

# BUILDING THE RIGHT INFRASTRUCTURE FOR MICROSOFT SHAREPOINT

EMC VNX, EMC FAST Suite, EMC Replication Manager, Microsoft SharePoint 2010, Metalogix StoragePoint, Metalogix Selective Restore Manager

- EMC<sup>®</sup> VNX<sup>®</sup> delivers automated tiered storage for SharePoint
- Metalogix StoragePoint enables file share content migration to SharePoint and VNX
- EMC FAST Suite optimizes Search efficiency and performance
- Metalogix StoragePoint enables remote BLOB Storage

## EMC Unified Storage Division

### Abstract

This white paper explains how EMC and Metalogix complementary products and technologies tightly integrated with SharePoint, help businesses expand the value of their SharePoint investment with file share consolidation, unified storage, automated storage tiering, and data protection delivering immediate improvements in operational efficiencies and lowering TCO. Together, EMC and Metalogix help companies bridge disconnected business processes by moving SharePoint into central IT.

October 2012



Copyright © 2012 EMC Corporation. All Rights Reserved.

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.

VMware, ESXi, VMware vCenter, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number: H11135

## Table of contents

The New Normal – Explosive Growth of Data .....	4
Challenges of Explosive Data Growth .....	4
Solution Essentials .....	4
EMC and Metalogix SharePoint Solutions.....	5
File Share Consolidation .....	5
SQL BLOB Externalization .....	5
Improved Asset Utilization .....	6
Flexible End Point Management .....	6
Unified Storage .....	7
Automated Tiering and Accelerated File Search .....	7
Reduced Time to Manage, Backup, and Recover .....	8
Virtualizing SharePoint Is Essential .....	8
Joint Solution Value .....	9
Savings and Efficiency .....	10
Summary .....	10
Resources .....	10

# Take control of your unstructured data, improve efficiency, lower TCO

## The New Normal – Explosive Growth of Data

We are in the “Content Explosion Era”. Organizations continue to create unstructured data at relentless rates causing people and business processes to become disconnected resulting in the creation of information silos (file shares, NAS filers, legacy ECM, etc) that don’t make data available to those who need it in real-time. From an IT perspective, data growth impacts operational efficiency and lowers TCO because assets are underutilized, performance and service levels cannot be met, backup windows are extended, recovery processes are suspect, and lack of version control is the rule – to list a few. Also contributing to the “new normal” are the regulations and compliance laws being placed on companies around data retention and audits. As Microsoft’s fastest growing software business, SharePoint emerges as an accessible and expandable content management platform with tight Microsoft Office integration that can help businesses connect people to the information they need, when they need it and wherever they are. However many IT challenges remain and the combination of EMC and Metalogix solutions are designed to fill the gaps created by the explosive growth of content.

## Challenges of Explosive Data Growth

SharePoint provides business-critical services to organizations allowing them to connect disparate business processes and empower people across boundaries, in order to rapidly respond to business needs. However, without the right infrastructure, moving SharePoint into central IT won’t be easy. When relentless data growth and content sprawl occurs, businesses cannot make the most of what they already have, budgets get stretched, resources become grossly underutilized, service level management (performance and availability) is significantly reduced, and opportunities to invest in new solutions are held back. Simply, data sprawl spawns complexity sprawl with content located across a variety of repositories. When this occurs, IT has two paths to follow – 1) add more complexity to their infrastructure with outmoded products or instead, 2) look for established, cost effective solutions that can deliver improvements in operational efficiencies affordably. Data that cannot be accessed and searched isn’t generally available to employees, it can’t be protected, and data integrity cannot always be guaranteed. The lack of visibility and access to data impacts business effectiveness in many other ways, including:

- Increased storage/management costs with inactive data stored on Tier 1 storage
- Regulatory compliance is difficult to achieve despite native tools
- Performance issues increase as customers scale their SharePoint infrastructures.
- Difficult to establish central backup and data protection schemes
- Inflexible infrastructures that cannot scale economically

## Solution Essentials

- Increase asset utilization **up to 80%** through file consolidation
- Reduce size of SQL DBs **up to 90-95%**
- Tune virtualized SharePoint SQL in **80% less** time
- Increase Search response times **by 30-40%**
- Realize **10–30X** SQL backup storage reduction
- Run virtualized SQL Server **4X faster**
- Lower SQL licensing costs by up to 75%

EMC and Metalogix solutions bring SharePoint to life in helping businesses build an effective and cost efficient infrastructure to consolidate unstructured data, reduce the size of SQL database BLOBs, improve operational efficiency, and reduce TCO.

