



EMC Solutions for Microsoft SQL Server EMC Celerra Unified Storage Platforms

Reference Architecture



EMC NAS Product Validation
Corporate Headquarters
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

Published May, 2009

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Microsoft SQL Server EMC Celerra Unified Storage Platforms

Reference Architecture

P/N H2970.4

Contents

Chapter 1	Solution Overview.....	7
	Executive summary.....	8
	Introduction.....	8
	Audience.....	8
	Business benefits.....	9
	Technology challenges.....	10
	System configuration.....	10
	Storage architecture.....	11
	Network architecture.....	12
	Virtualization layer.....	13
	Application architecture.....	13
	High availability and failover.....	14
	Conclusion.....	15
	Related documents.....	15
Chapter 2	Storage Sizing and Scaling.....	17
	Introduction.....	18
	Reference configuration.....	18
	Workload.....	18
	Server configuration.....	18
	Storage configuration.....	19
	Applying the reference.....	20
	Workload evaluation.....	21
	Server memory sizing.....	21
	Storage sizing.....	21
	Conclusions.....	21
Chapter 3	Virtualization and Consolidation Options.....	23
	Introduction.....	24
	Virtualization.....	24
	VMware Infrastructure 3.....	24
	Microsoft Windows Server 2008 with Hyper-V.....	25
	Additional virtualization options.....	26
	Consolidation.....	26
	Multiple virtual servers per physical server.....	26
	Multiple instances per server.....	27
	Multiple databases per instance.....	28
	Conclusions.....	28
Chapter 4	Backup and Restore Options.....	29
	Introduction.....	30
	SQL native backup.....	30

	Additional backup options.....	30
	EMC NetWorker.....	31
	Third-party applications.....	31
	Conclusions	32
Chapter 5	Data Protection and Replication.....	33
	Introduction	34
	Testing, development, and reporting	34
	Array-based options.....	34
	Server-based options.....	35
	Disaster Recovery.....	35
	Array-based options.....	35
	Server-based options.....	36
	Conclusions	37

Figure 1	System overview	10
Figure 2	Storage layout	11
Figure 3	EMC Celerra NS-120 blade ports	13
Figure 4	Disk layout.....	20
Figure 5	VMware Hypervisor	24
Figure 6	Microsoft Windows Server 2008 with Hyper-V	25
Figure 7	Multiple virtual machines per physical server	27
Figure 8	Multiple instances per server	27
Figure 9	Multiple databases per instance	28
Figure 10	SQL native backup configuration	30
Figure 11	EMC NetWorker.....	31
Figure 12	Testing and development	34
Figure 13	Celerra Disaster Recovery	36
Figure 14	Microsoft SQL Server Database Mirroring.....	37

Chapter 1 Solution Overview

This chapter presents these topics:

Executive summary	8
Introduction	8
Business benefits	9
Technology challenges	10
System configuration.....	10
Conclusion.....	15
Related documents	15

Executive summary

Microsoft SQL Server is an environment that serves other applications. This simple fact substantially increases the complexity in planning for a SQL Server deployment. In the vast majority of SQL Server installations, the Database Server, or SQL Server system is accessed by users through another application. This application can be a commercially available tool, or something that was custom-designed for the needs of a particular business. However, in both cases an understanding of the application is a critical component in understanding the system as a whole.

Even with the vast differences that can exist between two deployments of SQL Server that depend upon the applications accessing the database, there are also common elements that will apply to most environments. Regardless of the application, the database still needs a place to store information and preserve backups of critical data. The database must be protected from disaster, and replicas of the data are required to periodically test or develop improvements to the system. This reference architecture (RA) provides proven solutions for each of those use cases that enhance the performance, reliability, and resiliency of the database system while allowing system consolidation and enhancing manageability.

Introduction

When you consider a SQL Server implementation, it is important to understand that the SQL Server itself has at most a small impact on the load imparted to any one component. Instead, it is the design of the data in the database, the structure of the relationships between that data, and the user applications that access the data that play a greater role in driving the system. This RA presents solutions for common use cases faced in database applications and uses a TPC-C like benchmark workload designed to simulate an Online Transaction Processing (OLTP) system. While this workload is designed to be similar to many real-world systems, it is not a real-world workload and is used as an arbitrary load to show that the solutions presented are effective with a load in general. The solution presented in this RA is broadly applicable across a wide variety of workloads. However, before you deploy any solution, test it appropriately to ensure that it performs as expected in your environment, and for your workload. The information as to which solutions are best for a specific environment and workload are beyond the scope of this document.

Audience

This document is intended for internal EMC personnel, EMC partners, and customers.

Business benefits

With limited resources and increasing demands, today's enterprises must address the following challenges:

- ◆ Consolidate multiple database applications scattered throughout the enterprise
- ◆ Ensure information access, availability, and continuity
- ◆ Maximize server and storage utilization, and deliver optimal system performance
- ◆ Manage upgrades and migrations
- ◆ Reduce the demands on limited IT resources and budgets
- ◆ Reduce the complexity of technology choices

In addition, these enterprises must manage IT costs and reduce the risk of business disruption.

This solution addresses each of these challenges using tested and proven solutions.

[Table 1](#) provides the details about the benefits of the solution.

Table 1 **Solution advantages**

Benefits	Details
Maintain service levels	This solution ensures that critical and revenue-generating Microsoft applications are available and running at peak performance. The redundancy built in to this solution ensures that there is minimal interruption in the event of a disaster.
Reduce costs	This solution reduces the cost of deploying and managing a SQL database environment by: <ul style="list-style-type: none"> • Providing sizing and implementation guidelines to enable rapid deployment • Providing best practices realized from real-world deployments • Optionally leveraging virtualization to increase resource utilization
Reduce risk	This solution reduces the risk of deploying and managing a SQL database environment by: <ul style="list-style-type: none"> • Providing a reference architecture that includes proven configurations that improve performance and scalability • Providing detailed step-by-step instructions for configuring a production database environment
Accelerate implementations	EMC Professional Services and ASN-certified EMC partners provide a rapid assessment and efficient implementation.

Technology challenges

Database systems are environments and not applications by themselves. As such there are many factors to consider when you deploy solutions into an existing environment or plan a new application deployment. EMC Proven Solutions allow you to start from a known reference configuration, and then customize it for your requirements by examining the similarities and differences between different methods of accomplishing common tasks in a database environment. This includes options for:

- ◆ Server virtualization
- ◆ Data backup and recovery
- ◆ Database replication and protection
- ◆ Application testing and development

System configuration

Figure 1 outlines the configuration of the tested solution. There are several permutations that are still considered part of the proven solution, and these are explained in the later sections.

