<ul style="list-style-type: none"> • Four on-board Intel 1GbE adapters • One Chelsio T420 10GbE adapter
Fibre Channel	One Brocade 8Gb FC Dual-port HBA for IBM System x

Table 1 IBM System x3650 M4 server configuration

X3550 M4 SQL Server Configuration	
Processors	Two Intel Xeon Processor E5-2670 with eight 2.6GHz cores
Total Physical Memory	96GB using DDR-3 RDIMMs
Operating System	Microsoft Windows Server 2012 Datacenter Edition
Network	<ul style="list-style-type: none"> • Four on-board Intel 1GbE adapters • One Chelsio T420 10GbE adapter

Table 2 IBM System x3550 M4 server configuration

SQL Server configuration

The standard method for installing SQL Server 2012 is followed; however, when setting up the data directories, direct the files toward the UNC path which reside on the x3650 M4 cluster. For SQL Server setup to work properly, the appropriate ACL permissions should be set on these UNC directories.

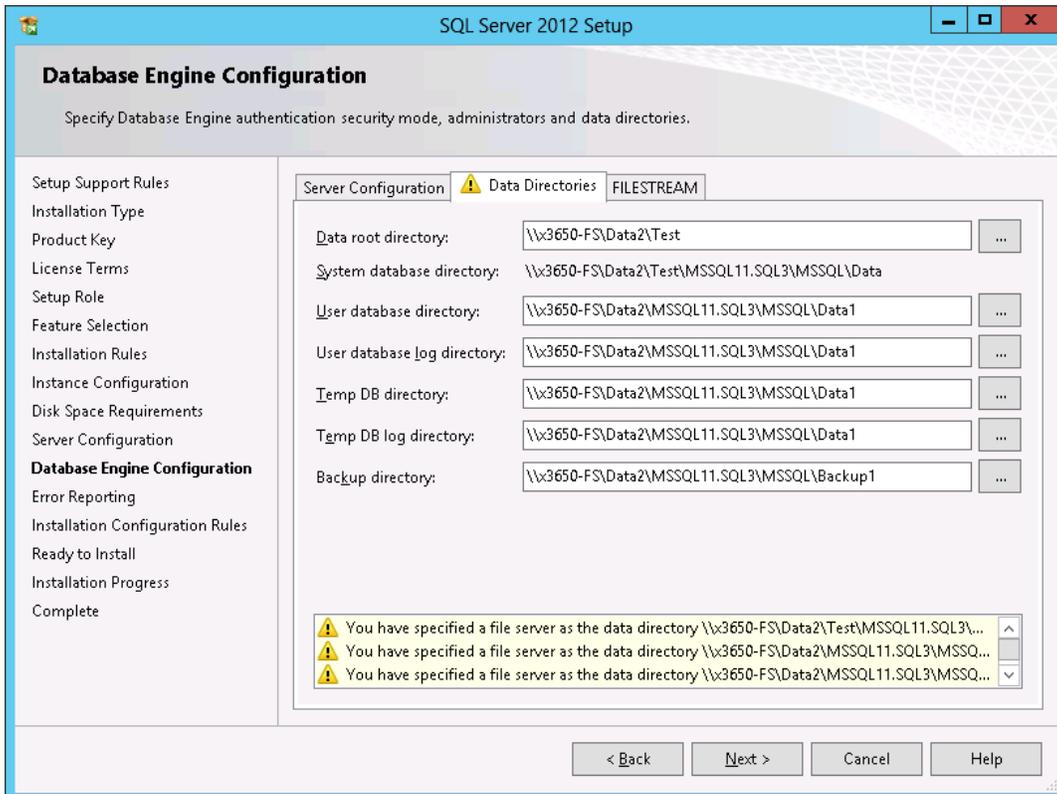


Figure 9 SQL Server 2012 Setup

Performance Monitoring

Performance Monitor was used to capture the relevant disk performance data for an SMB Client and an SMB Server and shown on Figures 10 and 11. These are the counters that should be checked to monitor IO performance. These are examples only and should not be used as a guide.

\\SQLSVR1	
Processor Information	_Total
% Processor Time	49.418
RDMA Activity	Chelsio T4 10GbE Adapter
RDMA Accepted Connections	0.000
RDMA Active Connections	2.000
RDMA Completion Queue Errors	0.000
RDMA Connection Errors	0.000
RDMA Failed Connection Attempts	0.000
RDMA Inbound Bytes/sec	0.000
RDMA Inbound Frames/sec	0.000
RDMA Initiated Connections	2.000
RDMA Outbound Bytes/sec	0.000
RDMA Outbound Frames/sec	0.000
SMB Client Shares	\\x3650-FS\Data1
Avg. Bytes/Read	32,304.974
Avg. Bytes/Write	59,531.062
Avg. Data Bytes/Request	39,396.013
Avg. Data Queue Length	3.436
Avg. Read Queue Length	0.612
Avg. sec/Data Request	0.011
Avg. sec/Read	0.003
Avg. sec/Write	0.005
Avg. Write Queue Length	2.824
Credit Stalls/sec	0.000
Current Data Queue Length	0.000
Data Bytes/sec	12,251,349.779
Data Requests/sec	310.979
Metadata Requests/sec	0.000
Read Bytes/sec	7,429,652.653
Read Requests/sec	229.985
Write Bytes/sec	4,821,697.126
Write Requests/sec	80.995

Figure 10 SMB Client Performance Monitor Data

\\X3650M4-2	
Processor Information	_Total
% Processor Time	0.546
SMB Server Shares	
\\X3650-FS\Data1	
% Persistent Handles	55.370
% Resilient Handles	0.000
Avg. Bytes/Read	32,283.744
Avg. Bytes/Write	68,101.626
Avg. Data Bytes/Request	49,215.834
Avg. Data Queue Length	5.342
Avg. Read Queue Length	0.578
Avg. sec/Data Request	0.014
Avg. sec/Read	0.003
Avg. sec/Request	0.014
Avg. sec/Write	0.005
Avg. Write Queue Length	4.764
Current Data Queue Length	2.000
Current Durable Open File Count	39.000
Current Open File Count	12.000
Current Pending Requests	2.000
Data Bytes/sec	18,947,752.725
Data Requests/sec	384.993
Files Opened/sec	0.000
Metadata Requests/sec	0.000
Read Bytes/sec	6,553,481.271
Read Requests/sec	202.996
Received Bytes/sec	12,443,950.554
Requests/sec	386.993
Sent Bytes/sec	6,584,280.713
Total Durable Handle Reopen Count	0.000
Total Failed Durable Handle Reopen Count	0.000
Total Failed Persistent Handle Reopen Count	0.000
Total Failed Resilient Handle Reopen Count	0.000

Figure 11 SMB Server Performance Monitor Data

Conclusion

Utilizing IBM hardware combined with Microsoft Windows Server 2012, it is possible to create a high-performance, interoperable, and reliable application file server. With Enterprise-class hardware, fault tolerance is achieved by configuring a robust collection of industry-leading servers, storage systems and networking components from IBM. The SMB direct framework promotes a standardized and highly manageable environment which helps satisfy even the most challenging business critical SQL Server storage demands.

