Metalogix

metalogix.com

WHITE PAPER

Optimize SharePoint Storage to Cut Costs and Dramatically Improve Performance

Michael Noel
Published: 2012

CONTENTS

UNDERSTANDING THE CURRENT STATE OF SHAREPOINT AND OUT OF CONTROL STORAGE GROWTH

SharePoint growth continues to outstrip growth in the overall content management market by a wide margin. According to Microsoft, if SharePoint were a standalone company, it would now rank among the 50 largest software companies in the world. Why? Organizations have found it an ideal document management, collaboration, intranet/extranet and web platform and SharePoint 2010's enhanced enterprise metadata structure and document controls have moved it from a departmental solution to a true alternative to legacy enterprise content management systems.

This growth in popularity has resulted in an explosion in both the amount of content created and uploaded into SharePoint and in the storage required to hold it. Nearly all of that content is documents and other types of unstructured data, which in database terminology are called Binary Large Objects (BLOBs). SharePoint stores documents in Microsoft SQL Server databases, an architecture that was never designed to efficiently store unstructured data. Couple the rapid expansion of content stored in SharePoint with an architecture that was designed to store relational (tabular) data rather than documents and you can quickly encounter performance and reliability issues.

Fortunately, remote BLOB storage (RBS) mitigates this issue by moving BLOBs outside SQL and replacing them with pointers to the actual content, without impacting SharePoint features or functionality. This approach can save money because RBS solutions, such as Metalogix StoragePoint, can store BLOBs on less expensive storage tiers and reduce overall storage requirements by eliminating the duplicate BLOBS that are common in SharePoint content libraries. This paper covers RBS in greater detail later, as well as alternative RBS solutions and other storage optimization strategies. The point is that SharePoint success requires storage tools and policies that can keep up with the rate of growth in SharePoint content and users.

A CLOSER LOOK AT THE PERFORMANCE CONSEQUENCES OF RAPID SHAREPOINT GROWTH

While SharePoint content growth is expected, as mentioned above, it can lead to a decrease in SharePoint performance if not handled correctly. Because more data requires more storage, environments that grow too large can quickly reduce performance. As we'll discuss, storage requirements actually grow disproportionately faster than the growth in SharePoint content itself. The result: pages load too slowly, searches bog down and adoption of the platform can slow as users tire of waiting for content to refresh. Microsoft recognizes this and has issued minimum and optimal storage guidelines for hardware performance and content database size.

Storage-related performance issues occur because SharePoint and its SQL underpinnings require a high performance storage infrastructure. Specifically, Microsoft recommends a minimum of 0.25 Input/Output Operations per Second (IOPS) per GB of content and 2.0 IOPS per GB of content for ideal performance. This means that your disk infrastructure must be robust enough to support a fairly high IOPS total, and the larger the data store, the higher the requirements. Table 1 illustrates how many IOPS are required for various data sizes.

Table 1: LOPS Required For Various Data Sizes

Total Content Database Size	IOPS required for minimum performance	IOPS required for optimal performance
100GB	25	200
500GB	125	1000
1TB	250	2000
5TB	3750	10,000

Considering that many disk drives provide only 150-300 IOPS each, SharePoint administrators may need to buy far more storage capacity than they require to achieve the recommended IOPS performance. This disparity will likely grow because the average capacity of a drive is increasing without an equivalent increase in the amount of IOPS.

Table 2 illustrates this challenge by listing disk architecture options that would deliver approximately 1000 total IOPS. As you can see, planning the disk infrastructure depends on the total number of IOPS per disk, the RAID type chosen, and the number of disks in the drive set.

Table 2: Example Disk Volumes to Achieve 1000 IOPS

Drive Type	IOPS per Disk	RAID	Capacity (GB) per Disk	# Disks	Usable Capacity (GB)	Max IOPS
7.2k RPM SATA	90	RAID 0+1	1024	14	7168	1008
10k RPM SATA	130	RAID 0+1	1024	10	5120	1040
10k RPM SAS	140	RAID 0+1	1024	10	5120	1120
15k RPM SAS	180	RAID 0+1	1024	8	4096	1152
7.2k RPM SATA	90	RAID 5	512	20	9216	1026
10k RPM SATA	130	RAID 5	512	14	6144	1037.4
10k RPM SAS	140	RAID 5	512	14	6144	1117.2
15k RPM SAS	180	RAID 5	512	10	4096	1026

UNDERSTANDING THE STORAGE CONCERNS OF RAPID SHAREPOINT GROWTH

In addition to the storage performance issues of SharePoint data growth, rapid growth can also have a disproportionate impact on storage capacity. To understand why this happens, it is first important to understand how information is stored in SharePoint.