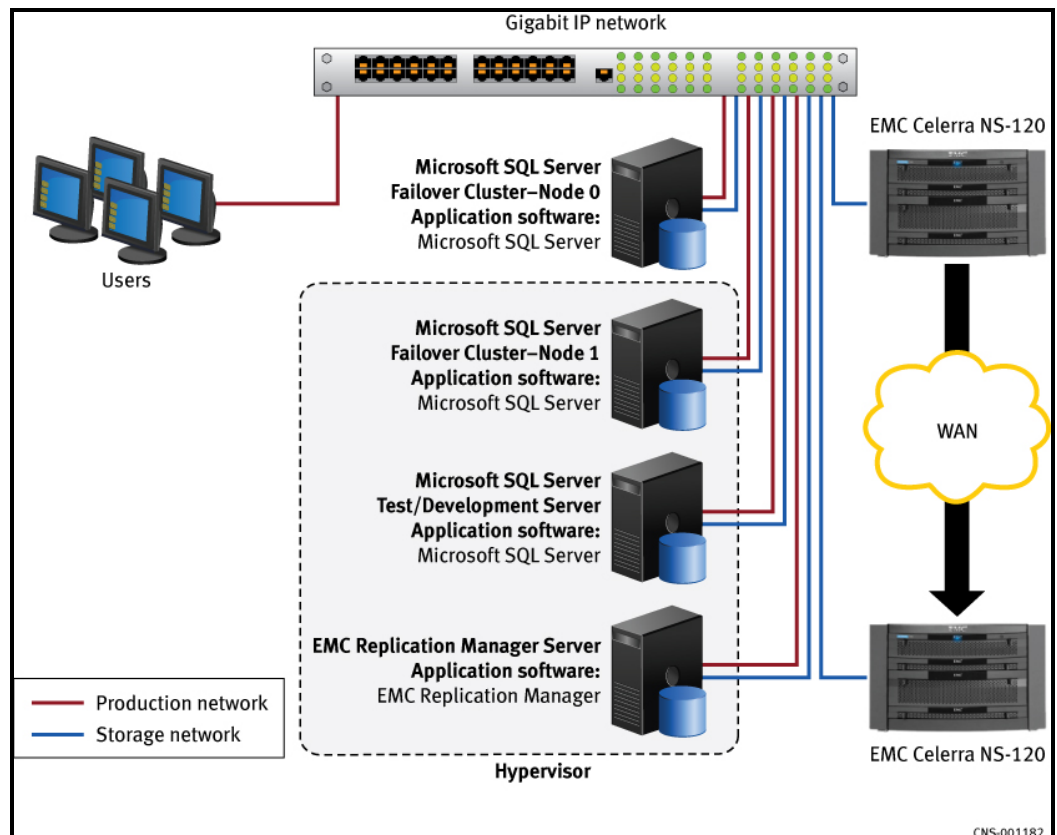


Figure 1 System overview

Storage architecture

Figure 2 outlines how storage is provisioned in one possible configuration of the validated solution.

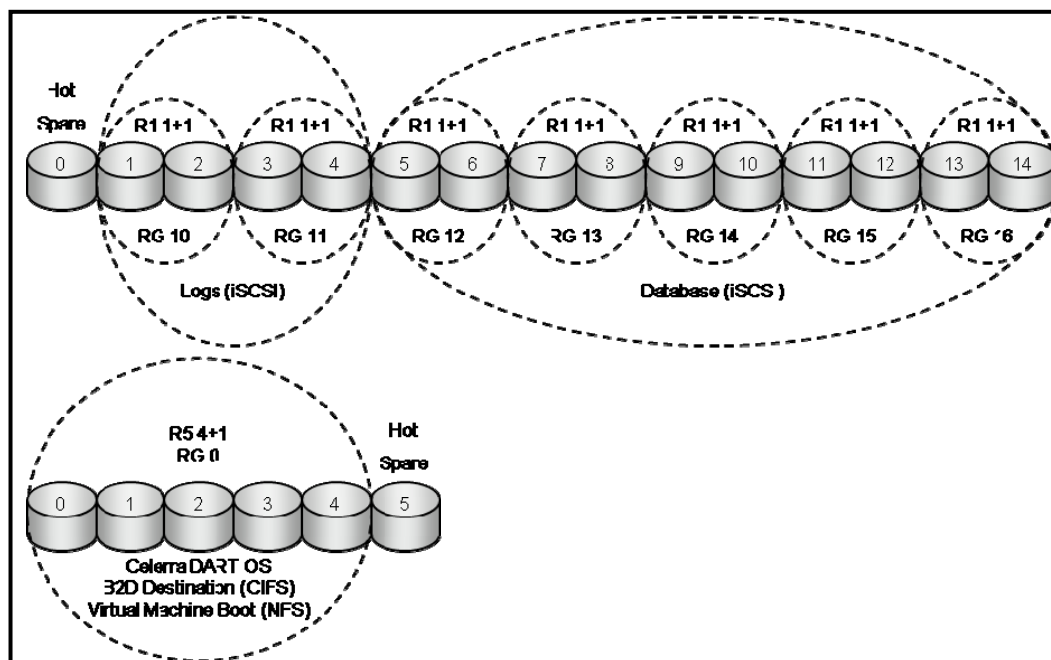


Figure 2 Storage layout

The validated solution can use storage through the Network File System (NFS), Common Internet File System (CIFS), and Internet SCSI (iSCSI) protocols. The NFS protocol is used to provide storage for the system disks of various virtual machines that are used in the solution. The CIFS protocol provides access to the area used for database backups. The iSCSI protocol is used to provide data storage to the primary database application. The underlying architecture of each storage area will be based on the anticipated workload that is applied to that area. These areas can have vastly different I/O performance, space, and protection requirements and should be maintained separately. For information on specific considerations for database area, refer to the [Storage Sizing and Scaling](#) chapter.

NFS Area – Virtual machine boot disks (if applicable)

The validated solution uses virtual machines in several roles, including a potential virtual SQL Server node on the primary database cluster. These virtual machines can be implemented as either physical or virtual servers with no impact on the database portion of the solution. If virtual machines are required, refer to several implementation options that are outlined in the [Virtualization and Consolidation Options](#) chapter.

If this solution is implemented with virtual machines, it is possible to use networked storage for the virtual machine boot disks. Therefore, the array-based methods such as snapshots, replication, and NDMP backups can be implemented to protect the virtual machines. It is also possible to protect these virtual machines using server-based techniques. The method of protection is not specified by the solution.

CIFS Area – Database backups

Database backups are critical to the administration of production applications. The CIFS area in the Celerra® provides storage space to accommodate full, differential, and transaction log backups as required.

iSCSI Area – Database and Log files

The validated solution allows both physical and virtual servers in the primary database server role. To accommodate a range of potential implementations, and to allow several array-based functions, the validated solution specifies that the primary application database should reside on Celerra iSCSI storage. This area must be sized according to the requirements of your specific database application. For information on sizing and scaling for specific considerations, refer to the [Storage Sizing and Scaling](#) chapter.

The validated solution also calls for a failover cluster. If this is implemented using a disk-based quorum, the quorum drive should reside on iSCSI storage.

Note: In a fully virtual environment, it is possible to host the database and log files by using an NFS datastore. This configuration is not part of this solution. The configuration is explained in *EMC Solutions for Microsoft SQL Server 2005 on VMware ESX Server EMC Celerra NS20 over NFS — Reference Architecture* available on Powerlink®.

Network architecture

System-wide network design and architecture are outside the scope of this document and solution. This section presents recommendations for proper functionality that are consistent with industry accepted best practices and should be compatible with the existing network infrastructure and policies.

Switches

In any highly available solution, there should be multiple network paths between each component in the system so that no single failure can disrupt communication. The solution architecture diagram (refer to [Figure 1](#) on page 10) defines a set of logical networks that are recommended for the solution. The networks do not provide failover protection. Therefore, the networks should be designed so that they are highly available on their own.

The validated solution uses Gigabit Ethernet (GbE) switches.

Virtual local area networks

When creating highly available networks, the requirement to have multiple switches on multiple networks can quickly become costly. Virtual local area networks (VLANs) allow each physical network switch to maintain multiple virtual networks that are logically isolated from each other. This solution was validated using VLANs within a highly available network and can be implemented with logical separation or with physical separation by using distinct physical switches.