Appendix: Demonstration System Setup

Figure 12 below illustrates the lab environment used to generate the results documented in this white paper.

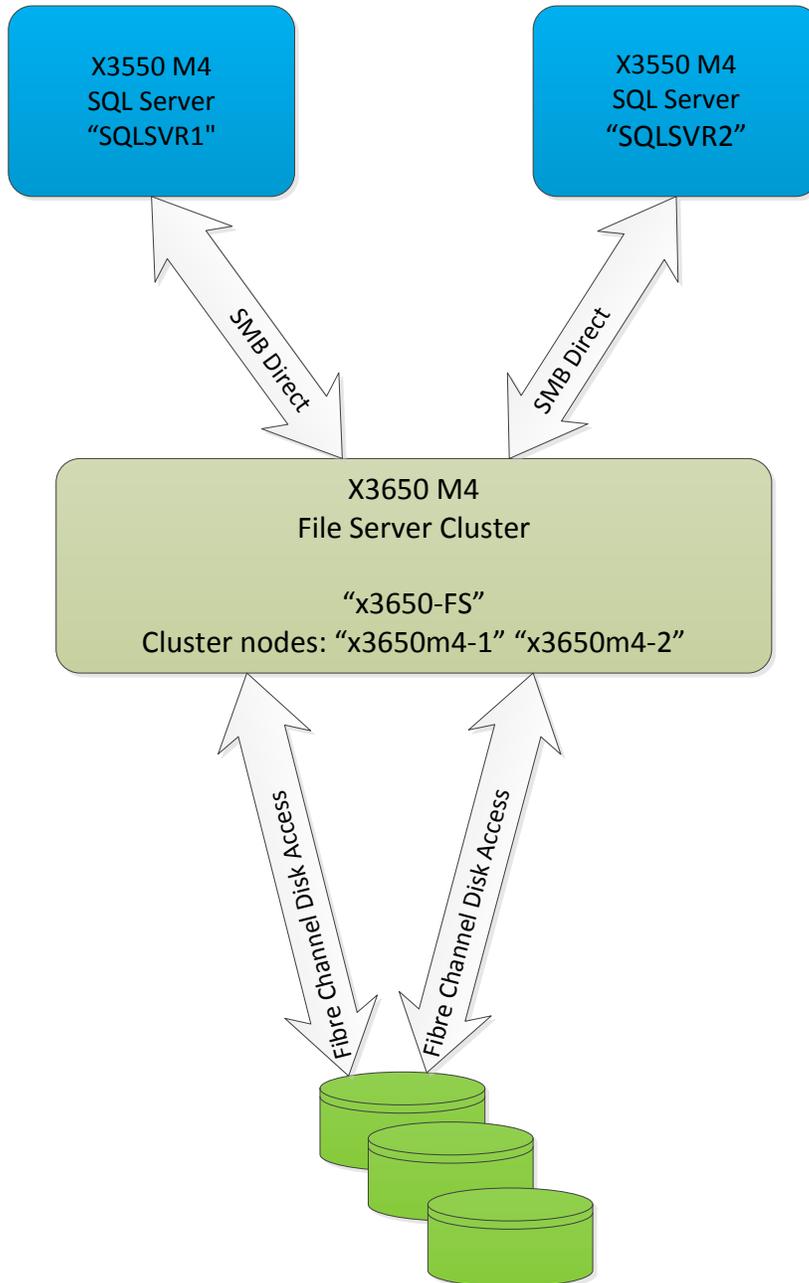
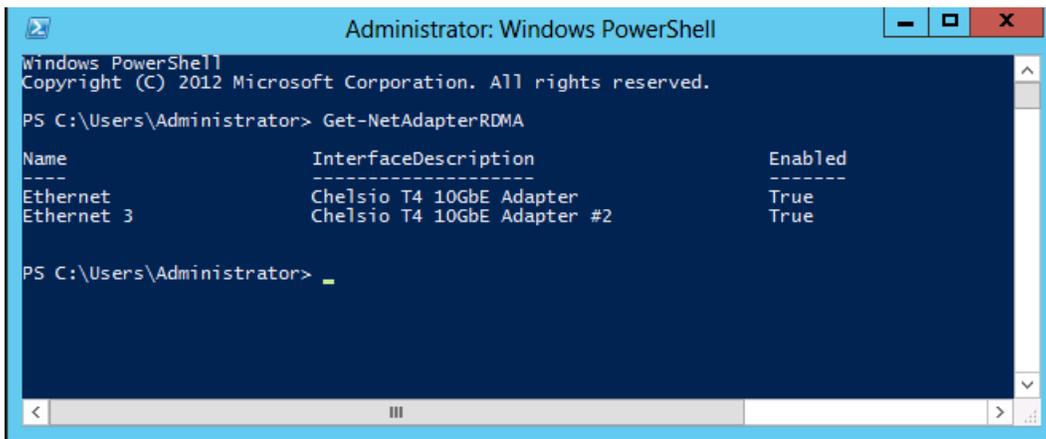


Figure 12 Overview of the demonstration environment

Install the Chelsio T420-CR 10 GigE Adapter

The Chelsio T4 Unified Wire driver for Windows 2012 Server can be downloaded from Chelsio's support website at <http://service.chelsio.com>. The driver is also available from Windows update, and if the user instructs Windows 2012 to find a driver for the Chelsio card in Device Manager, it will be able to locate and install it.

SMB Direct is installed and enabled by default when the driver is installed. To determine if the driver was installed correctly and that RDMA is properly enabled, use the *Get-NetAdapterRdma* command in PowerShell.



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2012 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> Get-NetAdapterRdma

Name                InterfaceDescription          Enabled
----                -
Ethernet            Chelsio T4 10GbE Adapter      True
Ethernet 3         Chelsio T4 10GbE Adapter #2   True