**EMC and Metalogix  
SharePoint  
Solutions**

The challenges for IT are clear when it comes to content explosion – centralizing and taking control of data to increase the reliability and availability of information throughout the organization. This challenge isn't as daunting when IT understands that they can extract more out of their SharePoint investments to consolidate file shares (from multiple sources) into a centralized resource to increase asset utilization, simplify management, and enhance service levels. Companies can easily maximize their investment in SharePoint as a document repository with infrastructure solutions from EMC and Metalogix (Figure 1). EMC and Metalogix solutions provide complementary products and technologies tightly integrated with SharePoint, helping businesses expand the value of their SharePoint investment with file share consolidation, automated tiered storage, flash technology, and data protection delivering immediate improvements in operational efficiencies and overall TCO.



Figure 1. EMC and Metalogix help companies plan and deploy global SharePoint architectures to meet scalability, security, and efficiency needs.

**File Share  
Consolidation**

When data is spread across the enterprise and inaccessible, it can't provide full value to the business. Metalogix StoragePoint Librarian improves file share utilization and operational efficiency by consolidating and cataloguing legacy unstructured data (CIFS/NFS shares, ECM files, existing file shares) across the enterprise into SharePoint collectively or on a selective basis determined through a governance discovery process. Users can select the files they would like to catalog and file shares can remain accessible as desired. Librarian places the associated metadata into the SQL database under the full control of SharePoint without having to move the physical file. At this point, data can be easily indexed, searched, managed, versioned and accessed like any other document within SharePoint while preserving the overall user experience.

**SQL BLOB  
Externalization**

With ever growing volumes of unstructured content being stored and accessed in SharePoint the challenge to efficiently and effectively manage the unstructured files (BLOBs) grows increasingly more difficult. On average, 90-95% of the data stored the SharePoint SQL Content Database is unstructured content such as Microsoft Office files, PDFs, and images. Storing BLOBs in the SharePoint SQL database slows operations such as search, upload/download, browse and modify particularly with larger files. Additionally, BLOBs end up consuming a disproportionate volume of expensive and vital database real estate which can adversely impact the time for backup and recovery. "Typically, as much as 80 percent of data for an enterprise-scale deployment of SharePoint Foundation consists of file-based data streams that are stored as BLOB data. These BLOB objects comprise unstructured data associated with SharePoint files. However, maintaining large quantities of BLOB data in a SQL Server database is a suboptimal use of SQL Server resources. You can achieve equal

benefit at lower cost with equivalent efficiency by using an external data store to contain BLOB data.” (Source: <http://msdn.microsoft.com/en-us/library/bb802976.aspx>). Metalogix StoragePoint’s Remote BLOB Storage (RBS) capabilities offers IT the option to externalize BLOB content to VNX tiered storage (Figure 2) This activity reduces the database size, increases the VM density per physical server and improves SQL’s ability to scale - without jeopardizing the user response time.