EMC Celerra Unified Storage Platforms

EMC® Celerra storage arrays contain at least two Data Movers, which can operate independently. They can also operate in the active/passive mode, with the passive blade serving as a failover device for the active Data Mover. In this solution, the Data Movers operate in active/passive mode.

EMC Celerra arrays can have a minimum of four network ports. [Figure 3](#) shows the ports on the rear of an EMC Celerra NS-120 blade.

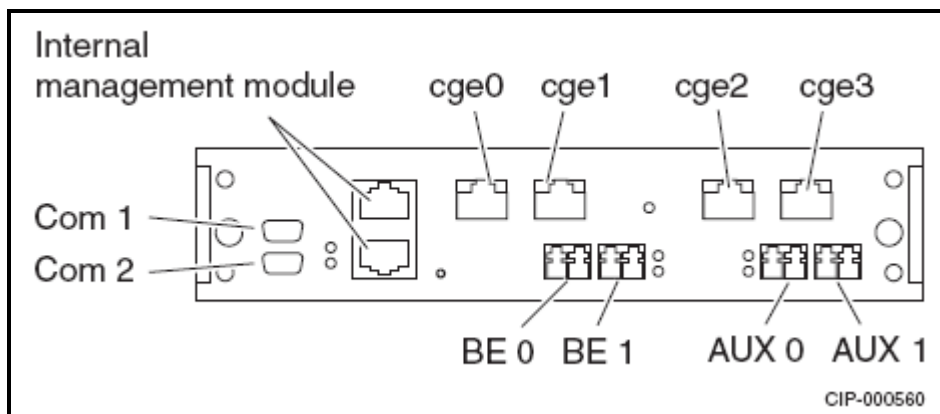


Figure 3 EMC Celerra NS-120 blade ports

Ports cge0 and cge1 handle the storage traffic. Ports cge2 and cge3 are left open for future requirements.

The Data Mover supports several types of link aggregation for Internet Protocol (IP) traffic. However, in this configuration, no link aggregations or Ethernet channels were configured, because Celerra also supports iSCSI Multiple Connections per Session (MC/S), which Microsoft recommends.

Note: As a best practice, the blade network ports connected to the storage network should be dedicated to storage traffic. However, if the ports are not heavily used, they can be shared with non-storage network traffic. EMC recommends monitoring the network to avoid bottlenecks.

Virtualization layer

The validated solution uses multiple virtual components to save space in the data center and also to reduce power and cooling costs. The solution is validated for both physical and virtual implementations as described in the [Virtualization and Consolidation Options](#) chapter.

As with all of the other layers in the solution, stack care should be taken to ensure that the impact of any virtualization technology on the performance, availability, and recoverability of the system is well understood before implementing it in a production environment.

Application architecture

The solution assumes a single database on a single database server instance that is running on a failover cluster. This simplifies the discussions around sizing and scaling for a workload. In a multi-database or multi-instance environment, each individual contributor to the workload should be considered and planned independently so that predictable performance can be obtained. As a best practice, all database workloads should be isolated from other workloads so that storage contention does not cause performance degradation.

The Microsoft SQL Server binary files are installed on the application servers. Database files, database log files, tempdb files, and other system database files reside on an EMC Celerra storage array in the iSCSI area as described earlier in this section.

The file system and logical unit numbers (LUNs) for the system database files, cluster quorum, and Microsoft Distributed Transaction Coordinator (MSDTC) files also reside in the Celerra iSCSI area. However, they should be physically isolated from the database to prevent contention.

The OLTP database is divided into four database files (based on the best practice guideline of one database file per processor). Each database file resides on a separate LUN. All of the LUNs reside on the same Celerra file system.

The database log files reside on a single LUN. The LUN resides on its own Celerra file system, and is physically isolated from the spindles used for the database files.

For the validated solution the system temporary database (tempdb) resides on a single iSCSI LUN on the same spindles as the user database. The test workload that is used for the validated solution has a minimal impact on tempdb. In real-world applications this usage pattern can be drastically different. It is a best practice to evaluate tempdb as a completely separate database and size it according to its measured workload.

High availability and failover

The validated solution provides protection at the storage layer, the connectivity layer, and the host layer.

Storage layer

Celerra can have multiple Data Movers to provide high availability and load balancing. In the Microsoft SQL Server EMC Celerra solution, primary and standby blades provide seamless failover capabilities for the Celerra storage. This minimizes the user disruption during routine Celerra maintenance such as upgrading the DART OS.

The redundant array of independent disks (RAID) disk configuration on the Celerra back-end provides protection against hard disk failures.

Connectivity layer

Using the Microsoft iSCSI Initiator software, the Celerra platform supports Multiple Connections per Session (MC/S). This provides high availability, failover, and load balancing. MC/S supports multiple TCP/IP connections from the initiator to the target within the same iSCSI session. If a connection fails, another connection can continue processing I/O without interrupting the application. In the validated solution, MC/S also provides “round-robin” load-balancing capabilities.

Note: There are other multipath technologies available with the Microsoft iSCSI Initiator Software. However, as a best practice Microsoft recommends using MC/S if it is supported by both the initiator and the target.

The solution configuration also includes separate network interface cards (NICs) at the source of each I/O path, a separate network infrastructure (cables, switches, routers, and so on), and separate target ports.

Host layer

The application hosts have redundant power supplies and network connections to reduce the impact of host hardware failure. In the validated solution, the application host is clustered using Microsoft Cluster Services, providing further redundancy.

Conclusion

This document provides a reference architecture for solutions based on Microsoft SQL Server with EMC Celerra. While several technologies, products, and considerations are presented as working together to create a system that can meet performance, reliability, availability, and recovery targets, a customer environment can have different requirements. For this reason, in each of the subsequent chapters, the reference configuration for specific functionality is discussed along with other options that can be used.

Microsoft SQL Server is a highly configurable and variable application environment, and every customer can have a different set of availability needs, performance goals, and recovery targets. This document attempts to provide a framework to implement these solutions that addresses the inherent complexity of many environments, while retaining the ease-of-use that characterizes EMC Celerra.

Related documents

The following documents, located on Powerlink.com, provide additional, relevant information. Access to these documents is based on your login credentials. If you do not have access to the following documents, contact your EMC representative:

- ◆ *EMC Solutions for Microsoft SQL Server on VMware ESX Server EMC Celerra NS20 over NFS – Reference Architecture*
- ◆ *EMC Solutions for Microsoft SQL Server EMC Celerra Unified Storage Platforms – Applied Best Practices Guide*

Chapter 2 Storage Sizing and Scaling

This chapter presents these topics:

Introduction	18
Reference configuration	18
Applying the reference	20

Introduction

Discussions about sizing and scaling of systems are critical to the deployment of this solution or any other solution for database environments. Sizing systems for database workloads is a complex task that deserves a great deal of thought and planning to ensure that a system behaves in a manner that is acceptable to the various stakeholders involved in a project.

Reference configuration

This section outlines a configuration that has been validated with the solutions modules presented in this RA. This is an example of a system that can be implemented, but it is not the answer to all customer issues. Database solutions can vary drastically based on individual customer requirements. So each customer environment must be evaluated, analyzed, and sized individually. The configuration presented in this chapter is used to have a common point of reference so that the interactions between the solution components can be explored.