PS C:\Users\Administrator> _
```

Storage Configuration

The following steps illustrate how quick and easy it is to configure DS3500 storage.

Create each of the arrays needed for the configuration. Assign an array name, and select the number of disks and RAID type to be used (Figure 13).

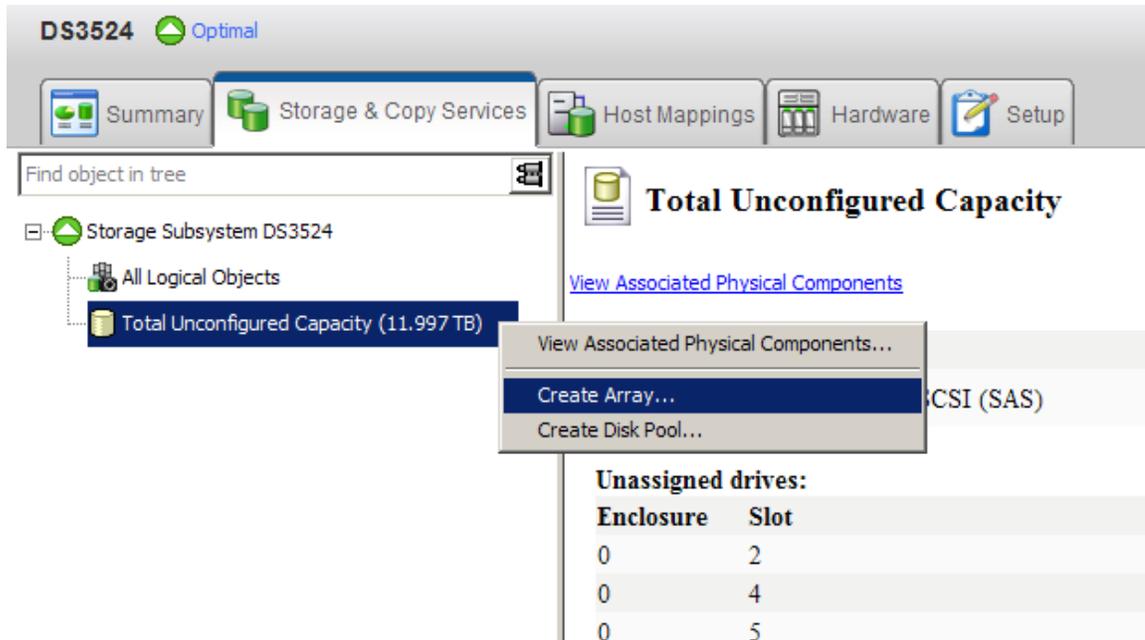


Figure 13 DS3500 Array Creation

Disks can now be created off each of the arrays (Figure 14)

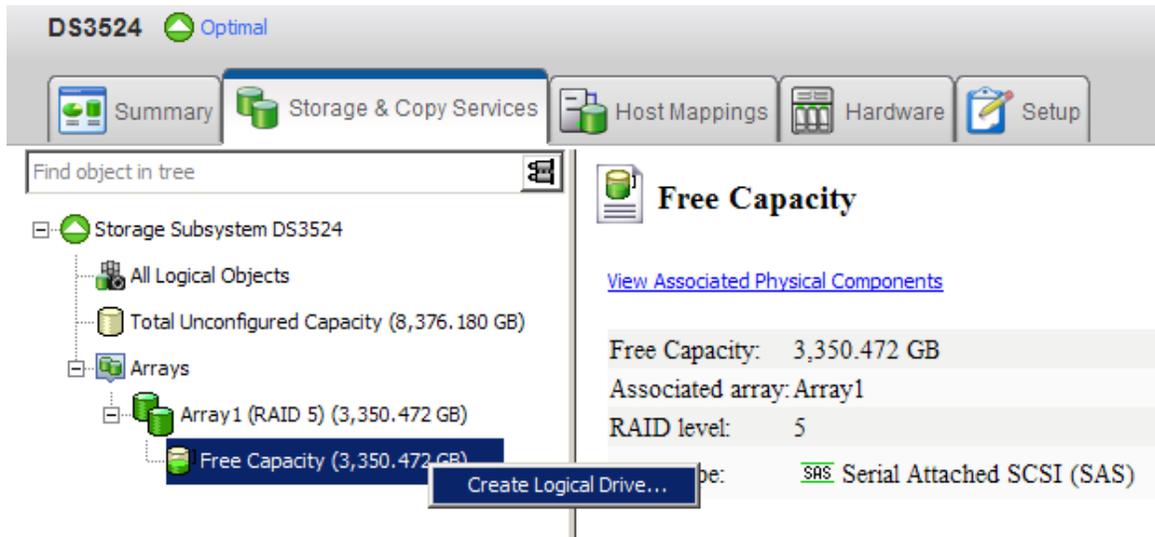


Figure 14 DS3500 Drive Creation

- A Host Group is a logical group that will contain the host servers that should all see the same storage volumes; in this case add the hosts x3650m4-1 and x3650m4-2.
- Create a Host Group to that will contain each of the host servers (Figure12)

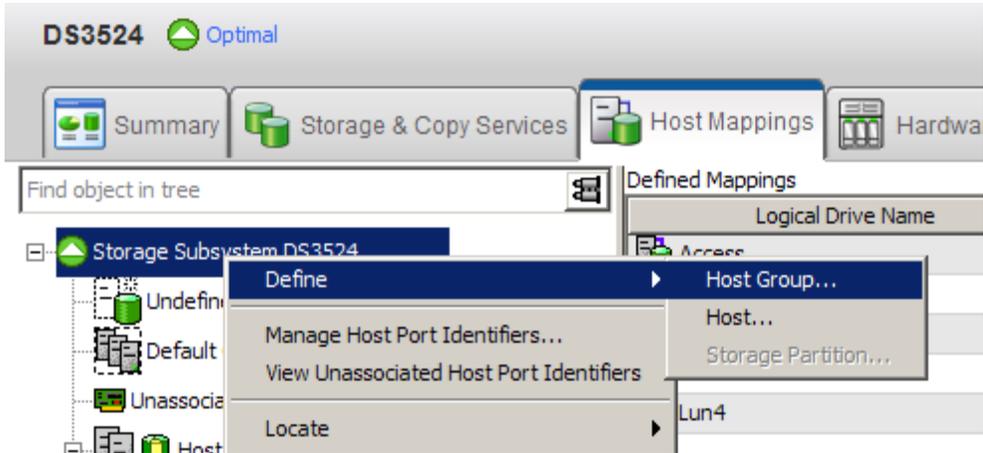


Figure 15 Host Group Creation

Multipath I/O (MPIO) Fault-Tolerance Driver

Multipath I/O is used to provide balanced and fault tolerant paths to DS3500. This requires an additional driver to be installed on the host servers before attaching the storage.

Install and Configure the File Server Cluster

In this section, steps to create and configure the file server cluster will be shown:

Install the Required Roles and Features

Both servers x3650m4-1 and x3650m4-2 were configured as files servers and cluster nodes. There are two ways to do this, PowerShell commands or via the GUI.

Using PowerShell Commands

Open the PowerShell windows:

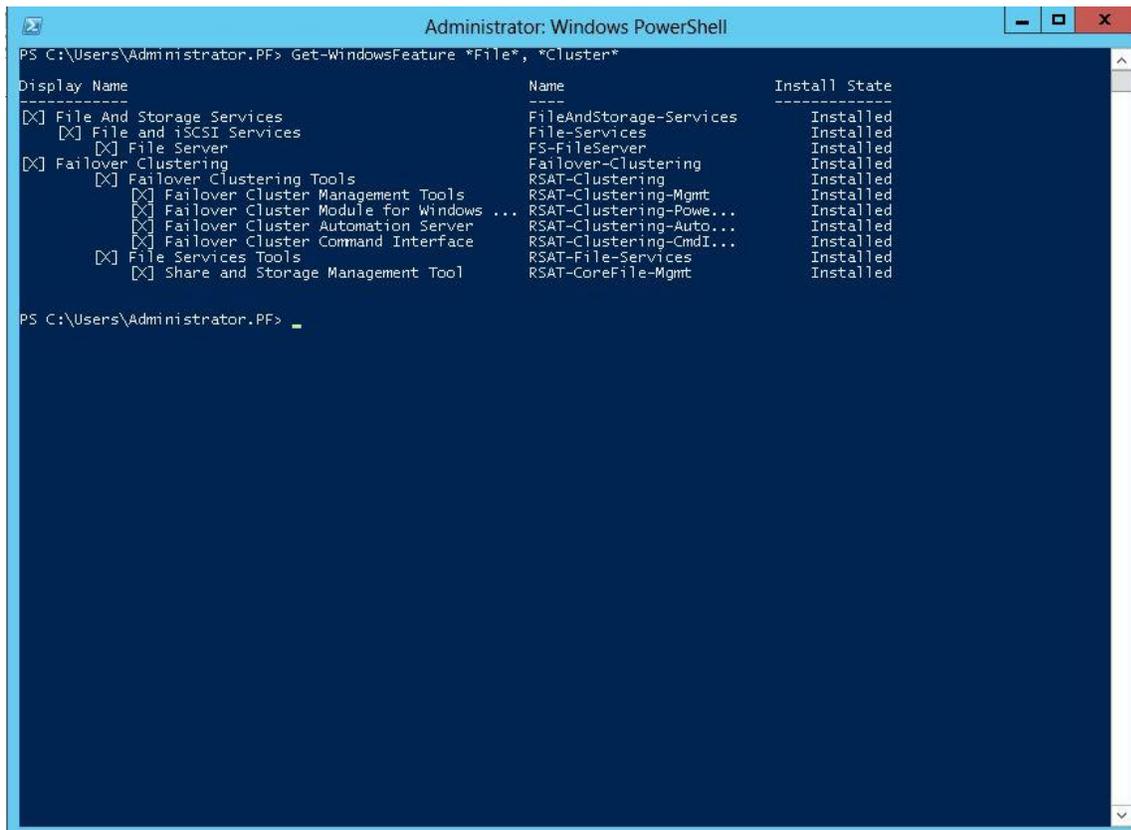
Type “Install-WindowsFeature File-Services, FS-FileServer, Failover-Clustering” <Enter>

Type “Install-WindowsFeature RSAT-Clustering –IncludeAllSubfeature” <Enter>

Type “Install-WindowsFeature RSAT-File-Services –IncludeAllsubFeature” <Enter>

Upon completing the previous three commands, execute the following command to display the status of the file server cluster:

Type “Get-WindowsFeature *File*, *Cluster*” <Enter>