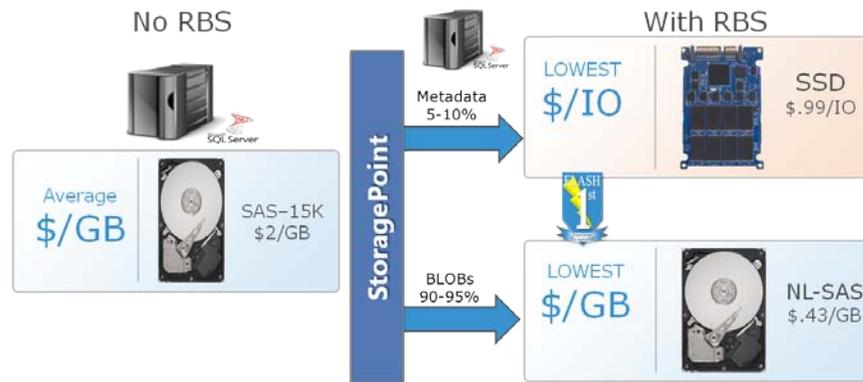


Figure 2. StoragePoint and EMC VNX lower both costs/IO and costs/GB

Improved Asset Utilization

Most file servers are poorly utilized even if they are virtualized. Users tend to keep older, duplicate, or non-compliant files on their already underutilized file servers. A discovery process can help identify what content needs to be kept and would make a good candidate for migration into SharePoint and onto the VNX. This leads to a smaller content body even before any migration begins. Files can then be stored on a much more efficient VNX reducing the number of file print servers and associated operating expenses. Native de-duplication and compression algorithms can further reduce the overall file content. VNX arrays can be used at very high utilization rates without impacting performance. Thresholds are typically set to 70-80% utilization before the user receives an alert that performance might level off. SQL servers are often underutilized as well with up to 70% of CPU cycles idle. Some customers still use physical SQL servers as part of their SharePoint environment due to concerns about possible performance impact. The challenge of keeping SQL on a physical infrastructure significantly limits the ability to provide a more elegant disaster recovery strategy as well as leading to much higher SQL licensing costs. SQL2012 Server Enterprise on a physical server is licensed per physical core. In virtual server environments it is licensed per vCPU in an ‘all you can use’ model. Virtualizing SQL servers and storing them on the VNX can improve VM density, lower SQL licensing costs, and dramatically improve performance (workload dependent). Assuming a vCPU : physical CPU ratio of 4:1, users can lower their SQL license costs by as much as 75%. Offloading BLOBs with StoragePoint reduces the SQL database corpus by 90-95% and lowers CPU load for further VM density and consolidation. At the same time VNX FAST Cache can reduce SharePoint user response times by 30-40% depending on the activity.

Flexible End Point Management

StoragePoint’s flexible end point management allows users to apply policies to cascade BLOB placement to less expensive media based on file age, type and size. Advanced filters can specify which data is placed on a VNX share that has File Level Retention enabled to ensure compliance. Native StoragePoint capabilities can encrypt sensitive data before it is moved to a CIFS share on the VNX platform.

## Unified Storage

EMC VNX is a simple, efficient, and powerful unified storage system that connects with SharePoint environments to deliver kind of performance, efficiency, and scalability expected in a SharePoint environment that may be supporting thousands of end users and file sources. As a multi-protocol system, VNX allows customers who made investment decisions based on block storage to easily add NAS connectivity at a later date. As data is added to the VNX, storage pool and load rebalancing technologies automatically initiate to maximize storage efficiency and performance across different drive tiers. When this occurs, VNX can provide up to 70% utilization on SAS & NL-SAS drives and up to 90% on SSDs. VNX also enhances data protection on the array with its support for mixed RAID types within a storage pool and flexible RAID options for pools that allow you to choose different RAID options for increased efficiency.

*Note: Microsoft achieved a record 200% increase in SharePoint scalability with VNX5700 at the SharePoint Conference 2011. (Scale Test Report for Very Large-Scale Document Repositories)*

## Automated Tiering and Accelerated File Search

Arguably, as data environment becomes more expansive, the more important it is to have an optimized infrastructure to accelerate user activities like “search”. Using EMC’s FLASH 1<sup>st</sup> strategy optimizes the use of flash technology significantly accelerating the performance of these activities and improving storage tiering efficiency.

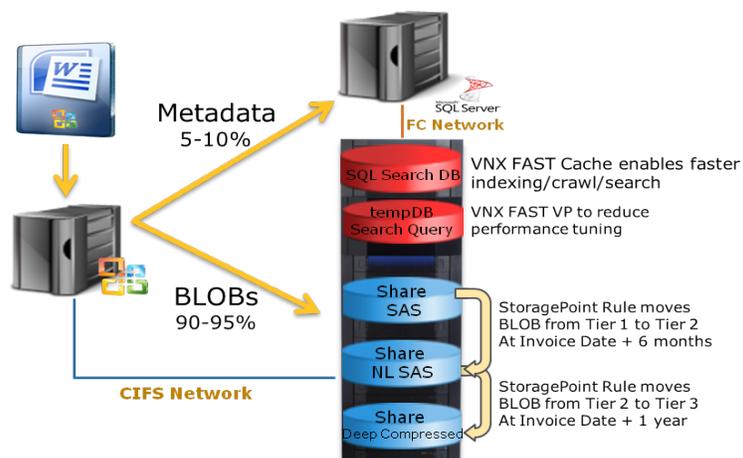


Figure 3. Using EMC FAST Suite optimizes performance and efficiency

FAST Cache is an ideal solution for improving the performance of extensive SharePoint SQL data searches (Figure 3). FAST Cache can be used for reads \*and\* writes, is non-volatile, and copies the hottest data to Flash as an extension of DRAM enabling Flash to be used as a secondary cache between the disks and the DRAM cache. This provides a real-time performance boost by ensuring that the hottest data is served from the highest-performing Flash drives for as long as needed. In doing so, FAST Cache provides an immediate performance benefit to the busiest workloads by holding a large percentage of the most frequently used data in high performance FLASH drives.