SharePoint runs as a three-tiered application, as shown in Figure 1. The first tier, the web tier, is the tier that runs Internet Information Services (IIS) and serves web content to clients. This tier connects via HTTP or HTTPS and is often load-balanced for availability.

The second tier is the service application tier, which runs services that are consumed by systems within a SharePoint environment, both within the immediate farm and sometimes outside of the farm itself. Service application examples include Excel Services, PerformancePoint, the Managed Metadata Service, and the User Profile Synchronization Service.

The third tier of SharePoint is the data tier, where the information stored within SharePoint is kept. By default, all of the documents, images, and other files stored within SharePoint are stored in one or more unique content databases at this data tier. These content databases are simply SQL databases that are used for storage of the content that is created and consumed in the SharePoint environment. This includes all content within SharePoint, including both structured content, such as the metadata and contextual information, plus unstructured content such as the actual documents themselves. Many of the service applications store their databases on this tier as well, so that information can be shared between various servers that may run a particular service application.

Figure 1: Three Tiers of SharePoint Architecture

One reason why SharePoint can grow very large very quickly is that every version of every document stored in SharePoint is stored in full within the content database as a BLOB. For example, if 100 versions of a 5MB document are stored within SharePoint, the total amount of space consumed by that single document is 500MB alone. Since document versioning is a major selling point with SharePoint document management, this often leads to a rapidly growing SharePoint environment. It is now common for SharePoint storage tier sizes to exceed a terabyte. In fact, this has become the norm for even mid-sized or smaller organizations. Because of the aforementioned performance issues, as well as the need to house the data, organizations need to pay close

attention to how they architect the data tier and how they manage their data. Thus it's important to understand Microsoft's guidance regarding SharePoint storage. We mentioned Microsoft's IOPs requirements above but that's just one piece of Microsoft's updated guidance. Let's now review the updated guidance in more detail.

MICROSOFT DATABASE AND SITE COLLECTION LIMITATIONS AND GUIDANCE

SharePoint content within a farm is housed within logical groupings of content known as web applications. Each web application in a SharePoint environment can house up to 300 content databases and each content database can house up to 2000 site collections. Site collections are the base site element in SharePoint. However, the downside to this architecture is that a single site collection can only exist in one content database; it cannot span multiple content databases.

SINGLE VS. MULTIPLE SITE COLLECTIONS

In some cases, organizations may choose to keep all or the majority of their content within a single site collection. This practice can lead to very large content database sizes, which can complicate an information governance plan and lead to some of the performance issues previously mentioned. Best practices for SharePoint environments have suggested that content should be distributed across a document management environment in multiple site collections that are stored in multiple content databases, a concept illustrated in Figure 2. This can help improve performance by decreasing the number of rows within the content database, which can have a positive impact on performance. Figure 2 shows a sample organization that deployed multiple site collections for each business unit and distributed their data across content databases in that manner, allowing for smaller overall databases and avoiding situations where all content is stored in a single database.

While this approach favors performance, it may not favor usability, because the default navigation structure is at the site collection level. Usability favors keeping all content related to a particular application within a single site collection. For example, an accounts payable application would be easier to use if it stores invoices and purchase orders in one site collection. A workaround that provides equivalent menu-based navigation across site collections requires the custom development of master pages in SharePoint. For this reason and due to organic growth of the SharePoint environment, many organizations end up with content database sizes that exceed the recommended best practices. However, Microsoft recently increased the recommended database sizes and this new guidance must be considered in the architecture of the data tier.

