



The **6 Key Points** of SharePoint Infrastructure Planning

Windows IT Pro

SharePoint products and technologies have readily found a place within the business environments of many organizations. SharePoint provides an ideal platform for business collaboration and other functions, therefore, it is becoming a critical tool that must be properly architected, maintained, and designed to be easily expanded. A large number of businesses and organizations are considering either deployment of new SharePoint environments or expansion of existing environments and they first need to consider design criteria for the platform.

Design and deployment of an infrastructure ideally suited for SharePoint is no small task, however. An improperly constructed SharePoint environment can experience significant growing pains or can leave an organization without a proper plan in place for recovery. Critical topics surrounding SharePoint architecture and infrastructure can be divided into six key points: Tiered/Unified Storage, Virtual Infrastructure, Backup/Recovery/Archive, Business Continuity, Enterprise Content Management, and Security. These points, if examined properly during design and deployment, can help make a SharePoint environment much more resilient, recoverable, and scalable.

Tiered/Unified Storage

SharePoint is a database-dependent technology that requires fast and reliable access to stable disk infrastructure. SharePoint uses Microsoft SQL Server database servers to store content, and many organizations are discovering that a significant bottleneck to SharePoint performance is oftentimes related to disk performance on the database server role. In addition to performance requirements, SharePoint can consume larger amounts of space within the SQL database environment when used as a document management platform, as all versions of all documents are stored in full in the SharePoint content databases. Subsequently, a scalable, dependable and highly performance-oriented storage system that can quickly adapt to data growth is recommended.

When designing a SharePoint environment, it is often common to deploy on multiple farms. In fact, it is best practice to deploy at least two farms, one for production and one for testing/staging, and it is also common to have other farms dedicated to extranet users, intranets, or for development work. Each of these environments may have separate performance requirements, subsequently, it is valuable to invest in both a disk infrastructure that is tiered, with high performance disks allocated to the most critical farms, and in a lower performance, yet cost-effective disk for other applications.

Virtual Infrastructure

Server virtualization technologies are fast becoming commonplace in many organizations, in some cases replacing a significant percentage of traditional physical servers with

virtualized replacements. Server virtualization can lower energy costs significantly and can also maximize server and storage utilization by reducing the amount of unused disk, memory, and processor consumed by servers.

SharePoint is no stranger to virtualization, and is fully supported running within Microsoft Server Virtualization Validation Program (SVVP), which includes many VMware® products and Microsoft's Windows Server 2008 Hyper-V. Indeed, virtualization of SharePoint server roles gives architects greater flexibility to design multiple SharePoint farms, highly available environments, and greater scalability for little overhead.

Commonly virtualized SharePoint server roles include the web role, the query role, and oftentimes the index and application roles. When using the database role, the disk IO and CPU requirements must be addressed, as they can be significant depending on the load. Allocating high-performing raw disk volumes and multiple processors to virtual machines can also be effective for the SQL database role.

Backup/Recovery/Archive

Backup and recovery of a SharePoint environment is notoriously difficult. The built-in SharePoint backup and restore tools don't address all aspects of SharePoint recovery, and can't handle individual item recovery. Many SharePoint administrators simply backup the SharePoint SQL databases, but this doesn't backup the SharePoint index, customizations made to web role servers, search settings, or solutions in the "12 hive" of SharePoint servers.

A good recovery plan for SharePoint subsequently takes into account recovery of all farm components, and ideally takes advantage of the built-in Volume Snapshot Service (VSS) writer, which allows for a snapshot of block-level changes within a farm at regular intervals. Products that integrate with the VSS writer can allow for unified protection of not just SharePoint, but other Microsoft VSS-aware applications as well. Ideal recovery solutions also allow for de-duplication of content, saving backup space and enabling faster backups at the source.

Archiving is another often-overlooked but highly valuable concept useful for optimizing a SharePoint environment. All content in SharePoint content databases, whether actively used or simply stored long-term, consumes the same high performance disk infrastructure. In addition, SharePoint document versioning can result in content databases growing exponentially over time. Because of this, designing an archiving technology into a SharePoint infrastructure that physically moves documents and items from near-term high speed storage to a long-term, slower disk makes sense for many organizations. Archiving also stores the BLOBs outside of SQL while at the same time keeping transactional consistency, allowing for better overall performance.