Workload

Real-world database installations can vary widely between customers or even between individual instances within a customer environment. This is because each database workload is different and is determined by the user applications accessing the database, the patterns of users making such requests, and the individual database design. All these factors make it impossible to strictly define a mapping between hardware and a number of “users” that can be supported in a customer environment without a detailed analysis of that environment.

Note: Unless otherwise noted, the testing for this solution uses an OLTP workload like the TPC-C benchmark test.

Server configuration

The validated solution uses a Microsoft Windows failover cluster to host the primary database instance. This cluster contains one node on a traditional physical system, and the second node as a virtual machine that is runs as a guest on a leading hypervisor. This allows the physical machine hosting the passive node of the failover cluster to also run other virtual machines to support testing and development operations and any other work that is required to support the solution. In a real-world implementation this system can host the passive node of multiple clusters simultaneously, if desired. The implications of virtualization in this space are detailed in the [Virtualization and Consolidation Options](#) chapter.

Both nodes of the failover cluster are configured with the same amount of memory and multiple processor cores. The physical node has eight processing cores in four physical processors, while the virtual node has four virtual CPUs (at the time of testing this was the maximum number of virtual cores supported).

Table 2 lists the hardware used for this solution.

Table 2 Server hardware

Name	Quantity	Description
Server 4U	2	SQL Servers (Production Cluster, Test/Dev)
Server 1U	4	Utility Servers (Load Generation, Replication Manager)
Gigabit Ethernet Switch	1	Network Switch
Storage	1	EMC Celerra NS-120 containing at least two Data Movers

Table 3 lists the software used for this solution.

Table 3 Server software

Name	Quantity	Description
Microsoft Windows 2003 Enterprise edition	4	SQL Server cluster (2) Test/Dev environment Replication Manager (RM) Server
VMware ESX Version 3.5	1	Hypervisor for Production Cluster Node 1, Test/Dev server and RM server
Microsoft SQL Server 2005 Enterprise edition	2	Production cluster Test/Dev environment
EMC Replication Manager agent	3	Production cluster (2) Test/Dev environment (1)
EMC Replication Manager server	1	RM server
EMC Celerra DART	1	DART version 5.6 and later

Storage configuration

The validated solution uses 14-disk spindles to host the database and transaction logs for the primary application with an additional five-disk spindle hosting a location for backup files. The database and log disks are set up as mirrors (RAID 1), and then striped together (RAID 10) to allow higher performance than the performance provided by a single pair. The backup area uses RAID 5 to provide protection.

Figure 4 shows the configurations used.

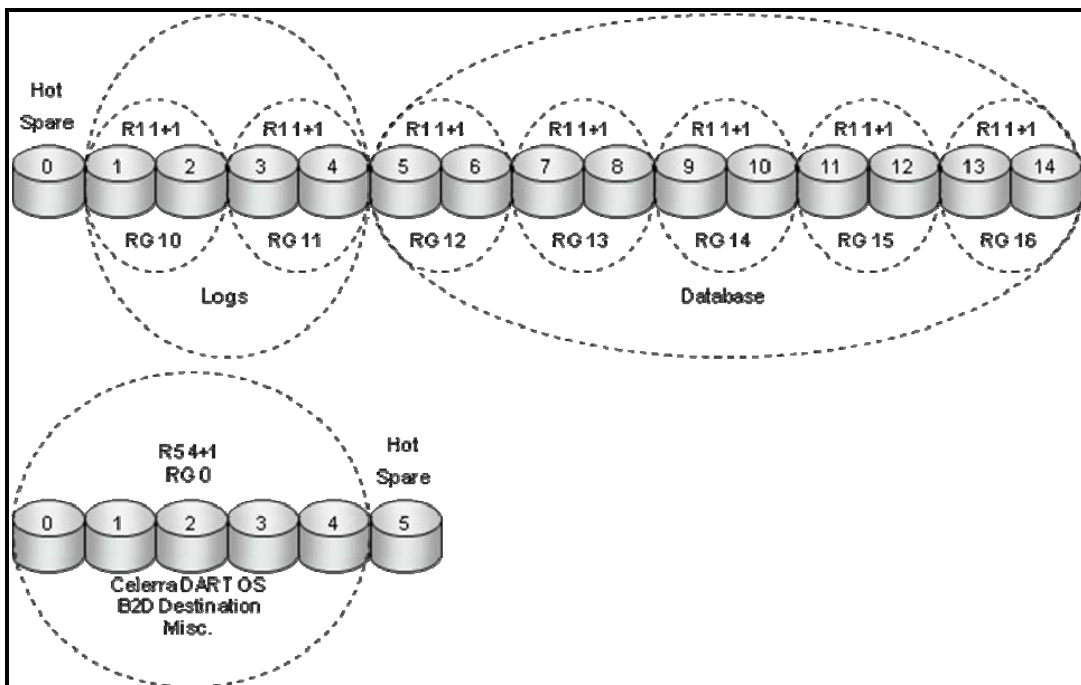


Figure 4 Disk layout

As with all of the other aspects of the system, it is important to understand how the storage architecture supports high availability and load balancing across various components.

The EMC Celerra platform is inherently highly available and has many architectural features to enable the system to continue running even when failures do occur. These include redundant power, I/O pathways, and network connections.

This is enhanced for databases by separating the data and log files onto dedicated physical spindles and using RAID 1 protection.

For a complete description of how this configuration is performed during validation testing, refer to *EMC Solutions for Microsoft SQL Server Data Storage for EMC Celerra — Validation Test Report*, available on Powerlink. This is an internal-only document available to EMC employees and partners on Powerlink.

Applying the reference

The reference configuration provides a place to start when you size a deployment in a real environment. By understanding how the workload relates to the test workload, and how changes to the test system impact that workload, you can gain some understanding of how changes should impact the real-world environment. This is not intended to be a guarantee that the changes and results indicated in this chapter can translate exactly into another environment, but rather as a starting point so that decisions that impact customers can be made with the best information available.

Workload evaluation

The first step in any sizing exercise is to evaluate or plan the workload that will be handled by the system in question. It is important to understand how the workload translates into the storage environment so that you can determine the I/O performance capacity. Based on the environment, there are numerous ways to accomplish this task.

Note: It is also important to evaluate the load at a granular level appropriate to the application. SQL Server provides a range of flexibility options to place the objects on the disk. These options may not be appropriate in all cases, but they should be evaluated.

Server memory sizing

In database environments, memory and I/O are intrinsically linked. Having a larger amount of memory in the buffer pool allows database pages to stay in memory longer and allows write I/Os to be coalesced more efficiently. Both these aspects change the I/O characteristics observed from the storage system. In general, as memory increases, I/O requests become larger and more sequential in nature. The reverse is also true. As memory is consumed by other processes and the buffer pool shrinks the database, I/O operations tend to become smaller and more random from the perspective of the storage system. When evaluating any storage changes it is important to understand how server memory influences the I/O behavior.

Storage sizing

There are many options to consider when you size database storage. Disks have two primary metrics of capacity and both must be considered in designing a storage system. They are storage capacity, typically measured in gigabytes (GB) and performance capacity, typically measured as Input/Output operations per second (IOPS). It is easy to determine the storage capacity in most environments. However, the performance capacity is often overlooked because it is difficult to discern the capacity.

The first item to consider when examining a disk is the performance capacity. After that number is known, it is a simple matter to evaluate storage capacity. However, going in the reverse order is not as intuitive. For more information, refer to *EMC Solutions for Microsoft SQL Server EMC Celerra Unified Storage Platforms — Applied Best Practices Guide* available on Powerlink.