Figure 2: Distributing Content across Multiple Content Databases

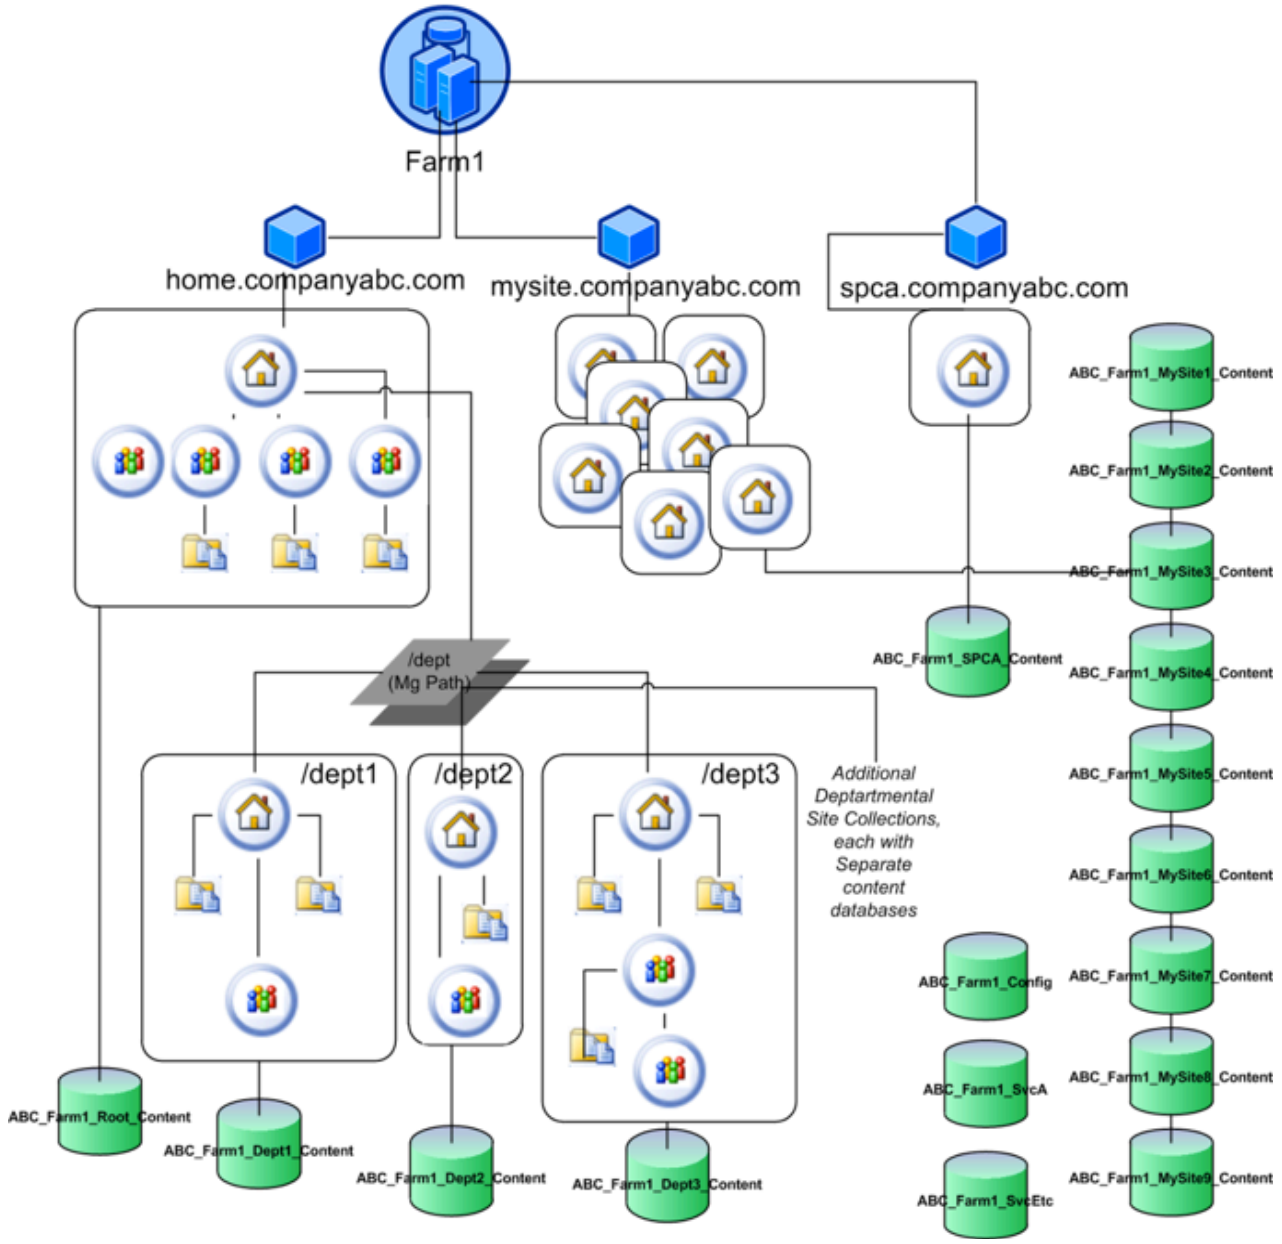

MICROSOFT SHAREPOINT 2010 SERVICE PACK 1 GUIDANCE ON DATABASE SIZES

With the release of Service Pack 1 for SharePoint 2010, Microsoft updated their recommendations for content database sizes. Previous recommendations capped database sizes at 200GB for each content database for collaboration sites and 1TB for document archives. The maximum recommended size has been expanded to a whopping 4TB for normal scenarios and unlimited for records management or archive scenarios. The only remaining limitation is the maximum of 60 million objects in any single content database, a massive amount of items but not an unprecedented number in some larger SharePoint deployments. You can read more about the recommendations in the white paper, "Optimize Your SharePoint 2010 Content Using Microsoft's New Storage Guidance."

Microsoft increased the recommended limits on content databases in Service Pack 1 after reassessing real-world usage — and the resulting IOPS requirements — and adding the ability to recover sites that have been deleted from a newly created 'Site Recycle Bin.' What Microsoft found is that one of the main reasons to restore databases was to recover deleted sites. Since that is no longer a major issue post-SP1, Microsoft considered this — and a more accurate estimate for IOPs performance — and updated their guidance to allow larger databases. This does not necessarily mean that all organizations should immediately go out and create massive content databases. In fact, very large content databases can still have a negative impact on overall functionality and design possibilities. For example, high availability and disaster recovery can be complicated with massive database sizes, because technologies such as SQL Mirroring and Log Shipping do not work well on large databases. Traditional restore techniques can also take too long for massive databases and it can be more difficult to move data around to rebalance the load on SQL Servers when databases are extremely large.

IMPROVING SHAREPOINT STORAGE PERFORMANCE

SharePoint data tier restrictions and high disk IOPS requirement have driven organizations to find a way to improve SharePoint performance without investing a massive amount in their storage infrastructure. Given the cost of high-performance SAN and NAS infrastructure combined with the IOPS requirements for a large document management repository, the storage cost alone can dwarf all other SharePoint implementation costs. Options are needed that increase SharePoint performance without a massive investment at the storage tier.