Business Continuity

Considering the criticality of most SharePoint environments, planning for disaster recovery and business continuity of the SharePoint data and applications is of the utmost importance. Business continuity of SharePoint content mainly centers on the concept of getting SharePoint data from one physical site location to another to allow it to be recovered and utilized from that second location. Ideally, the recovery solution would allow for a short or near zero RPO (Recovery Point Objective), or the maximum amount of data that could be lost in the event of a failure.

From a SharePoint perspective, business continuity requires at a minimum the SharePoint content databases where the documents, lists, and site structure are housed. Some site replication solutions, such as SQL Server 2005/2008 asynchronous database mirroring, support only replication of the content databases, requiring that search settings be replicated and the index be recreated from scratch which is a very manual process. However, solutions that synchronously replicate all SharePoint databases simultaneously or that allow for geographically dispersed clusters can avoid these extra steps and provide quicker RTO (Recovery Time Objective), which defines how long it takes to bring the environment to full recovery.

Enterprise Content Management

Enterprise Content Management (ECM), encompassing both Document Management and Web Content Management, is a central component of SharePoint and is often the driving factor behind the technology. ECM requires a stable, robust, and flexible infrastructure for content storage, metadata capture, records management, litigation, and policy-based archival of older content—concepts that aren't fully developed with out-of-the-box SharePoint functionality.

SharePoint provides ECM capabilities out of the box but they are often extended to allow for more comprehensive document management and customization of ECM capabilities so that ECM will match needs of the organization deploying it. For example, custom workflows and solutions can help automate document and records management, litigation, and e-discovery of content.

Security

The final key point when implementing a SharePoint infrastructure is security - a critical, yet often overlooked element of a properly designed SharePoint environment. SharePoint users need to trust the system, and unmitigated information leakage is of major concern to many organizations. A Data Loss Protection platform based on Microsoft Rights Management Services (RMS) with SharePoint can allow for rights policies on documents in

SharePoint, limiting their contents from being printed, cut/pasted, or saved elsewhere. However, certain out-of-the-box limitations exist with RMS. For example, administrators can override rights management protections. In this case, additional data loss solutions may need to be examined. Additionally, a reporting and monitoring infrastructure that examines and verifies end user activities is a critical element for ensuring secured access to SharePoint content.

Addressing the Six Points with EMC® Technology

EMC Corporation offers end-to-end solutions for SharePoint infrastructures that scale from small standalone servers up to enterprise-scale SharePoint farms with multiple web front-ends and SQL databases. Whether a business is consolidating, protecting, virtualizing, or securing SharePoint, EMC has the right combination of technology and people to solve the key, pressing technical and business problems discussed here across all Microsoft servers and applications. For example:

- **Tiered / Unified Storage:** EMC Symmetrix® V-Max™, DMX, CX, NS, Iomega® storage hardware
- **Virtual Infrastructure:** EMC Proven Solutions for SharePoint using Hyper-V and VMware
- **Backup/Recovery/Archive:** EMC NetWorker® Microsoft Module, EMC Replication Manager, EMC Avamar™
- **Business Continuity:** EMC RecoverPoint, SRDF®, MirrorView™
- **Enterprise Content Management:** SharePoint Repository and Content Services, SourceOne™
- **SharePoint Security:** RSA® Secure View, RSA Data Loss Prevention

To learn more about EMC offerings for SharePoint, please visit www.emc.com/sharepoint

Michael Noel is an internationally recognized technology expert, bestselling author, and well known public speaker on a broad range of IT topics. He has authored several major best-selling industry books that have been translated into over a dozen languages worldwide. Significant titles include SharePoint 2007 Unleashed, Windows Server 2008 Unleashed, ISA Server 2006 Unleashed, Exchange Server 2007 Unleashed, SharePoint 2003 Unleashed, and many more. Currently a partner at Convergent Computing in the San Francisco Bay Area, Michael's writings and extensive public speaking experience across six continents leverage his real-world expertise helping organizations realize business value from Information Technology infrastructure.



EMC², EMC, EMC Proven, EMC SourceOne, MirrorView, NetWorker, SRDF, Symmetrix, and where information lives are registered trademarks or trademarks of EMC Corporation in the United States and other countries. Iomega is a registered trademark of Iomega Corporation. RSA is a registered trademark of RSA, The Security Division of EMC. VMware is a registered trademark of VMware, Inc. All other trademarks used herein are the property of their respective owners.