```
Administrator: Windows PowerShell
PS C:\Users\Administrator.PF> Get-WindowsFeature *File*, *Cluster*

Display Name                                     Name                                     Install State
-----
[X] File And Storage Services                   FileAndStorage-Services                 Installed
[X] File and iSCSI Services                     File-Services                           Installed
[X] File Server                                 FS-FileServer                           Installed
[X] Failover Clustering                        Failover-Clustering                     Installed
[X] Failover Clustering Tools                  RSAT-Clustering                         Installed
[X] Failover Cluster Management Tools           RSAT-Clustering-Mgmt                    Installed
[X] Failover Cluster Module for Windows ...    RSAT-Clustering-Powe...                 Installed
[X] Failover Cluster Automation Server         RSAT-Clustering-Auto...                 Installed
[X] Failover Cluster Command Interface        RSAT-Clustering-CmdI...                 Installed
[X] File Services Tools                        RSAT-File-Services                      Installed
[X] Share and Storage Management Tool          RSAT-CoreFile-Mgmt                       Installed

PS C:\Users\Administrator.PF> _
```

Using the GUI

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

Page 19

Open “Server Manager”

For both servers x3650m4-1 and x3650m4-2, select “Add Roles and Features” and check the “File and Storage Services” box.

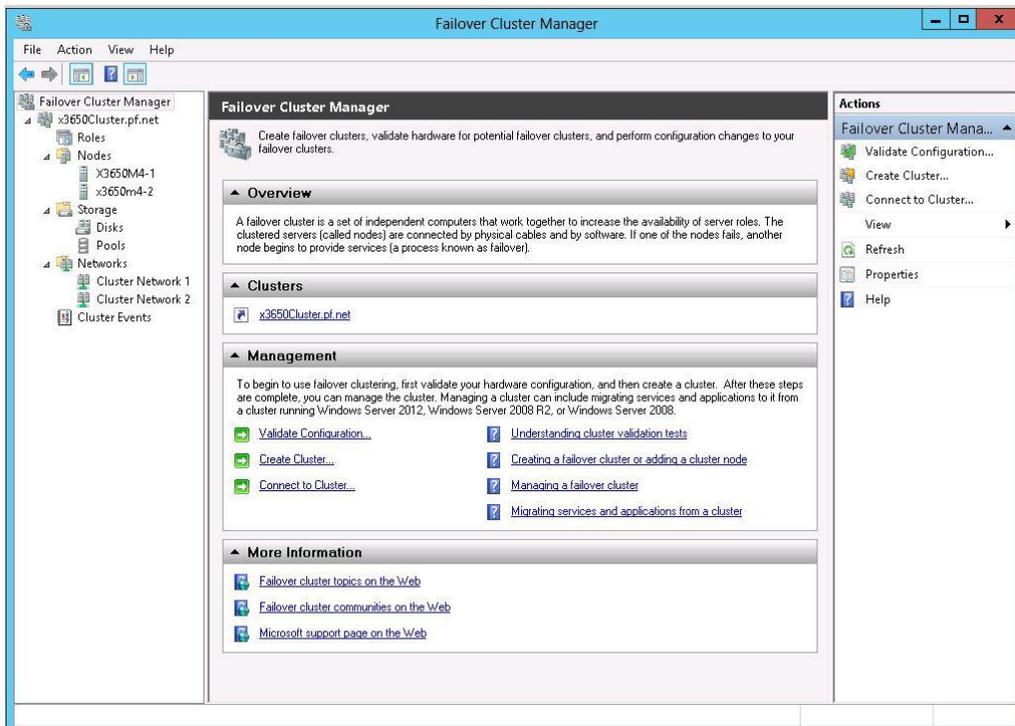
On the next screen, select “Add Feature” and check the “Failover Clustering” box.

Click “Install”

Validate the Cluster

The next step is to validate the newly created cluster.

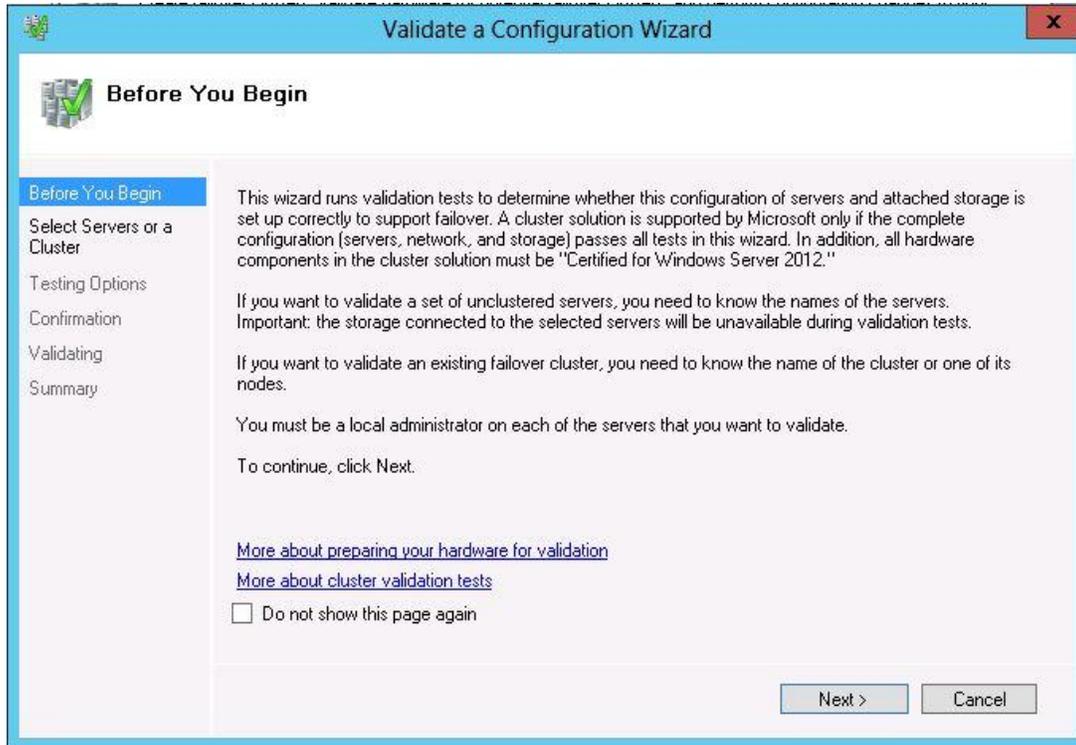
Open the “Failover Cluster Manager”



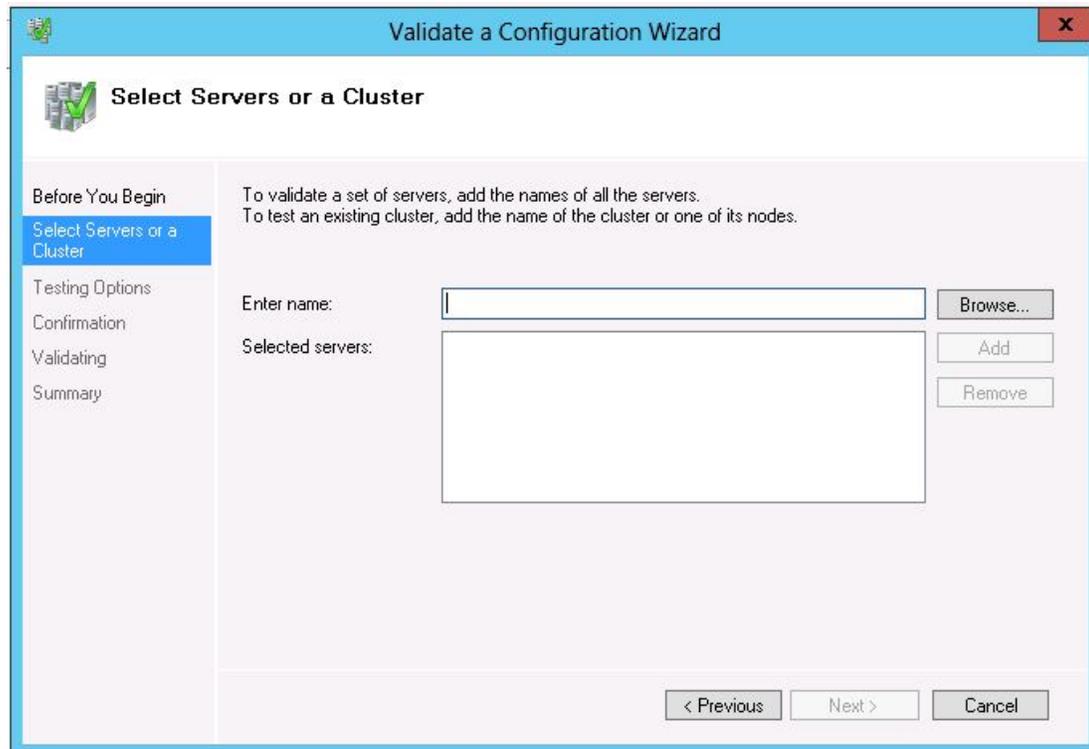
Click on “Validate Configuration”

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

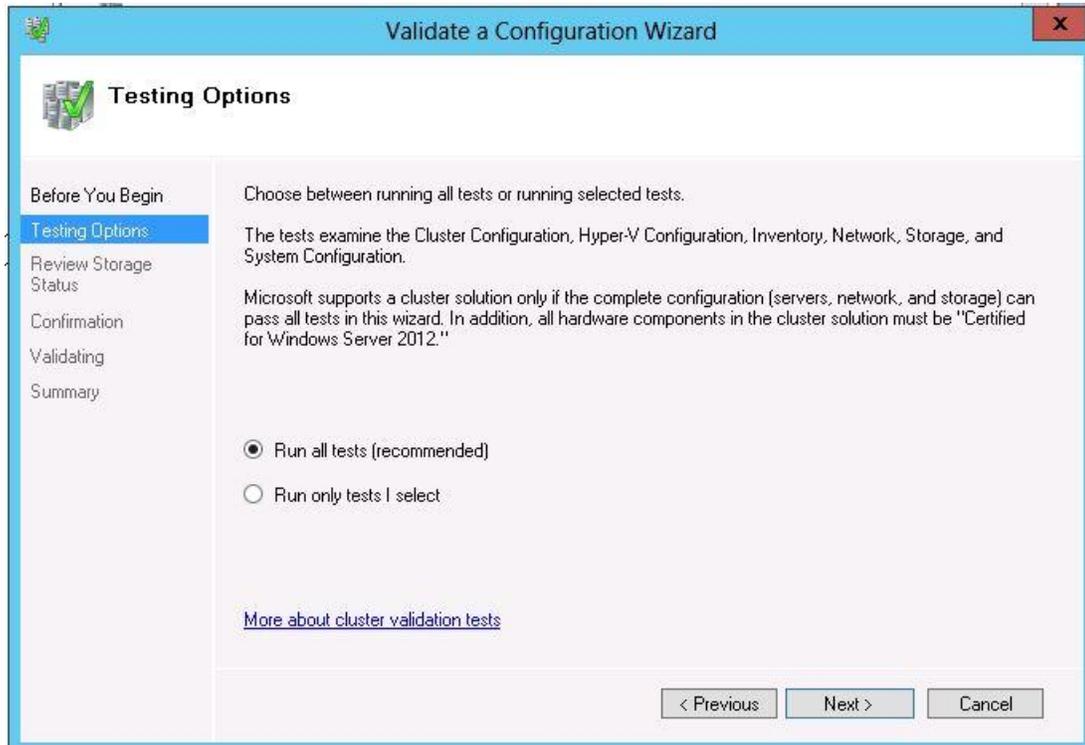
Page 20



Click "Next"



Enter the names of the servers, one at a time followed by clicking "Add."



Select “Run all tests” (recommended), click “Next”

On the next screen, click “Next” to confirm

The validation process will take a few minutes to complete. It should not return any errors. If there are errors, address the issues listed in the validation report.

Create a Failover Cluster

There are two ways to do this, PowerShell commands or via the “Failover Cluster Manager”

Using PowerShell Command

Type “New-Cluster –Name x3650Cluster –Node x3650m4-1, x3650m4-2” <Enter>

In this example, the name of the cluster will be x3650Cluster. It has two nodes, namely x3650m4-1 and x3650m4-2.

Using the GUI – Failover Cluster Manager

From one of the systems, open “Server Manager”

On the “Tools” menu, select “Failover Cluster Manager”