Conclusions

You can use the reference configuration as a starting point to evaluate the data storage options available in your environment. This configuration is intended as a reference only and should not be substituted for proper evaluation and testing of the actual environment and workload.

Chapter 3 Virtualization and Consolidation Options

This chapter presents these topics:

Introduction	24
Virtualization.....	24
Consolidation	26
Conclusions.....	28

Introduction

An increasing number of companies are exploring ways to virtualize and consolidate SQL Server environments to reduce costs, increase flexibility, and use resources more efficiently. This section describes how various virtualization and consolidation options can be used with the solution.

Virtualization

Virtualization technology has the ability to separate the operating system from the underlying hardware by using software called a hypervisor. The operating system still behaves like it is running on physical hardware, but in reality it is interacting with a set of virtual hardware provided by the hypervisor. There are many advantages to such an approach, which are outlined in the materials available from the leading hypervisor vendors. The validated solution uses virtualization to allow a physical system to host more than one logical system. This is intended as a proof of concept, and not as a requirement. There are no functions inherent in the solution that requires virtualization. The proven solution has been tested with the two leading Hypervisor products VMware Infrastructure 3 and Microsoft Windows Server 2008 with Hyper-V.

Note: There are other methods for creating and running virtual machines on a physical system including an approach called hosted virtualization. The proven solution has only been validated with VMware ESX and Microsoft Hyper-V hypervisor-based virtualization methods. Other methods are not part of the validated solution.

VMware Infrastructure 3

VMware Infrastructure 3 allows guest operating systems to connect to storage using several methods. This reference architecture specifies that database storage should be provided to the server by using the iSCSI protocol without translation at the Hypervisor level. The remaining connection methods explained in this section may support some or all the modules in this solution.

This section summarizes the options available for each connection method. For more information, refer to *EMC Solutions for Microsoft SQL Server Virtualization for EMC Celerra — Validation Test Report*, available on Powerlink. This is an internal-only document available to EMC employees and partners on Powerlink.

Figure 5 shows the configuration tested for VMware Infrastructure 3.

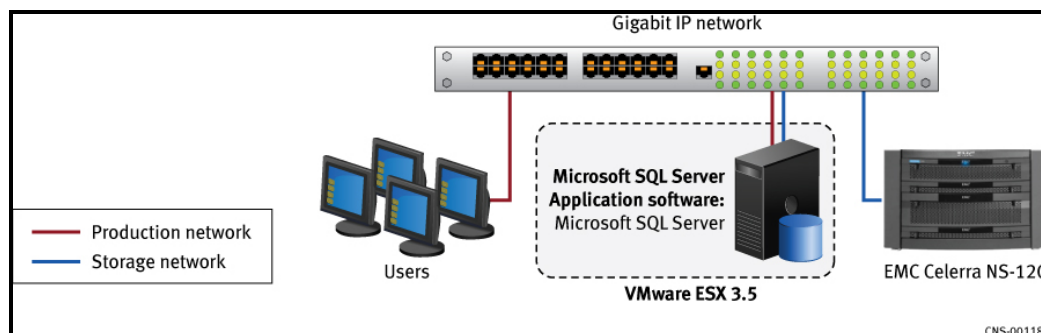


Figure 5 VMware Hypervisor

Guest OS Software iSCSI Initiator (MSI)

A Windows guest that runs within the server that hosts VMware ESX can connect to Celerra iSCSI storage using the Microsoft iSCSI initiator just like a physical system. By using this connection method, all the modules in this solution are supported.

Software or hardware initiator from VMware ESX (VMDK on iSCSI VMFS)

Using a Virtual Machine File System (VMFS) datastore drastically changes the semantics of interacting with storage. This solution does not explicitly cover this scenario. However, there are several modules that can be applicable to this scenario that are not validated.

Software or hardware initiator from VMware ESX (Raw Device Mapping)

This connection method is very similar to the Guest OS initiator method, and many of the same solution modules can be applied. This solution does not explicitly support Raw Device Mapping (RDM) devices, but can be used as a reference when designing such a system.

NAS datastore (VMDK on NFS)

The EMC Celerra platform can become a NAS datastore for a server that hosts VMware ESX using the NFS protocol. This solution is covered in the document *EMC Solutions for Microsoft SQL Server 2005 on VMware ESX Server EMC Celerra NS20 over NFS — Reference Architecture*, available on Powerlink.

Microsoft Windows Server 2008 with Hyper-V

In the Hyper-V environment, there are several methods available to connect storage to a virtual guest operating system. The differences in those methods impact the supportability of various solution modules that are being used in that mode. For more information, refer to *EMC Solutions for Microsoft SQL Server Virtualization for EMC Celerra — Validation Test Report*, available on Powerlink. This is an internal-only document available to EMC employees and partners on Powerlink.

Figure 6 shows the configuration tested for Microsoft Windows 2008 with Hyper-V.

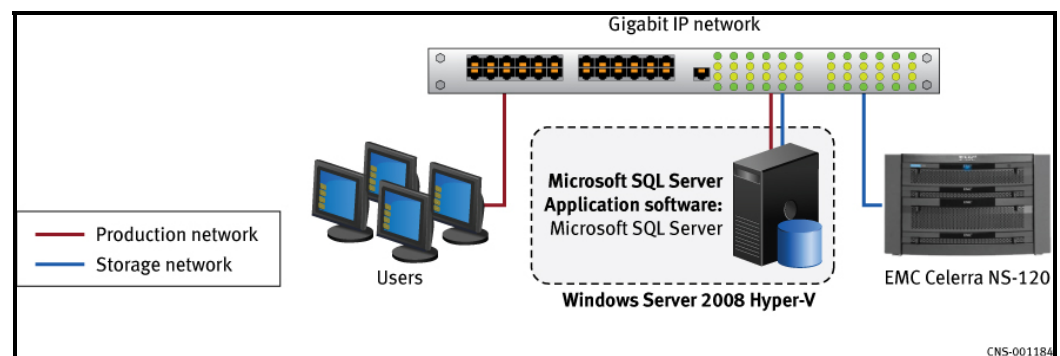


Figure 6 Microsoft Windows Server 2008 with Hyper-V

Guest OS Software iSCSI Initiator (MSI)

A Windows guest that runs within Hyper-V can connect to Celerra iSCSI storage by using the Microsoft iSCSI initiator just like a physical system. By using this connection method, all the modules in this solution are supported.

Software or hardware initiator from the host OS (VHD on iSCSI NTFS)

Using an NTFS datastore introduces significant changes into how the guest OS interacts with storage. This method is analogous to the VMware VMFS scenario, which is not covered in this reference architecture.

Software or hardware initiator from the host OS (Pass-Through)

This connection method is similar to the guest OS software initiator, and the VMware RDM methods. Many solution modules may be applicable using this method, but they have not been validated as part of the solution.

NAS datastore (VHD on NFS or CIFS)

At this time, Microsoft Windows 2008 with Hyper-V allows CIFS and NFS file system shares to serve as datastores for virtual machines; however, this configuration is not supported. This connection method was not validated as part of this solution, and specific solution documentation for such a configuration has not yet been published by EMC at this time. It is likely that such a solution would look similar to the VMware solution for Celerra NFS, which is described in the document *EMC Solutions for Microsoft SQL Server 2005 on VMware ESX Server EMC Celerra NS20 over NFS — Reference Architecture*, available on Powerlink.

