



White Paper

SQL Server Consolidation with Server Virtualization on NetApp Storage

Generosa Litton and Richard Preston, NetApp
March 2010 | WP-7095

FLEXIBLE AND EFFICIENT STORAGE CONSOLIDATION FOR SQL SERVER USING VIRTUALIZATION

NetApp® solutions provide storage efficiencies that offer numerous advantages for Microsoft® SQL Server® database consolidation, especially when virtualization is a part of the architectural strategy. As a result, companies can create flexible, secure, cost-effective, and efficient storage environments that complement and enhance the efficiencies and management benefits of SQL Server databases.

TABLE OF CONTENTS

1	INTRODUCTION	3
2	CHALLENGES IN DATABASE ENVIRONMENTS	3
	OPTIONS FOR DATABASE STORAGE ARCHITECTURES	3
3	VIRTUALIZATION INCREASES DEMANDS ON STORAGE	4
4	THE NETAPP ADVANTAGE: STORAGE EFFICIENCY CAPABILITIES	5
5	FLEXIBLE IT ARCHITECTURE	5
6	COMPLETE DATA PROTECTION	7
	INSTANTANEOUS AND EFFICIENT BACKUP FOR SHARED STORAGE	7
7	MICROSOFT TOOLS AND NETAPP VIRTUALIZATION	8
8	EFFICIENT STORAGE	9
9	CONCLUSION	9
	REFERENCES AND MORE INFORMATION	9

1 INTRODUCTION

Consolidation of SQL Server database infrastructures results in lower cost of ownership through several mechanisms. A smaller footprint in the data center reduces operational costs for power, cooling, and hardware. Licensing benefits for virtualized servers result in lower capital expenses. In addition, reduction of SQL Server sprawl allows for greater IT management control.

Microsoft proposes three strategies to consolidate SQL Server deployments in its “SQL Server Consolidation Guidance” white paper (<http://msdn.microsoft.com/en-us/library/ee819082.aspx>):

- Multiple databases may be consolidated onto a single SQL Server instance.
- Multiple SQL Server instances may be consolidated onto a single server.
- Multiple servers running SQL Server instances can be consolidated onto a single physical server running in virtual machines (VMs) using server virtualization such as Windows® Server 2008 Hyper-V™.

Each of these approaches produces increased demands on the storage infrastructure, but none more so than server virtualization. Selecting the correct storage solution is critical to reaping the maximum benefits from consolidating SQL Server deployments. NetApp provides solutions that support all SQL Server consolidation strategies and further provides storage and data management that enables effective server virtualization environments, further extending the benefits of SQL Server consolidation.

This paper reviews NetApp solutions for Microsoft SQL Server consolidation using server virtualization.

2 CHALLENGES IN DATABASE ENVIRONMENTS

Database environments present a number of distinct challenges. It becomes very difficult to manage environments consisting of multiple databases, each with its own storage, especially as they grow in size and in number of instances. This can be particularly so in the case of SQL Server. With its ease of use and affordability, SQL Server can quickly sprawl throughout companies outside the control of central IT. As these database applications grow, they begin to challenge the ability of IT departments to manage them. In addition, providing data protection for these silos of SQL Server storage poses additional challenges, particularly as it relates to meeting application availability.

OPTIONS FOR DATABASE STORAGE ARCHITECTURES

There are multiple options for database storage architectures (see Figure 1). For example, direct-attached storage (DAS) is often encountered for departmental applications with small databases (part of the server sprawl). DAS challenges include lower storage utilization; difficulty in meeting anticipated SLAs; and higher demand for space, power and cooling.

Moving applications into storage area networks (SANs) is an opportunity for IT to regain control of the storage environments, which can improve unreliable or inconsistent backups and reclaim unused capacity. However, many SAN implementations continue to be deployed as islands of storage. Although utilization of storage resources is improved over DAS, reclaiming unused storage remains a challenge in disparate SANs. The silos of SAN storage remain complex to manage, protect, and scale. In addition, the need for more space and demands for higher power and cooling persist.