In “Failover Cluster Manager”, click “Create Cluster...”

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

Page 22

Click “Next” on the “Create a Cluster Wizard”

Add the names of the servers to be members of the cluster (x3650m4-1 and x3650m4-2). Click “Next”

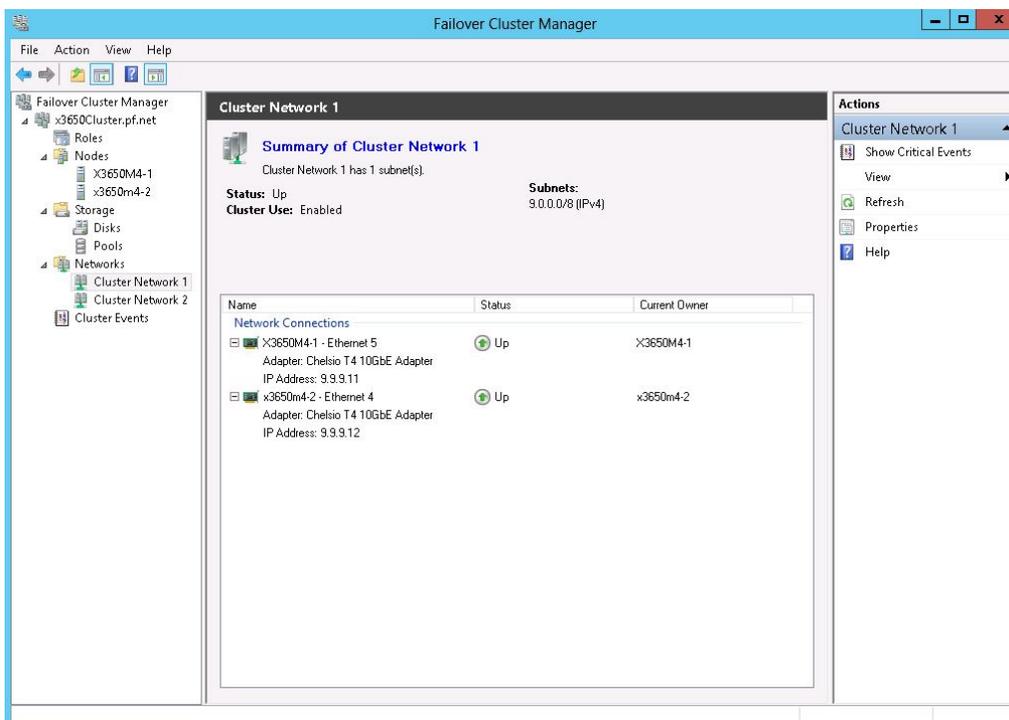
Enter the name of the cluster “x3650m4Cluster”, click “Next”

On the Confirmation screen, Click “Next”

Click “Finish”

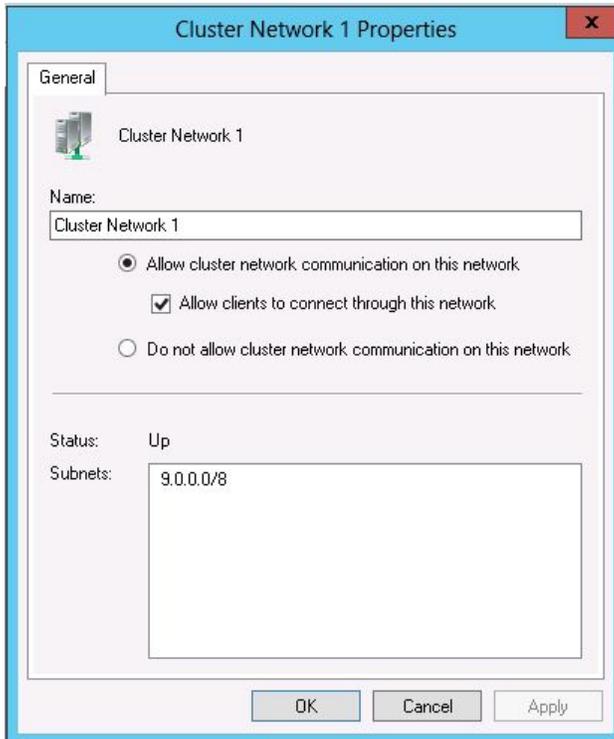
Configure the Cluster Networks

During the creation of the cluster, active network adapters in the individual nodes are assigned to the cluster networks. In our lab example, the Chelsio T4 10 GigE adapters are assigned to Cluster Network 1 and the onboard 1 GigE adapters are assigned to Cluster Network 2. These cluster network names can be changed to match the configuration names in your environment.

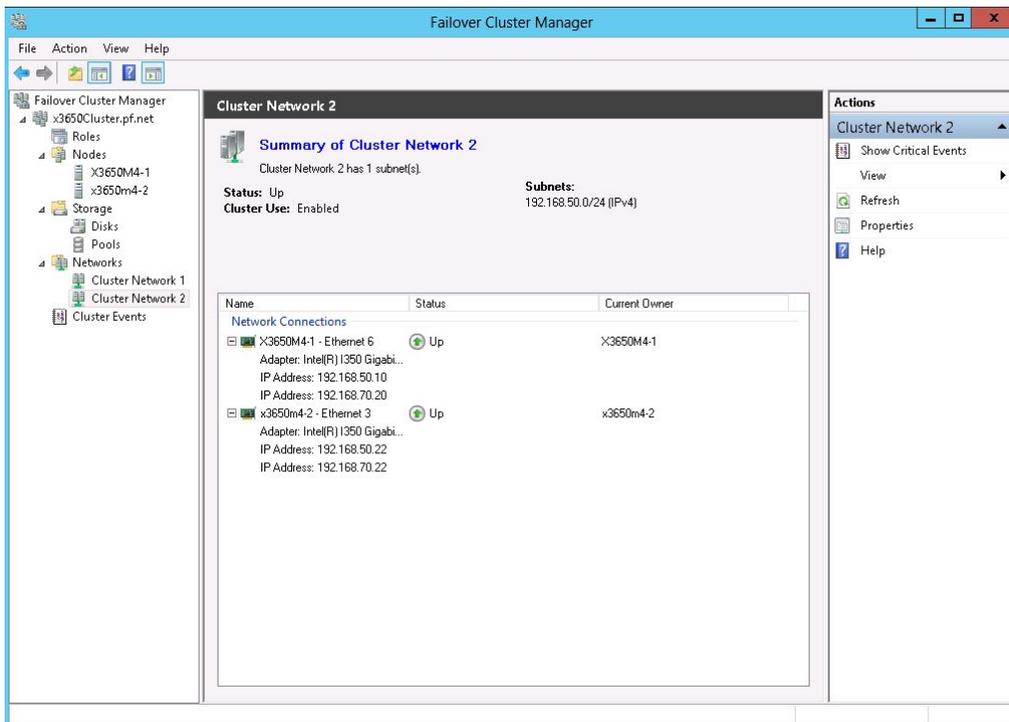


Right click “Cluster Network 1”

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers



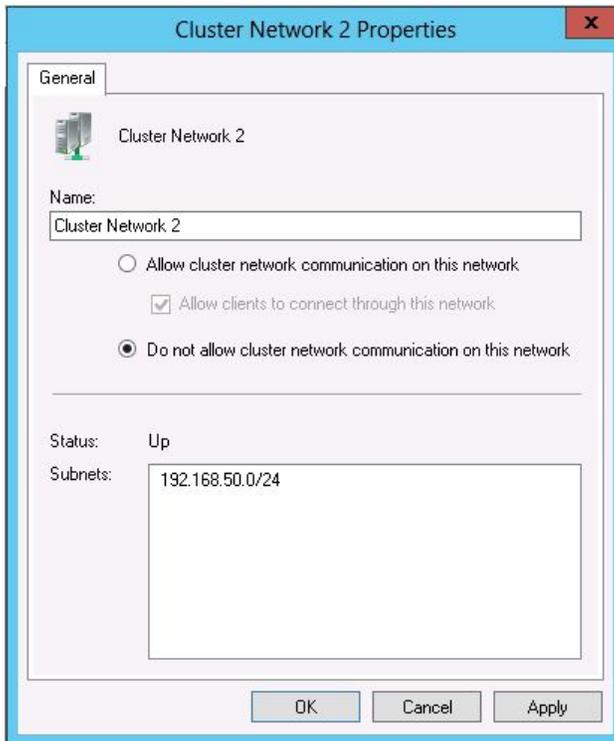
Select "Allow cluster network communication on this network"
Check the box "Allow clients to connect through this network"
Click "OK"



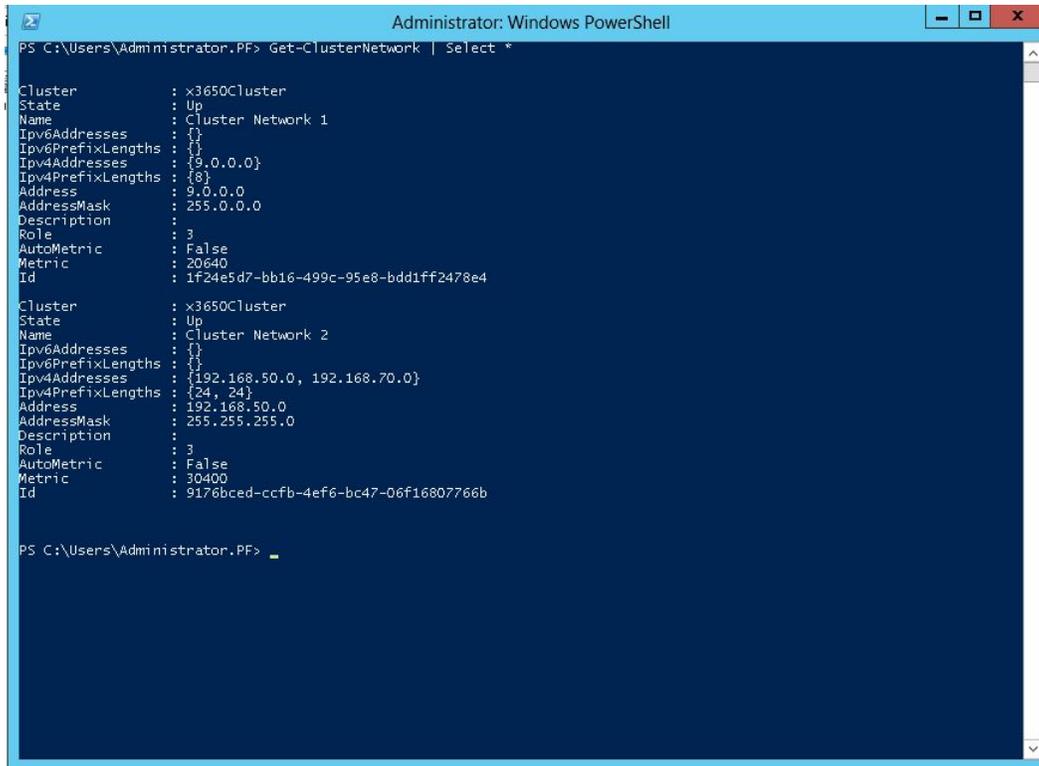
Right click "Cluster Network 2"

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

Page 24



Select "Do not allow cluster network communication on this network"
Click "OK"