IMPROVING STORAGE PERFORMANCE BY DIVIDING CONTENT DATABASES INTO MULTIPLE FILES

One simple way to improve the performance of the data tier is to break a content database into pieces. This can be done by creating multiple distinct files for each content database within SQL Management Studio and distributing those files across different storage volumes, as illustrated in Figure 3. Assuming each volume runs on its own unique disk aggregate, you can improve performance at the data tier, resulting in faster page loads.

Figure 3: Distribution of Content Database Files Across Volumes

In Figure 3, a total of four files were created for each database. For example, DB-A is broken into four files distributed across four volumes, as is DB-B. It is also important to do the same for the SQL tempdb file, which is a critical file for SharePoint performance. As illustrated in Figure 3, the rough calculation to determine how many files to create is directly related to the number of physical processors in use on the SQL Server used by SharePoint. In this case, there are four physical processors, which is why the total number of files created is four. You may run into guidance that dictates that the number of files to create should equal the number of processor cores, but this guidance originated with SQL 2000 testing and is not as accurate for today's modern multi-core processors. While there is no perfect equation for this and creating too many files does not necessarily lead to performance issues, the best practice in this case is to distribute by number of physical processors, not the number of cores.

Each of the storage partitions that house SharePoint content should also be write-optimized, to enable them to handle the increased number of write operations that are performed by increasing the number of files. The ideal RAID level for SharePoint content is RAID 0+1, which provides the highest performance and availability options. RAID 5 can be a cost-effective option that results in larger drive sizes, but be cautious about ensuring that the number of IOPS required is maintained as RAID 5 does not provide the same level of performance as RAID 0+1.

IMPROVING STORAGE PERFORMANCE WITH REMOTE BLOB STORAGE (RBS)

As mentioned earlier, another option for improving SharePoint storage is to take advantage of RBS technologies that move documents and files (BLOBs) outside of the content database, thereby keeping the overall size much lower.