Centralizing islands of SAN onto a single multiprotocol platform represents the ultimate in realizing storage efficiency and is the first step toward deploying a dynamic SQL Server database infrastructure. Utilization in this architecture is higher; SLAs become predictable as well as achievable; management complexity is reduced; and demand for space, power, and cooling is lower.

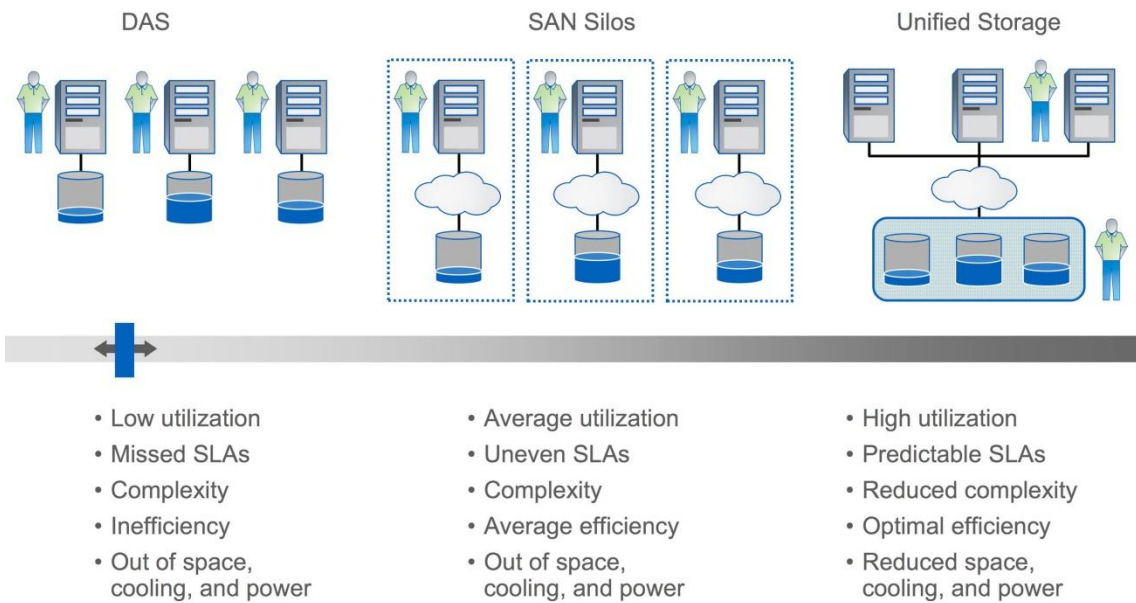


Figure 1) Database storage architecture options.

3 VIRTUALIZATION INCREASES DEMANDS ON STORAGE

In order to derive maximum efficiency from consolidation using virtualization, the decisions regarding the storage architectures are of paramount importance. The storage architecture deployed should be carefully considered as it can drastically affect the benefits gained from consolidation using virtualization. Server virtualization indeed allows dramatic levels of server consolidation; a consolidation ratio of 10:1 is very feasible. While this eliminates the old silo model of one server per application, hardware failures, including storage, can now take down multiple applications, instead of just a single instance. This situation drives the need for more reliable storage.

This increased risk calls into question the efficacy of RAID 5 as an approach to storage reliability. RAID 5 protects against failure of a single disk in an array. However, should a second disk in the array fail before the data of the first failed disk has been restored, the whole data set will need to be reloaded, resulting in significant downtime. Considering how long it takes to restore a single disk after failure because of the size of today's drives, this means that RAID 5 offers insufficient protection for most virtualized server environments.

Longer backup times also become a critical problem, as data per server increases by 10X or more. The higher utilization rate of the servers running the VMs results in insufficient resources to back up all the VMs simultaneously. In any case, the applications on a server might not have a common backup window, especially given today's worldwide 24x7 demand for application services.

Additionally, as IT operations become ever more critical, disaster recovery (DR) increases in priority. It is extremely difficult to provide DR in a consolidated DAS environment, and it becomes even more complex when virtualization enters the picture. Where disaster recovery is a priority, new approaches to storage are needed for consolidation with virtualization.

Finally, while virtualization greatly enhances server provisioning, storage provisioning is generally much slower, unless new methods are brought into the picture. In order to realize the benefits of consolidation using virtualization, the storage deployed must meet more stringent requirements.