FAST VP helps lower TCO by tiering data onto less expensive drives (NL-SAS) optimizing drive utilization through a relative ranking of the data over time. FAST VP moves that data across drives or between tiers automatically as a scheduled activity helping to reduce manual SQL database tuning by as much as 80%.

Having the FAST Suite allows IT to focus more on SharePoint capacity and service management issues, not on disk technologies. Together they deliver a policy-based, fully automated data placement strategy for achieving the highest system performance at the lowest cost attainable simultaneously across your SharePoint environment. Externalizing BLOBs using Metalogix StoragePoint actually compounds the benefits of FAST Cache. A 90-95% reduced SQL instance offers a smaller working set which helps accelerate FAST Cache warming and acceleration with the highest benefits occurring in document searches.

So having StoragePoint and FAST Suite together means you create an optimized solution that delivers automation (storage tiering) in the array resulting in fewer drives, lower energy bills, and lower \$/IOPs and \$/GB. The combination of Metalogix and EMC technologies can drive utilization to 80% to 90% depending on the drive type that's being utilized.

Reduced Time to  
Manage, Backup,  
and Recover

Proper protection of the SharePoint farm and the information that lives within it is essential. Most users rely on native backup tools like Central Administration or STSADM. Unfortunately these are not designed for scale. Microsoft recommends using external recovery tools for content databases larger than 100GB or site collection larger than 15GB (Source: <http://technet.microsoft.com/en-us/library/cc263427%28office.12%29.aspx>). Which tool is appropriate to do the job depends on several factors:

- Speed of backup: Can the tool perform within the maintenance window for your databases?
- Completeness of recovery.
- Granularity of objects that can be recovered.
- Backup type supported (full, differential, or incremental).
- Complexity of managing the tool.
- Requirements for also protecting the host OS and distance requirements.

The combination of VNX snapshot technology, EMC Replication Manager, and Metalogix Selective Restore Manager offers unique advantages over many of the available technologies:

- Reduced backup and recovery times. Testing revealed that the time to restore an entire farm with content database and BLOB data took around 30 minutes compared to many hours using native tools.
- Complete and synchronized backup and recovery process.
- Reduced management overhead through automated creation of point in time copies and intuitive recovery workflows.

**Virtualizing  
SharePoint Is  
Essential**

Virtualization continues to accelerate and adoption rates are rising to record levels for both servers and application portfolios. What's driving this is the need to take costs out of IT and create IT as a service so IT organizations can appear more like cloud service providers and be more cost effective. This trend is occurring because it has become abundantly clear that physical application deployments restrict IT agility and flexibility including wasting time waiting for servers to arrive and software to load, performing effective backup/recovery operations, testing complex business continuity schemes, and keeping pace with dynamic application and data movement. SharePoint is an ideal application environment for virtualization because of its intricate multi-platform deployment structure of index servers, web front ends, database servers, and more to support workgroup collaboration, content

management, and workflow activities. Virtualizing SharePoint will increase productivity and make it easier to scale SharePoint resources like CPU, memory, and capacity and exhibit flexibility by moving workloads to different servers and rapidly recover from outages – all on the fly and without application downtime. The right storage infrastructure plays a significant role in delivering the benefits of a virtualized SharePoint environment, including:

- Consolidation: Achieve 2-10x consolidation ratio, especially for larger deployments
- Performance: Improved front end performance with more, smaller WFEs rather than few large WFEs
- Availability: VM-based protection for SharePoint provides homogeneous HA (WSFC, VMware® HA)
- Business Continuity: Simplified DR management (Geo-Clustering, vCenter™ Site Recovery Manager)
- Maintenance: Live migration of virtual machines (Hyper-V Live Migration, VMware vMotion)
- Load Balancing: Maximized performance with balanced cross farm HW utilization (SCVMM PRO, VMware DRS)
- Costs: Lower TCO of SharePoint and simplify management through virtualization of farms

EMC VNX storage is optimized for virtualized application workloads and is consistently recognized as the primary storage platform for virtualization environments. The VNX is a unified storage system that consolidates file and block technologies and multiple protocols into a single platform to support virtualized SharePoint environments with built-in simplicity, efficiency, and power. VNX increases IT flexibility, reduces management costs by up to 50%, is highly efficient with built-in compression and thin provisioning, and is very powerful with leading INTEL processors and advanced features like the FAST Suite, integrated snapshots, and continuous protection functionality.

