

**EMC Backup and Recovery for
Microsoft SQL Server**

Enabled by EMC NetWorker Module for
Microsoft SQL Server

Reference Architecture

EMC NAS Product Validation



Copyright © 2010 EMC Corporation. All rights reserved.

Published February, 2010

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

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly.

EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. 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.

Part number: h6896

Table of Contents

Reference architecture overview.....	4
Solution architecture	6
Key components	9
Validated environment profile.....	11
Hardware and software resources	12
Conclusion.....	13

Reference architecture overview

Document purpose

EMC's commitment to consistently maintain and improve quality is led by the Total Customer Experience (TCE) program, which is driven by Six Sigma methodologies. As a result, EMC has built Customer Integration Labs in its Global Solutions Centers to reflect real-world deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges currently facing its customers.

This document describes the reference architecture of the EMC[®] Backup and Recovery for Microsoft SQL Server - Enabled by EMC NetWorker[®] Module for SQL Server solution. This was tested and validated by EMC Global Solutions.

Solution purpose

The purpose of this reference architecture is to build and demonstrate the functional, performance, and scalability aspects of the Backup and Recovery solution for Microsoft SQL Server using EMC NetWorker.

This reference architecture validates the performance of all aspects of the solution and provides guidelines for building similar solutions.

This reference architecture is not intended to be a comprehensive guide to every aspect of the EMC Backup and Recovery for Microsoft SQL Server – Enabled by EMC NetWorker Module for SQL Server solution.

The business challenge

Databases are an integral part of every business enterprise. From Customer Relationship Management and Human Resource systems, Payroll, Business Intelligence, web content and more, designing and building a database infrastructure that supports these functions and protects the data is a major challenge.

These systems are always on and accessed by users. At the same time, they store more and more data. In many cases, it is not acceptable to take the database offline for nightly backups. If a restore is required, it must be restored quickly with minimum impact to users.

The technology solution

It is important to work with an experienced provider of backup and recovery solutions — one with a strong track record and the ability to deliver. EMC has designed solutions for protecting Microsoft SQL Server 2008 database environments using EMC NetWorker and EMC Celerra[®].

This solution uses EMC NetWorker to take disk-based backups of the SQL Server database. The backup destination for this solution is a CIFS share created on an EMC Celerra unified storage platform.

The solution benefits

Simplifies the backup management: EMC NetWorker reduces the complexity of backup management by centralizing and scheduling the SQL database backups.


Server-initiated backup: EMC NetWorker with NetWorker Module for Microsoft SQL Server (NMSQL) eliminates the need for administrators to access the SQL server.

Improves storage and power efficiency: EMC Celerra deduplication helps in reducing the storage footprint for SQL backups. Low-power SATA drives help in reducing the power consumption to store backups.

Solution architecture

Architecture diagram

The following illustration depicts the overall physical architecture of the solution.



Reference architecture overview

The validated solution is built with a virtualized SQL Server 2008 environment on EMC Celerra unified storage.

The key components of the reference architecture are:

- Microsoft SQL Server 2008
- EMC Celerra unified storage
- NetWorker server


Microsoft SQL Server 2008 is installed on a database server and the NetWorker server is created on another machine. An EMC Celerra system is used to store the production database, log files, and database backups.

The different connectivity methods used in the solution are as follows:

- The SQL database drives are accessed by the SQL servers through the Fibre Channel (FC) protocol.
- The backup target is a file share accessed through the Common Internet File System (CIFS) protocol.

Storage layout

The following illustration depicts how the storage is provisioned in one possible configuration of the validated solution.



Storage layout overview

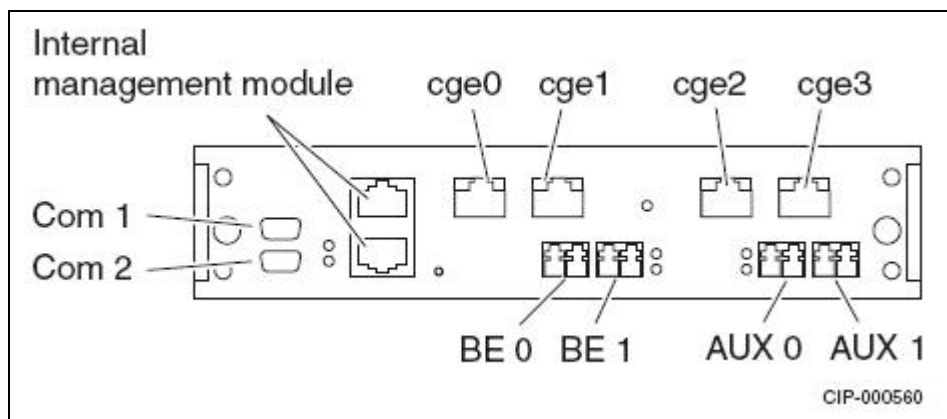
The validated solution uses storage through the FC, iSCSI, and CIFS protocols. The FC protocol is used to provide the storage for SQL database and log files. The area used for database backups can be accessed through the CIFS or iSCSI protocols.

To satisfy the performance requirements and to allow several array-based functions, the validated solution specifies that the primary database should reside on FC storage. The validated solution uses 14 spindles to host database and transaction logs. The protection level used is RAID 1/0 for high performance.

The CIFS area in Celerra provides space to accommodate full, differential, and transaction log backups as required. Five low-cost, low-power SATA drives with RAID 5 protection level are used for the CIFS area.

Network layout

The following illustration shows the ports on the rear of an EMC Celerra.