4 THE NETAPP ADVANTAGE: STORAGE EFFICIENCY CAPABILITIES

NetApp offers a number of storage efficiency capabilities that provide numerous advantages to consolidation and virtualization. NetApp offers:

- Unified architecture: An approach to storage that supports multiprotocol and multivendor storage deployments.
- Thin provisioning: Optimizes storage utilization in the multiapplication environment typical of virtualized servers.
- Simple and rapid provisioning and cloning: Designed to match the server provisioning characteristic of server virtualization.
- End-to-end deduplication: A fundamental component of NetApp's core operating architecture, Data ONTAP®. NetApp deduplication is the first that can be used broadly across many applications, including primary data, backup data, and archival data.
- Space-efficient snapshots and rapid recovery: NetApp Snapshot™ and SnapRestore® software enables frequent backups and rapid recoveries of entire file systems or data volumes in seconds, regardless of capacity or number of files.
- Simplified disaster recovery: By replicating data at high speeds over a LAN or a WAN, SnapMirror® software provides the highest possible data availability and fastest recovery by application.

NetApp provides a viable storage architecture for consolidation and virtualization, and its advantages are categorized as:

- Flexible IT architecture
- Complete data protection
- Efficient storage

5 FLEXIBLE IT ARCHITECTURE

NetApp's flexible unified architecture forms the foundation for designing the storage architecture for virtualized environments.

The unified architecture allows the selection and deployment of the most suitable SQL Server storage environment: FC SAN or iSCSI SAN. And NetApp's single operating system provides all of NetApp's protection, provisioning, and management capabilities which are available to all NetApp systems and storage architectures. NetApp's data management capabilities, including end-to-end deduplication, Snapshot copies, and thin replication, allow improvement in storage efficiencies. In addition, the ability to create virtual, writable copies of Snapshot copies provides fast and space-efficient provisioning of VMs and copies of databases for development and testing. Accordingly, the same skill sets and processes can be used to manage and protect all storage architectures.

For customers who want to use their existing storage from other storage vendors (including IBM, HP, Hitachi Data Systems, and EMC), NetApp offers the V-Series open storage controllers, which allow customers to manage disk arrays from other major storage vendors as if they were NetApp storage (see Figure 2). The current generation of NetApp disk shelves can also be used with V-Series controllers, creating more choice and flexibility for capacity expansion. The V-Series runs virtually all of the NetApp management and protection software, bringing the power of NetApp to an existing infrastructure.

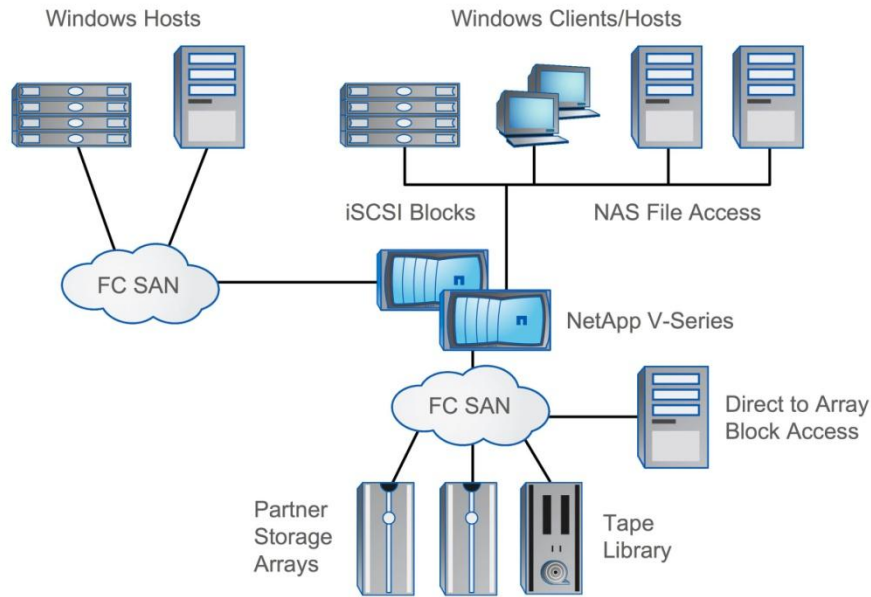


Figure 2) V-Series network diagram.