```
Administrator: Windows PowerShell
PS C:\Users\Administrator.PF> Get-ClusterNetwork | Select *

Cluster      : x3650Cluster
State        : Up
Name         : Cluster Network 1
IPv6Addresses : {}
IPv6PrefixLengths : {}
IPv4Addresses : {9.0.0.0}
IPv4PrefixLengths : {8}
Address      : 9.0.0.0
AddressMask  : 255.0.0.0
Description  :
Role        : 3
AutoMetric  : False
Metric      : 20640
Id          : 1f24e5d7-bb16-499c-95e8-bdd1ff2478e4

Cluster      : x3650Cluster
State        : Up
Name         : Cluster Network 2
IPv6Addresses : {}
IPv6PrefixLengths : {}
IPv4Addresses : {192.168.50.0, 192.168.70.0}
IPv4PrefixLengths : {24, 24}
Address      : 192.168.50.0
AddressMask  : 255.255.255.0
Description  :
Role        : 3
AutoMetric  : False
Metric      : 30400
Id          : 9176bced-ccfb-4ef6-bc47-06f16807766b

PS C:\Users\Administrator.PF>
```

From the PowerShell window, type “Get-ClusterNetwork | Select *” displays the contents in the above screen.

Create the Folders and shares

In this section, steps to create shares for database and log files will be shown. Since there are two SQL servers, two data shares (Data1 and Data2) and two log shares (Log1 and Log2) will be created.

Using PowerShell Commands

From one of the cluster nodes (SQLSVR1)

```
Type “md c:\ClusterStorage\Volume1\Shares\Data1”
Type “New-SmbShare –name Data1 –Path c:\ClusterStorage\Volume1\Data1 –FullAccess Pf.net\Administrator, Pf.net\SQLService, NT AUTHORITY\SYSTEM”
Type “(Get-SmbShare Data1).PresetPathAcl | Set-Acl”
```

Repeat the same steps for SQLSVR2, change Data2 for Data1.

From one of the cluster nodes (SQLSVR1)