A SharePoint environment, virtualized and supported by EMC and Metalogix infrastructure solutions, offers IT significant advantages including performance at parity with physical implementations, the facilitation of rapid growth, high availability, elimination of scalability constraints, rapid provisioning, and effective data protection.

## Joint Solution Value

Powerful solutions to consolidate file share content into SharePoint and VNX

EMC	Metalogix
VNX unified storage as the ideal home for metadata and BLOBs	StoragePoint externalizes BLOBs to make SQL lean and responsive.
Flash 1 <sup>st</sup> , Compression, De-Dupe	StoragePoint's File Share Librarian makes data searchable & accessible without ever having to move a file
FAST Suite to enable 90+% faster access to critical information	StoragePoint's File Share Librarian brings stranded File Shares and legacy content into the SharePoint fold
Leverage VNX File Level Retention (FLR) for compliant retention	Encrypts BLOBs to protect sensitive data at rest and in transit.
VNX FAST Cache for faster indexing/search/crawl	StoragePoint applies policies to store content in the right repository at the right time

EMC	Metalogix
VNX FAST to automate performance management & tuning	StoragePoint Endpoint Management for tiered storage and migration to private / public cloud
VNX compression to reduce storage costs	Metalogix Selective Restore Manager allows granular recovery of files, lists, list items and libraries without requiring farm recovery
EMC Storage Integrator to streamline provisioning and ease of management, monitoring, and orchestration	Migration tools to enable journey to the private and public cloud
System Center Mgmt Packs for monitoring / management in the private cloud	StoragePoint's File Share Librarian connects teams / information silos to improve collaboration
EMC RecoverPoint + Cluster Enabler for farm-level protection	StoragePoint eliminates SQL bloat by reducing the size of SharePoint content databases by up to 95%

## Savings and Efficiency

EMC and Metalogix solutions for SharePoint together create enormous opportunities for savings and efficiency – for IT and the business. The immediate benefits start with better asset utilization across the enterprise where unstructured data can be economically and efficiently shared and protected. With these complementary solutions, IT can experience increased utilization rates on tiered storage by up to 70% on SAS/NL-SAS and up to 90% on SSD. Also, when the FAST software is implemented into the SharePoint environment, more efficiency is introduced into the solution resulting in fewer drives, lower energy bills, lower \$/GB, and lower \$/IO. And when SharePoint is virtualized in combination with RBS, customers will see lower SQL licensing costs while being able to reduce database sizes and increase VM density per physical server.

## Summary

Content explosion is a reality and so are the challenges that come with it. SharePoint's potential to reach more of the enterprise as a content platform for data consolidation and improve its ROI to the business is at hand. EMC and Metalogix are uniquely positioned to consolidate unstructured information into SharePoint, bridge gaps between information silos, boost productivity by making information easier and faster to find, and improve data protection. EMC and Metalogix solutions for SharePoint radically improve asset utilization, dramatically improve operational efficiency, significantly lower SQL costs, and optimize virtualized SharePoint environments. EMC and Metalogix create the right infrastructure allowing businesses to increase revenue opportunities, remain competitive, and extend the value of their SharePoint investments deeper into the enterprise in order to enable the knowledge worker of the future.

## Resources

[Technical Whitepaper: EMC Efficient BLOB Storage Management for Microsoft SharePoint](#)  
The EMC white paper explains how Metalogix StoragePoint is used to externalize Microsoft SharePoint BLOBs to EMC VNX for file storage.

[Metalogix StoragePoint BLOBulator:](#)

A tool to help you estimate how much StoragePoint can shrink SharePoint content databases

[MetaLogix e-book: RBS for Dummies](#)

With this book, you see how Remote BLOB Storage (RBS) is used to optimize the storage of your content.

[EMC Everything Microsoft Community](#)

Comprehensive web site for content, messages, announcements, documents, and more.