Additional virtualization options

Additional virtualization options exist, but have not been tested as part of this solution. Many of the concepts, practices, and recommendations that apply to this solution will be applicable in those systems. However, they have not been tested and are not supported as part of this solution.

Consolidation

Consolidation refers to the idea that many database servers are underutilized, and that by combining workloads an organization can realize improved utilization, reduced costs, and greater management efficiency. There are several ways to implement consolidation by using Microsoft SQL Server. In addition to the frequently discussed virtualization options, there are two additional often-overlooked features of SQL Server that can be used for consolidation. These features are explained in this section.

Multiple virtual servers per physical server

A common method to consolidate SQL Server installations is by using virtualization technology. This allows you to combine the advantages of consolidation with the advanced features and benefits that you get from virtualization.

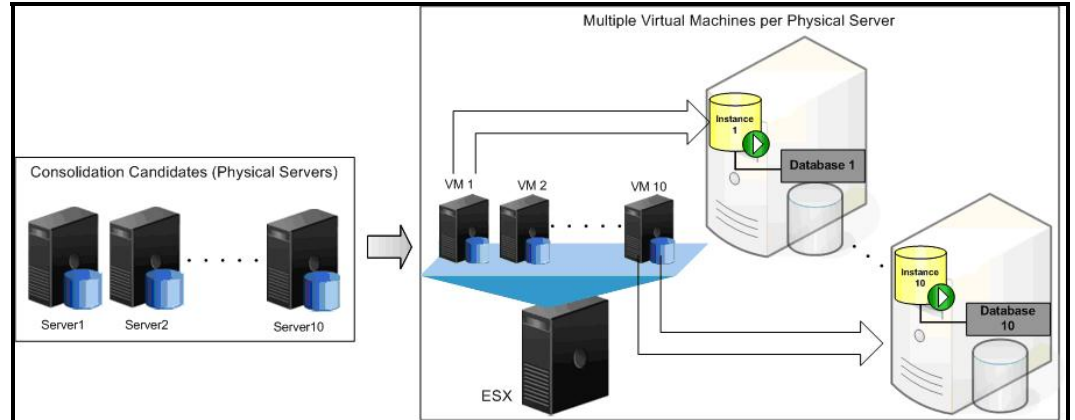


Figure 7 Multiple virtual machines per physical server

In general, by consolidating with virtualization you need not make many changes to your production applications and the entire consolidation project can be confined to the data center. However, virtualization may not be the best fit in all cases. It is often the first option considered due to the recent interest in the industry on the topic of virtualization, but it should not be the only option that you consider.

Multiple instances per server

It is an often overlooked fact that when you are managing a Microsoft SQL Server machine by using the management tools that Microsoft provides, you are not connecting to the physical server, but to a SQL Server instance on the physical server. If you do not specify an instance, the tool connects to the default instance on the physical machine. One easy way to consolidate servers is to create multiple instances on the same underlying hardware. These instances can be managed separately, and can even run different versions of SQL Server.

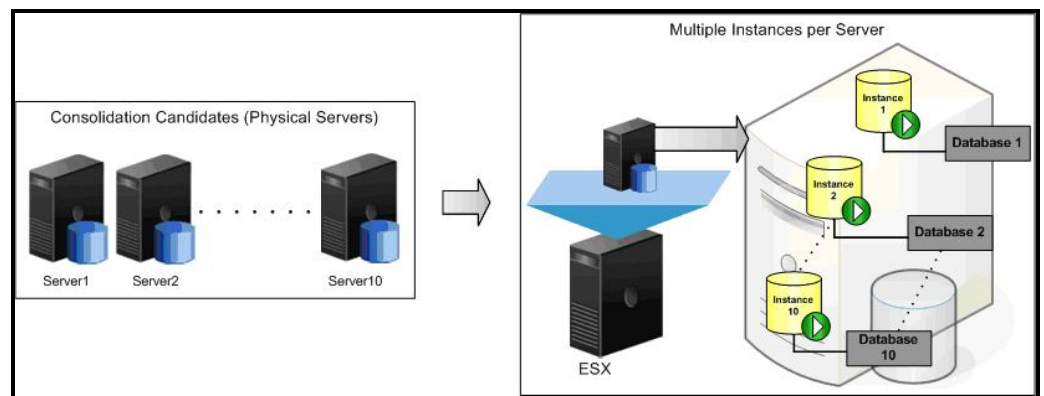


Figure 8 Multiple instances per server

Additional information about creating, managing, and maintaining instances is available from Microsoft SQL Server Books Online.

Multiple databases per instance

It is common to find database server machines in an environment that hosts only one database. If this one database has not been fully utilizing the computing resources available, the database server may be a good entity for consolidation.




Figure 9 Multiple databases per instance

If there are several application databases that share management, security, maintenance, or other requirements, it is a good idea to consider hosting multiple databases on a single instance of SQL Server. This has numerous advantages in terms of reducing management complexity and administrative overhead for these databases.

Conclusions

Consolidation provides opportunities to reduce costs, increase flexibility, and more efficiently use resources. Consolidation can be implemented using SQL Server features, or by introducing virtualization into the solution. Many of the components in this solution can be virtualized depending on your specific needs. However, there are no components in this solution that explicitly require virtualization to be considered as a part of the validated solution.

Chapter 4 Backup and Restore Options

This chapter presents these topics:

Introduction	30
SQL native backup	30
Additional backup options.....	30
Conclusions	32

Introduction

Data backups are an essential part of any production environment. Regardless of the RAID protection level, hardware redundancy, and other high-availability features present in EMC Celerra storage arrays, conditions exist where you may need to be able to recover a database to a previous point in time.

SQL native backup

Figure 10 shows the SQL native backup configuration.

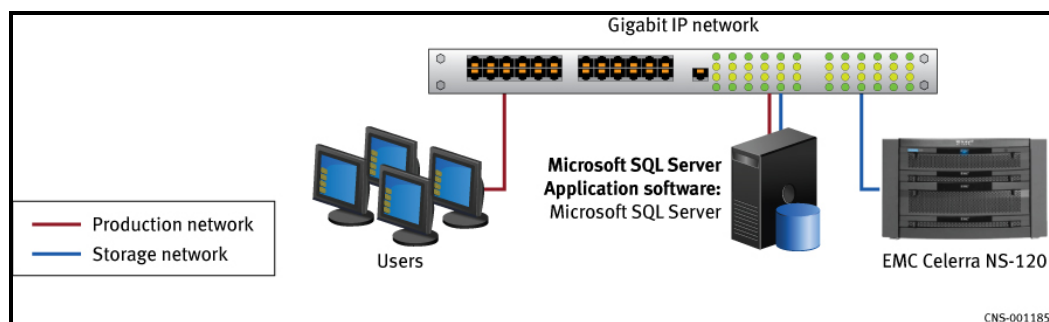


Figure 10 SQL native backup configuration

As a no-cost option that is included in every SQL Server installation, the native backup functionality is identified as the preferred backup method for the validated solution. It provides full, differential, and transaction log backups to local disk and network file share locations, which can be scheduled like any other database maintenance job using the SQL Server agent.

The native backup tool is managed on a per-server basis using the same interface as other SQL Management Operations. The disadvantage is that as additional servers are added to the environment, they must continue to be managed independently. For larger environments, one of the other backup options may be more suitable.

Additional backup options

In addition to the SQL native backup tool that is included in the validated solution, other backup tools exist that may be more suitable for your needs. Any of the options listed in this section can be exchanged for the SQL native backup module in the solution without impacting other solution modules.