To reiterate from page one, SQL Server isn't the best spot to store this unstructured content. SQL Server works better if it is given structured data to work with, such as metadata and the context of files. Using RBS, organizations can extract the BLOBs from the SQL databases and store them on alternative storage, as shown in Figure 4.

Figure 4: Understanding Remote BLOB Storage (RBS)

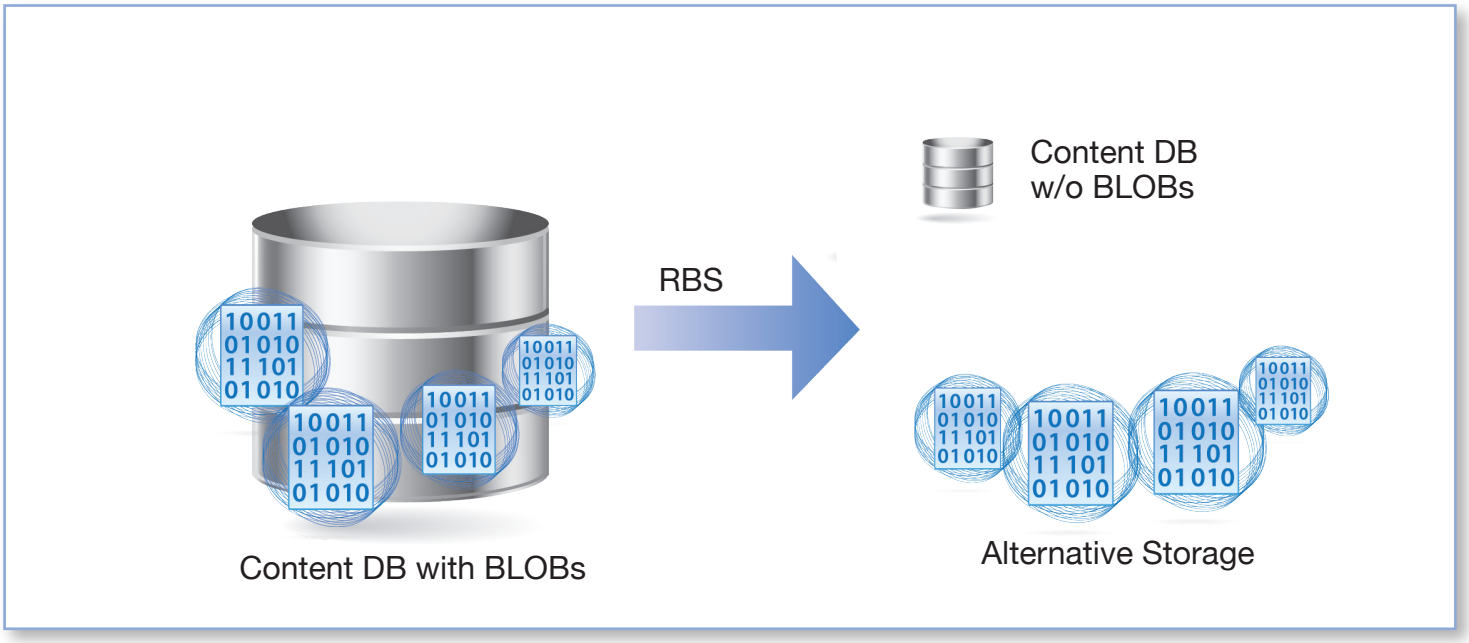

By storing the content outside of the SQL content database, the size of the database itself can be reduced by as much as 95 percent. That's because the overwhelming majority of content in the average SharePoint content database is BLOBs. This enables a wide range of money-saving and/or performance enhancing options for SharePoint, such as the ability to tier or segment the storage, or the ability to create complex archiving policies.

One additional note about RBS: it isn't the only externalization option for SharePoint. RBS's predecessor is known as External BLOB Storage (EBS). Microsoft introduced EBS support in SharePoint 2007/WSS 3.0 SP1 and it remains supported by SharePoint 2010.

Microsoft provides a sample RBS solution, called the RBS FILESTREAM provider, that enables a simple RBS implementation for SharePoint, but there are some significant limitations to the use of the FILESTREAM provider. Microsoft itself does not recommend FILESTREAM for enterprise deployments. It's offered as a feature pack of sample code to give customers an example of the potential benefits of RBS. Its limitations include:

- Available only for SQL standard edition. Must write custom code to use with Enterprise Edition.