A highly valued advantage of server virtualization is the ability of the virtualized IT infrastructure to quickly respond to change. As new applications become available, it is very important to move rapidly from development through test and deployment to achieve faster time to market. Traditional storage systems struggle to support rapid rollout. Test, development, and deployment activities typically require generation of physical copies of data or VMs, which can be costly in terms of time and disk space.

Using FlexClone® technology, NetApp enables rapid creation of virtual copies of SQL Server databases and VMs. Application and database developers can thus use FlexClone to create database copies and VMs to run SQL Server instances in a few seconds, using minimal additional storage and reducing both storage and administrative costs for supporting development and test efforts (see Figure 3)

Traditional Storage Provisioning for VMs



NetApp Storage Provisioning for VMs



Figure 3) Instant storage provisioning.

6 COMPLETE DATA PROTECTION

NetApp uses a combination of next-generation RAID technology and built-in backup and replication technologies to offer complete data protection.

The previous standard for data storage was RAID 5. As discussed before, the increased risk posed by the data needs of virtualized servers has made the use of RAID 5 problematic. With RAID 5, protection is provided for single disk failures only. A number of factors have combined to require data centers to prepare for double disk failure scenarios. RAID 5 is inadequate to address this eventuality.

A number of storage vendors have responded to this issue by promoting RAID 10 (mirroring), which offers protection for some, but not all, double disk failures and provides much higher performance than RAID 5. However, these improvements come at a steep price, since RAID 10 effectively doubles the raw capacity of the array and hence doubles the cost of the solution.

Another potential solution lies in RAID 6. It is designed to tolerate two distinct drive failures by storing two sets of distributed parities. However, the complex mechanism is likely to compromise data rate in and out of the system. Accordingly, while RAID 6 delivers costs comparable to RAID 5 and considerably higher resiliency, most implementations suffer performance lags compared to RAID 10 and in some cases to RAID 5 as well.

NetApp offers RAID-DP[®], a standard Data ONTAP feature that safeguards data from double disk failure. Integrated with the NetApp WAFL[®] (Write Anywhere File Layout) file system, RAID-DP delivers data protection plus performance. In addition, RAID-DP delivers the high performance of RAID 10 with the low cost of RAID 5. With RAID-DP, there are no tradeoffs among performance, cost, and protection.

INSTANTANEOUS AND EFFICIENT BACKUP FOR SHARED STORAGE


Virtualization poses new problems for servers tasked with backing up all the data volumes associated with a single server (which might support on the order of 10 VMs and their data). The amount of data can tax the server utilization to the point that backup windows cannot be met. Further, the backup requirements result in reduced CPU resources for running applications. In addition, tape backups are slow, manual processes, with a growing inventory that adds to cost and complexity. Traditional disk-to-disk backup methods, while faster than tape, require large amounts of storage space and are subject to similar time and server utilization constraints. Recovery is likewise slow and prone to error, especially recovering from tape, and recovery points are limited.

NetApp addresses all these problems with Snapshot technology. Snapshot is integrated into the core operating system of NetApp storage systems, Data ONTAP. Snapshot technology enables customers to create point-in-time copies of file systems, from a single file to a complete disaster recovery solution. The technology operates while applications are running and creates Snapshot copies in less than a second, regardless of volume size or level of activity. It does this with no server overhead, freeing server CPU resources to focus on application efficiency, further enhancing the benefits of consolidation and virtualization.

NetApp SnapRestore software uses stored Snapshot backups to recover entire file systems or data volumes in seconds, regardless of capacity or number of files. Restores can be made virtually instantaneously, providing a broad range of point-in-time options. In contrast, alternative storage solutions copy all of the data and require much more time and disk storage for the backup and restore operations.

NetApp offers tools to extend the backup capability to full, built-in data protection. MetroCluster combines array-based clustering with synchronous mirroring to deliver continuous availability and zero data loss at the lowest cost. MetroCluster delivers automatic storage failover between arrays within 100 km of each other. Using SnapMirror, full deduplication-aware disaster recovery can be achieved with no distance limitation (see Figure 4).