EMC NetWorker

Figure 11 shows an EMC NetWorker® configuration.

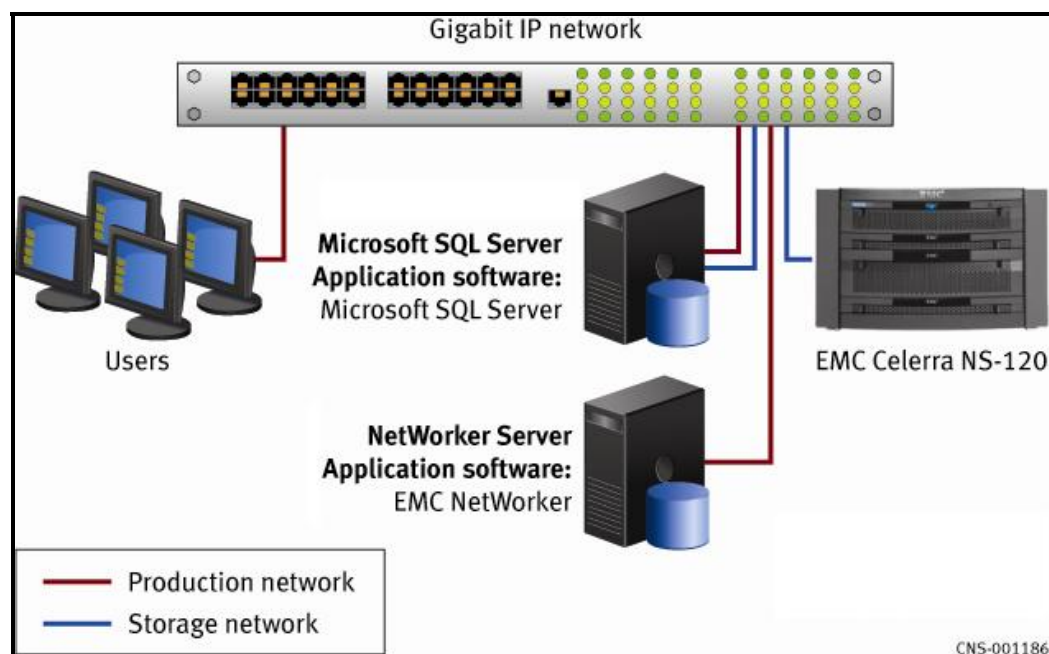


Figure 11 EMC NetWorker

The EMC NetWorker product allows you to manage all your backups from various systems with one interface. The modular architecture allows you to incrementally extend functionality to various applications as they are introduced to your environment.

EMC NetWorker Module for SQL Server

The SQL Server module allows database-consistent backups to minimize the time spent in recovery in the event that you need to restore from a backup. The module integrates with EMC Celerra Replicator™ to create an application-consistent snapshot of the iSCSI LUNs behind the database and uses those as the source for the backup operation. This allows the production database to have only a very small impact in performance due to the checkpoint operation.

EMC NetWorker Module for Microsoft applications

The Microsoft Applications module provides a common interface for managing backups of SQL Server, Exchange, and SharePoint. The EMC NetWorker for Microsoft Applications module is independent of the SQL Server module. This module also integrates with EMC Celerra Replicator to minimize the performance impact of the backup operation.

Third-party applications

In addition to the offerings from EMC, there are other backup offerings from other companies. For a complete list of the software packages that have been tested, refer to *EMC Solutions for Microsoft SQL Server Data Backup for EMC Celerra — Validation Test Report* available on Powerlink. This is an internal-only document available to EMC employees and partners on Powerlink.

Conclusions

A thoughtful and complete backup strategy is an essential part of database maintenance in a production environment. The validated solution uses SQL native backup functionality to ensure recoverability by using the included, no-cost tool. The native backup tool is suitable for many types of customer environments, but is not appropriate in all cases. Some environments may have needs that are not addressed in the software. For those environments, several additional options are presented that can replace a specific solution module without impacting the other modules that are required in your environment.

Chapter 5 Data Protection and Replication

This chapter presents these topics:

Introduction	34
Testing, development, and reporting	34
Disaster Recovery	35
Conclusions	37

Introduction

Data protection and replication involve creating online copies of data that can be used for functions other than the primary production application. These functions are dependent on the applications in use, but can be broadly categorized as testing, development, reporting, and Disaster Recovery. These are not absolute distinctions but are useful to categorize replicas and discern differences in replication technologies.

Testing, development, and reporting

At some point during an application lifecycle testing, developing changes to an application becomes incompatible with the production workload on the application. At this point a dedicated testing and development system becomes a necessity. In such case, business reports about the data in the production application can introduce a load, which cannot be supported on the production system without adversely impacting service level agreements, which introduces a need for a dedicated reporting system.

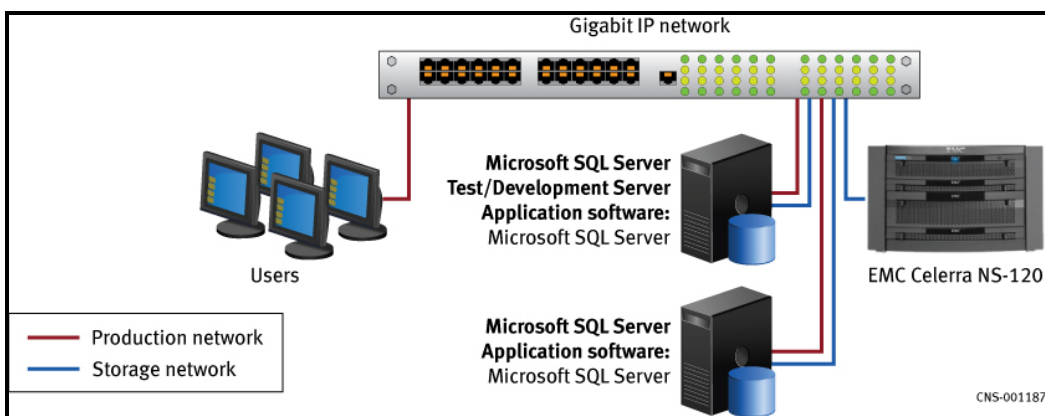


Figure 12 Testing and development

In both cases, even though we need a new database server, we need it to have a copy of the production database, and may need to have this copy updated periodically.

Array-based options

A primary benefit of array-based replication is that the processing and I/O overhead associated with creating and maintaining the replica is contained within the storage system. This eliminates any impact to your server hardware that can be associated with the operation. There are two types of array-based replicas, snapshots and clones which are discussed in this section.

Snapshot replicas

A snapshot is a logical, point-in-time view of a source storage object whose state is maintained by keeping track of the changes between the current production object and the snapshot object. These changes can be tracked in various ways, which are not important in this context. The important point is that snapshots depend on their original source, so any load against the snapshot object can cause contention with the original object in some way. If the original object is corrupted, the snapshot will be corrupted as well.

The advantage of snapshots is that they do not require as much storage space as a full copy. An additional space is required for changed data only.

Clone replicas

Clones are complete physical copies of data that require as much space for the replica as the original object. The advantage of clones is that I/O requests to the clone do not cause I/O requests to the original object and therefore there is no contention between them.

Note: It is possible to place a clone on the same physical spindles as the source, in which case there would be no logical contention between the two objects, but there would be physical contention against the underlying disk spindles.