Network layout overview

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

EMC Celerra storage arrays contain at least two Data Movers that can operate independently. Each Data Mover can have a minimum of four Ethernet ports. For high-availability purposes, the Data Movers can be bound together as aggregated links or have multiple physical connections serving a single logical connection.

EMC Celerra comes with an integrated CLARiiON® storage system that has two storage processors (SPs). The front-end ports on the SPs may be connected to a SAN switch or directly connected to a host bus adapter (HBA) on a host or Data Mover. The validated solution uses a SAN switch for FC connectivity between the storage array, Data Movers, and servers. Port 0 and Port 1 on each SP are used for host or server connectivity. Port 2 and Port 3 are used to connect to the Data Movers.

Key components

Introduction

This section briefly describes the key components of this solution.

For details on all the components that make up the reference architecture, see [Hardware and software resources](#).

EMC Celerra unified storage platform

The EMC Celerra unified storage platform is a dedicated network server optimized for file and block access, delivering high-end features in a scalable, easy-to-use package. For maximum scalability, Celerra unified storage platforms leverage both the innovative EMC CLARiiON FC RAID storage, delivering best-in-class availability and data protection, and the availability, performance, and ease of management of EMC Celerra.

Celerra unified storage systems deliver a single-box block and file solution offering a centralized point of management for distributed environments. This makes it possible to dynamically grow, share, and cost-effectively manage multi-protocol file systems and provide multi-protocol block access. Administrators can take advantage of simultaneous support for NFS and CIFS protocols by enabling Windows and Linux/UNIX clients to share files using the Celerra's sophisticated file-locking mechanisms and by leveraging iSCSI or FC for high-bandwidth or latency-sensitive applications.

The usage of low-power SATA II drives in Celerra reduces the power consumption because they require 32 percent less energy per terabyte than traditional 1 TB SATA II drives. Also, Celerra deduplication compresses all inactive files and then single-instances them to remove duplicate data. Celerra data deduplication supports all Celerra functionality and provides an intuitive, single-click start-up with an automated policy.

EMC NetWorker

EMC NetWorker provides the ability to protect an enterprise against data loss. As the enterprise grows, so does the complexity and importance of protecting data. NetWorker software provides the power and flexibility to meet these challenges.

The NetWorker software is a cross-platform, client/server application that provides the ability to remotely manage all NetWorker clients and servers from a web-enabled, graphical interface. The NetWorker server backs up client data regularly by using scheduled backups. They are preferred over manual backups because the backups occur automatically and data can be recovered more easily.

The EMC NetWorker Module for Microsoft SQL Server is a NetWorker add-on module that provides application-consistent backup and restore for SQL database and transaction logs. The NetWorker software provides backup and restore capabilities for file system data only. However, a file system backup does not save SQL Server data in a recoverable form. With NetWorker Module for Microsoft SQL Server, the NetWorker software can back up and restore Microsoft SQL Server data on a daily basis.

The relationship between the NetWorker server, NetWorker Module, and other network components are shown in the following figure.

Validated environment profile

Profile characteristics The solution was validated with the following environment profile.

Profile characteristic	Value
SQL 2008 database size	207 GB
Instances and databases	Single instance and single database
Number of database files	Four files, each file on a different LUN
Workload	OLTP - based on TPC-C
Storage for SQL database	FC storage
Storage for SQL backups	CIFS file share
Production SQL 2008 databases RAID type, physical drive size, and speed	RAID 1/0, 300 GB SAS drives (15k rpm)
Backup area RAID type, physical drive size, and speed	RAID 5, 1 TB, SATA drives (7.2k rpm)

Hardware and software resources

Hardware The following table lists the hardware used to validate the solution.

Equipment	Quantity	Configuration	Purpose
Storage	1	EMC Celerra unified storage Two Data Movers 20 GB to 300 GB SAS drives (15k rpm) Six 1 TB SATA drives (7.2k rpm)	
Enterprise-class FC network switch	1	4 GB FC switch	
Enterprise-class IP network switch	1	Gigabit Ethernet switch	
HP ProLiant DL585 D5 server	1	4 quad-core 2.31 GHz AMD Opteron processors, 72 GB RAM, four 1 Gigabit Ethernet adapters, two FC HBAs	SQL Server
HP ProLiant DL385 D5 servers	2	2 quad-core 3 GHz Intel Xeon processors, 20 GB RAM	Load generation and NetWorker servers

Software The following table lists the software used to validate the solution.

Software	Version
VMware vSphere	4
Microsoft Windows Server	Windows 2008 x64 Enterprise Edition SP2 Windows 2003 x32 Enterprise Edition R2 SP2
Microsoft SQL Server	SQL Server Enterprise Edition 2008 SP1
EMC Celerra DART	5.6.46.4
EMC CLARiiON FLARE®	02.23.050.5.705
EMC NetWorker	7.5 SP1
EMC NetWorker Module for SQL Server	5.2 SP1

Conclusion

Summary

SQL Server backup and recovery are integral parts of an organization to ensure high availability of data and recovery in case of data corruption. This reference architecture depicts a validated design for Microsoft SQL Server using an EMC Celerra unified storage system and EMC NetWorker.

This solution reduces the complexity of backup management by centralizing and scheduling the SQL database backups using NMSQL. Also, EMC NetWorker with NMSQL eliminates the need for backup administrators to access the SQL server.

This solution uses EMC Celerra deduplication to reduce the storage footprint for SQL backups. Additionally, low-power SATA drives reduce the power consumption to store backups.

Next steps

EMC can help accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of a backup and recovery solution for a Microsoft SQL Server 2008 environment.

To learn more about this and other solutions, contact an EMC representative.