```
Type “md c:\ClusterStorage\Volume1\Shares\Log1”
Type “New-SmbShare –name Log1 –Path c:\ClusterStorage\Volume1\Log1 –FullAccess Pf.net\Administrator, Pf.net\SQLService, NT AUTHORITY\SYSTEM”
Type “(Get-SmbShare Log1).PresetPathAcl | Set-Acl”
```

Repeat the same steps for SQLSVR2, change Log2 for Log1

Using the GUI, Failover Cluster Manager

Select the “Roles” node on the left pane

Click on x3650-FS role and then click on “Add File Share” on the actions menu on the right

The “New Share Wizard” will start

On the “Select Profile” page, select “SMB Share – Applications” and click “Next”

On the “Share Location” page, select “c:\ClusterStorage\Volume1\” as the location. Click “Next”

On the “Share Name” page, type “Data1” as the share name, click “Next”

On the “Other Settings” page, click “Next”

On the “Permissions” page, click on “Customize permissions...”

Ensure user “NT AUTHORITY\SYSTEM” has “Full Control” permission

Click on “Add”, then click on “Select a Principal”

Enter “pf.net\administrator”, click on “Check Names”, click “OK”

Click “Full Control” and click “OK”

Click on “Add”, then click on “Select a Principal”

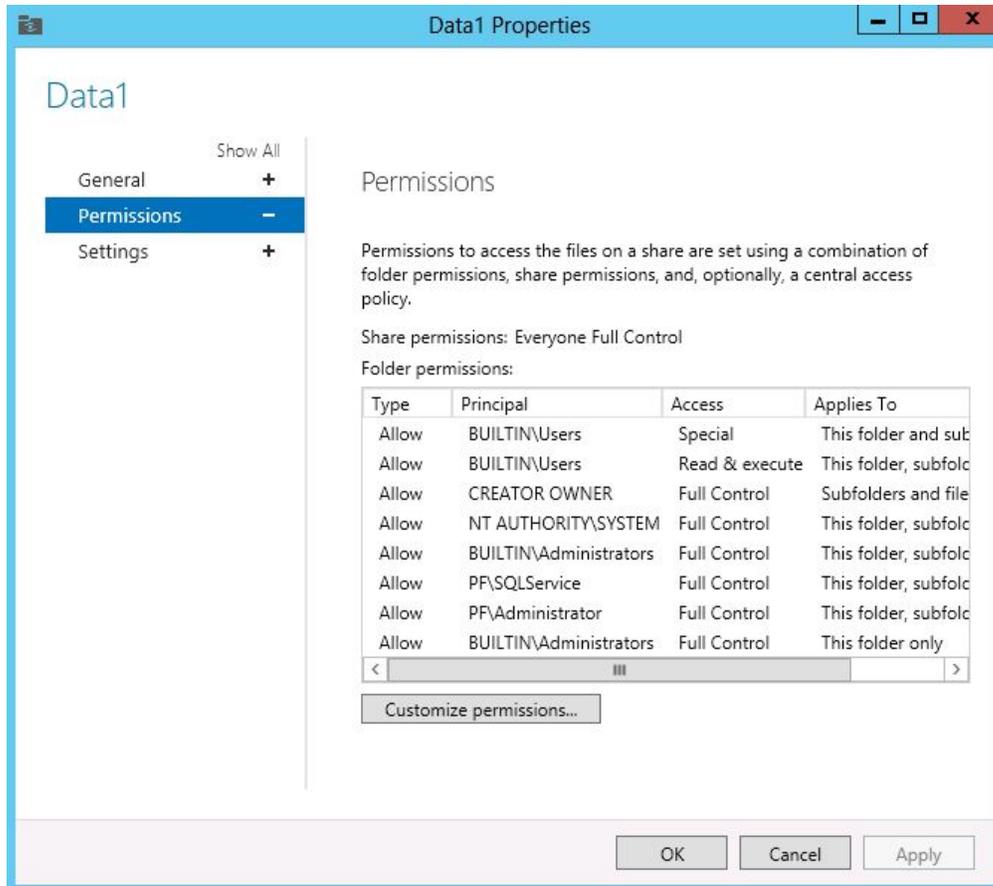
Enter “pf.net\sqlservice”, click on “Check Names”, click “OK”

Click “Full Control” and click “OK”

Click “OK”, then click “Next”, then click “Create”

Click “Close” after the wizard finishes creating the share.

Repeat the same steps for shares “Data2, Log1 and Log2”



Ensure “Administrator”, “SQLService” and “NT AUTHORITY\SYSTEM” users have “Full Control” privileges

Install SQL Server

SQL Server 2012 DataCenter Edition was installed on two servers SQLSVR1 and SQLSVR2. These two servers will access their database and log files via the file server cluster.

Run SQL Server 2012 Setup from the DVD

In the “SQL Server Installation Center”, click “Installation”, then select “New SQL Server stand-alone...”

SQL Server Setup Support Rule runs its verification tests. Ensure all test are passed. Click “OK”

Select the “Enter the Product Key” button. Enter the appropriate key in the space below

Click “Next” on the Licensing page

“SQL Server 2012 Setup” will start. Setup Support Rules will be verified. When passed, click “Next”

In the “Feature selection” page, select only the “Database Engine Services” and the “Management Tools”

Running Microsoft SQL Server 2012 on a Scale-Out File Server Cluster via SMB Direct Connection – Solution Utilizing IBM System x Servers

Page 28

Use the default directories. Click “Next”

In the “Installation rules” page, click “Next”

In the “Instance Configuration” page, click “Next”

In the “Disk Space Requirements” page, click “Next”

In the “Server Configuration” page, enter “pf.net\sqlservice” as the account name for the SQL Server Database Engine and the SQL Server Agent, set bot services to start automatically. Click “Next”

In the “Database Engine Configuration” page, click “Add Current User”

Click the “Data Directories” tab. Enter \\x3650-fs\data1 as the “Data Root Directory”

Change the two path for Log directories to use \\x3650-fs\log

Click “Yes” when prompted to confirm the correct permissions

Click “Next” on the “Error reporting” page

Click “Next” on the “Installation Configuration Rules” page

Click “Install” on the “Ready to Install” page

Installation of SQL will complete



For More Information

IBM System x Servers
IBM System x and BladeCenter Power Configurator

<http://ibm.com/systems/x>
<http://ibm.com/systems/bladecenter/resources/powerconfig.html>
<http://ibm.com/systems/x/hardware/configtools.html>
<http://ibm.com/systems/info/x86servers/serverproven/compat/us>
<http://ibm.com/server/support>
<http://ibm.com/systems/support>

IBM ServerProven Program
Technical Support
Other Technical Support Resources

Legal Information

© IBM Corporation 2012
IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
November 2012
All rights reserved.

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven® or ClusterProven®. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, ibm.com, System Storage, and System x are trademarks of IBM Corporation in the United States and/or other countries. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. For a list of additional IBM trademarks, please see <http://ibm.com/legal/copytrade.shtml>.

Intel, the Intel logo, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows Server, the Windows logo, Hyper-V, and SQL Server are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS

OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.

Intel® Virtualization Technology may require a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. For more information, visit <http://www.intel.com/go/virtualization>