EMC Replication Manager

EMC Celerra Replicator for iSCSI, when used with EMC Replication Manager, allows you to create application-consistent snapshot and clone replicas with a single interface for SQL Server, Exchange, and file systems on both Celerra and CLARiiON® arrays.

Server-based options

SQL Server has the ability to create a database snapshot within the database server and present this as a separate database for use. This does not require additional agent software as with the Replication Manager case, but does require additional processing overhead on the database server.

Disaster Recovery

In the [Testing, development, and reporting](#) section, all the potential use cases can be satisfied with data on a single Celerra system. Disaster Recovery implies some amount of physical distance between copies of data so that an event at one site will not impact the copy of data at another site. The distance between these sites and the amount of data that can be lost is a primary consideration when discussing Disaster Recovery options.

[Figure 13](#) on [page 36](#) shows the architectural diagram of the setup of Celerra Disaster Recovery.

Array-based options

In the same way that Celerra Replicator for iSCSI allows local clone replicas of a database using the EMC Replication Manager software, it can also provide remote replicas. These remote copies of data are updated asynchronously on a schedule set by the administrator and controlled from Replication Manager.

For more information, refer to the *EMC Replication Manager Administrator's Guide* available on Powerlink.

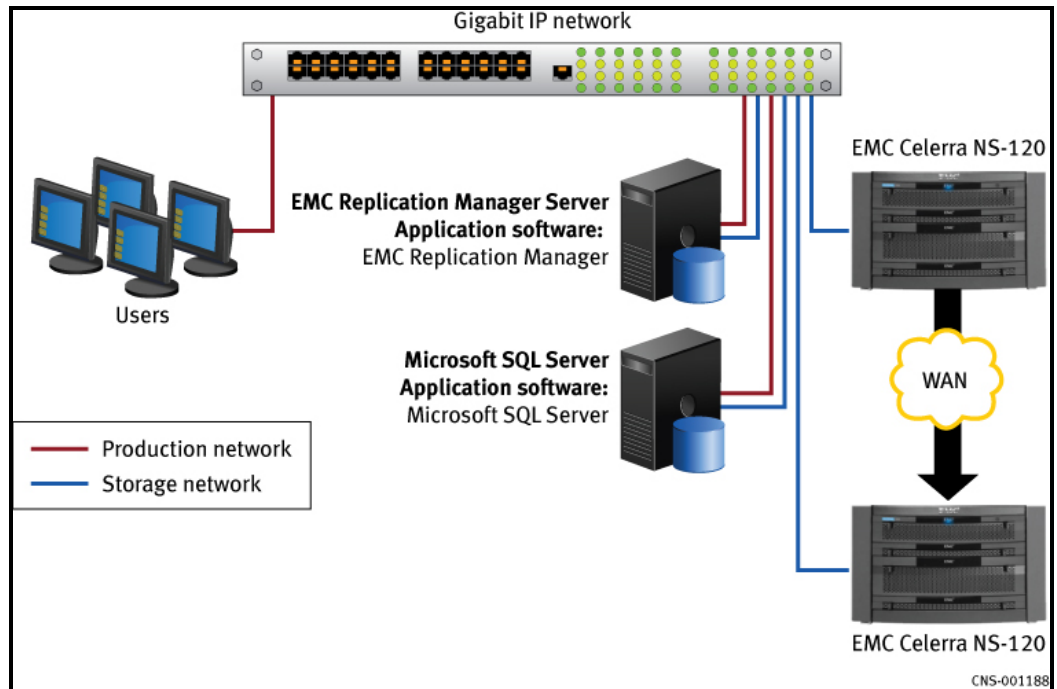


Figure 13 Celerra Disaster Recovery

By using asynchronous replication combined with the snapshot functionality provided by the Celerra platform, this protection method requires changed blocks of data to be sent across the network only once per update cycle, and does not require any additional server resources. The process also leverages Microsoft Virtual Device Interface (VDI) to allow no application downtime, and to allow inter-database consistency during replication cycles.

Server-based options

In addition to an array-based replication method, Microsoft SQL Server allows a function known as Database Mirroring between two database servers. Database Mirroring replicates the transaction log data to a remote database server, and then copies the log into the remote database. In this way the remote database is kept up to date with recent changes made on the source as shown in [Figure 14](#).

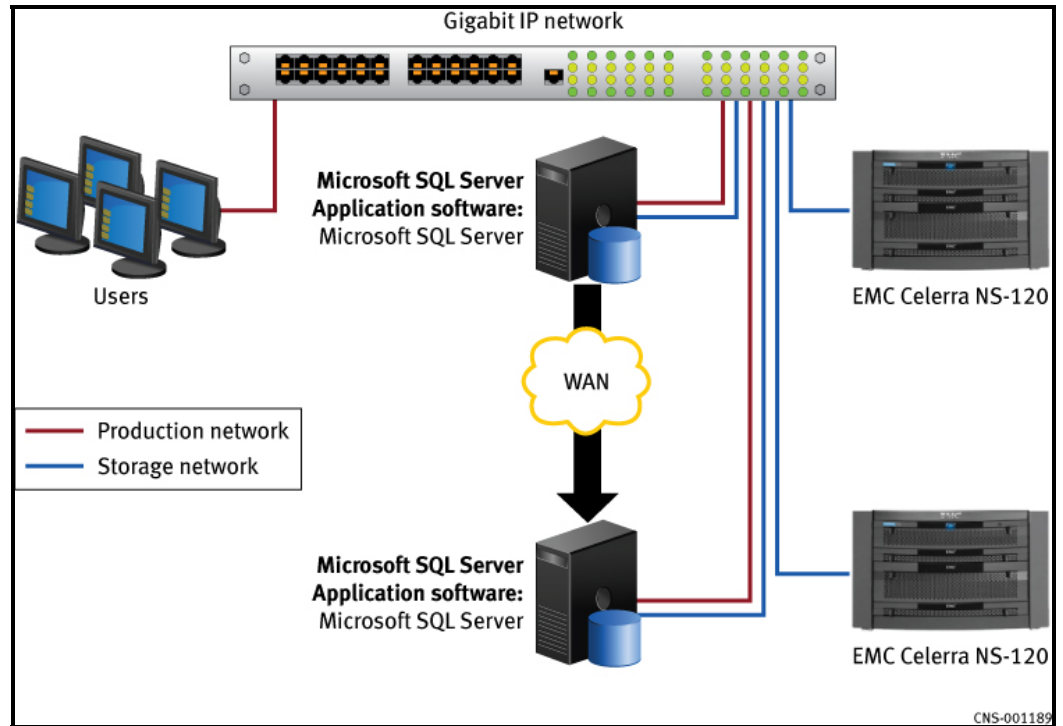


Figure 14 Microsoft SQL Server Database Mirroring

Database Mirroring supports both asynchronous and synchronous mirroring options, but does not support keeping two or more databases transactionally consistent with each other. The feature also requires server resources to manage the replication activity.

Conclusions

As database environments grow, they tend to increase both the requirements for daily operation and their criticality to the business. Even if testing, development, reporting, and disaster recovery may not be the requirements during an initial application deployment, it is prudent to understand how those functions can be added to a system in the future. This solution provides options to create local and remote replicas of application data that are suitable for testing, development, reporting, and disaster recovery and many other operations that can be important in your environment. Additional information on replication and protection technologies is detailed in *EMC Solutions for Microsoft SQL Server Data Protection and Replication for EMC Celerra — Validation Test Report*. This is an internal-only document available to EMC employees and partners on Powerlink.