- Command line interface.
- No tiered storage support: Requires expensive Tier 1 storage for BLOBs, negating the substantial cost savings that can be achieved with RBS.
- Lack of support for remote and Content Addressable Storage (CAS). BLOB storage must be physically attached to the SQL server. Network Attached Storage (NAS), CAS (for speed and compliance) and cloud storage are not supported.
- No outbound BLOB cache to accelerate externalization and make RBS transparent to users.
- No filtering of storage endpoints by lists or content type to leverage lower-cost storage tiers.
- No compression.
- No encryption.
- No multi-threading for RBS Maintainer.
- No orphaned BLOB RBS Maintainer.
- Does not bypass SQL Server for BLOB Processing. It pulls the BLOB out and redirects right back to SQL Server using FILESTREAM Column Type.
- Backing out of FILESTREAM is difficult and time consuming.

For these reasons, it is typically recommended to use a third-party tool, such as Metalogix StoragePoint, which can leverage the full potential of RBS or EBS.

StoragePoint offers the following advantages over FILESTREAM:

- Works with SQL Enterprise Edition.
- May access through the familiar SharePoint Central Administration interface.
- Provides full Hierarchical Storage Management (HSM) support to leverage lower-cost storage tiers to dramatically reduce storage costs and to support archiving scenarios. May define multiple endpoints per content database to balance performance and cost.
- Supports remote storage: Enables cost-savings benefits of NAS or cloud storage.
- Supports Content Addressable Storage (CAS) for performance and regulatory compliance.
- Encrypts data (256-bit) for security and compliance.
- Filters endpoints by list, content type, aging rules, versioning rules or by metadata change rules.
- Compresses data.
- Caches outbound BLOBS targeted to remote storage to eliminate user delays.
- Supports PowerShell to automate activities using business rules.
- Supports shallow copies: Keeps data on file servers but can access it in SharePoint via RBS to enhance performance. Enables data de-duplication options that are not possible when BLOBs are stored within the content databases. Since 90 percent of one version of a document may be identical to

the last version in SharePoint, by storing the nearly identical BLOBs on a de-duplicated SAN volume, the overall space taken up by the files themselves can be reduced.

CONCLUSION

Because SharePoint requires a minimum of 0.25 IOPS per GB and recommends 2.0 IOPS per GB of content, many organizations with rapidly growing SharePoint environments can find themselves trapped between limiting adoption of SharePoint or requiring a significant and perhaps unfeasible investment in storage. Fortunately for these organizations, there are methods that can be used to rein in storage growth while maintaining performance. These methods include simple techniques such as creating multiple files for each SharePoint content database or advanced options such as RBS, which can shrink SharePoint content databases by 95 percent. An advanced RBS solution can dramatically reduce storage costs while improving SharePoint performance, enhance data security and enable robust archiving schemes. Organizations planning for growth within their SharePoint environment should consider these options or run the risk of storage performance issues eventually sabotaging their plans for SharePoint.

To find out more about the Metalogix StoragePoint RBS solution,
visit http://www.metalogix.com/Products/StoragePoint.aspx.

ABOUT THE AUTHOR

Michael Noel – Partner/Principal Consultant – Convergent Computing

Michael Noel (MCITP) is an internationally recognized technology expert, bestselling author, and well known public speaker on a broad range of IT topics. He has authored multiple major industry books that have been translated into over a dozen languages worldwide. Significant titles include SharePoint 2010 Unleashed, Exchange Server 2010 Unleashed, Windows Server 2008 R2 Unleashed, ISA Server 2006 Unleashed, and many more. Currently a partner at Convergent Computing (www.cco.com) 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.

ABOUT METALOGIX

Metalogix is the trusted provider of innovative content lifecycle management solutions for Microsoft SharePoint, Exchange and Cloud platforms. We deliver high-performance solutions to scale and cost-effectively manage, migrate, store, archive and protect enterprise content. Metalogix provides global support to thousands of customers and strategic partners and is a Microsoft Gold Partner, a managed partner in Microsoft's High Potential ISV Group and GSA provider. Metalogix is a privately held company backed by Insight Venture Partners and Bessemer Venture Partners.

METALOGIX
5335 Wisconsin Ave NW, Suite 510, Washington DC 20015
sales@metalogix.com | www.metalogix.com | 1.877.450.8667