Continuous Availability with MetroCluster



- Built-in instantaneous backups
- Automatic storage failover
- Efficient VM data replication across unlimited distances

Figure 4) Built-in data protection.

7 MICROSOFT TOOLS AND NETAPP VIRTUALIZATION

To match the ease of migration of VMs across virtualized servers, NetApp offers tools that work in concert with Microsoft virtualization capabilities. NetApp Data Motion™ allows easy and quick data migration across multiple storage systems while maintaining continuous access to applications. The result is that customers can keep their shared storage infrastructure running as they add capacity, update infrastructure technology, and balance performance. As Microsoft Live Migration allows real-time migration of VMs across physical servers, NetApp Data Motion allows data stores containing VMs to move between storage controllers for load balancing or planned maintenance.

SnapManager® for Hyper-V protects virtual machines in Windows Server 2008 R2 Hyper-V environments. It can back up and restore VMs running on standalone servers or in Windows failover clusters. SnapManager for Hyper-V integrates with the Microsoft Volume Shadow Copy Service to create application-consistent Snapshot copies of VMs. This allows recovery of entire VMs in seconds.

SnapManager for Microsoft SQL Server is tightly integrated with Microsoft technology to help streamline database storage management while simplifying storage layout planning, backup, and restore operations for SQL Server databases. SnapManager enables DBAs to perform granular backups and restores as well as cloning of SQL Server databases. SnapManager can dramatically reduce SQL Server data recovery times from hours to minutes. SnapManager can also be used to automate backup, recovery, and database cloning.

8 EFFICIENT STORAGE

VMs store large quantities of identical data, such as operating systems, patches, and application software. With NetApp deduplication, these redundancies are eliminated and reduced to a single instance on disk, which in some cases can reduce storage space requirements by substantially more than 50%. Moreover, NetApp can provide deduplication across all tiers of storage: primary, backup, and archival.

NetApp's thin provisioning strategy drastically increases storage utilization, with a consequent reduction in power, cooling, and space costs. Thin provisioning allows storage to be oversubscribed, while only deploying the physical storage in use. Storage can be added transparently to the servers and applications as required. A recent study¹ estimated that using thin provisioning with NetApp storage can result in more than 70% utilization and in purchases of 50% less storage for virtualized environments. Using NetApp V-Series open storage controllers with third-party storage similarly results in up to 35% savings in reduced power, cooling, and storage space. NetApp offers a guarantee that using NetApp for a virtualized environment will result in using less storage².

As previously mentioned, NetApp SnapManager for SQL Server and Hyper-V improves administrator efficiency, enabling them to perform data management on NetApp storage using application-centric tools.

9 CONCLUSION

NetApp storage supports all strategies for SQL Server consolidation. If virtualization is chosen as the preferred consolidation strategy, NetApp provides storage efficiency capabilities that will result in lower costs, higher flexibility, and greater reliability. NetApp delivers these benefits through three major advantages: a flexible IT architecture, complete data protection, and a highly efficient storage architecture.

REFERENCES AND MORE INFORMATION

NETAPP MICROSOFT SQL SERVER SOLUTIONS

www.netapp.com/us/solutions/applications/microsoft-sql/

SNAPMANAGER FOR SQL SERVER

<http://www.netapp.com/us/products/management-software/snapmanager-sql.html>

FLEXCLONE

www.netapp.com/products/software/flexclone.html

V-SERIES OPEN STORAGE CONTROLLERS

<http://media.netapp.com/documents/v-series.pdf>

METROCLUSTER

www.netapp.com/us/products/protection-software/metrocluster.html

NETAPP DEDUPLICATION

<http://media.netapp.com/documents/virginia-credit-union.pdf>

www.netapp.com/us/company/news/news_rel_20071127.html

¹ <http://www.netapp.com/us/library/analyst-reports/ar-1054.html>, Oliver Wyman Study, "Making Green IT a Reality," November 2007

² www.netapp.com/guarantee.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.



www.netapp.com

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, NetApp Data Motion, RAID-DP, SnapManager, SnapMirror, SnapRestore, Snapshot, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, Windows, and SQL Server are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7095-0310