

SharePoint 2016 Edition

The SharePoint MVP Guide to Optimizing Storage and Performance



Written by Michael Noel,
SharePoint MVP

Table of Contents

- 3 | Understanding the Current State of Share-Point and Out-of-Control Storage Growth
- 9 | Microsoft Guidance on Storage Sizes and Limitations
- 12 | A More Efficient Route to Improving SharePoint Storage Performance

- 16 | Conclusion
- 17 | About Metalogix

Understanding the Current State Of SharePoint and Out-Of-Control Storage Growth

SharePoint is an invaluable tool for document management and collaboration and the ideal platform for web-based environments, including intranets and extranets. As such its popularity has grown exponentially, 80 percent of Fortune 500 companies use SharePoint and the latest data suggests that Microsoft adds

more than 20,000 users each day. SharePoint's document management capabilities are of particular interest, with enhanced document and information controls that allow for a better experience. In fact, according to a survey by Collaboris , users cite document management as the major use of SharePoint within their organizations.

With this growth, however, comes storage challenges. Despite recent improvements in the way SharePoint stores documents, SQL Server remains a poor choice for storing the unstructured content like documents, images and videos that dominates SharePoint. If not addressed, these storage limitations can lead to performance issues—making it even more important to implement controls on SharePoint content growth and devise a plan to manage it.

This whitepaper will explore SharePoint storage constraints and suggest steps that organizations can take to limit the impact that content growth has on their mission-critical environments.

How Rapid SharePoint Growth Inhibits Performance

Simply put, the larger the content – rich media, large repositories of content, etc. – the higher the performance requirement. Subsequently, environments that grow too large have a negative impact on overall performance. Pages can take longer to load and users quickly tire of waiting for content to refresh, which, in turn, has a negative

effect on platform adoption.

These storage concerns stem from the fact that SharePoint demands a high degree of performance from its storage infrastructure. These performance metrics are directly related to the total amount of space that is taken up by the databases. Specifically, Microsoft has

determined that SharePoint content databases typically use between 0.05 IOPS/GB and 0.2 IOPS/GB. For optimum performance, Microsoft recommends 0.5 IOPS/GB. This means that any disk infrastructure must be robust enough to support a fairly high IOPS total. The bigger the data store, the higher the requirements.

Table 1 illustrates how many IOPS are required for various content database sizes.

Total Content Database Size	IOPS required for minimum performance (0.2 IOPS/GB)	IOPS required for optimal performance (0.5 IOPS/GB)
500GB	100	250
1TB	200	500
2TB	400	1000
5TB	1,000	2,500
20TB	4,000	10,000

Table 1: IOPS Required for Various Data Sizes

Considering that many disk drives can provide 150-300 IOPS apiece, it is obvious that a large number of disk spindles are required to support a growing environment. If the number of IOPS doesn't match the growth, there will be a gradual degradation in performance over time. In order to keep up with the growth, administrators may need

to over-provision disks to maintain the recommended number of IOPS, ending up with a larger number of disks with unused capacity on each. This is especially true today as the average size of a disk is increasing without an equivalent matching rise in the amount of IOPS capable per each disk.

For example, Table 2 shows sample disk architecture options that would provide approximately 1000 IOPS. Considering that the maximum recommended IOPS required for 2TB of storage is 1000 IOPS, many of these configurations result in a large amount of wasted space: up to 7TB in total. So we can see that planning the disk infrastructure is highly dependent on the total number of IOPS per-disk, the RAID chosen, and the number of disks in the drive set.

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

Table 2: Example Disk Volumes to Achieve 1000 IOPS

How Rapid SharePoint Growth Challenges Impact Storage

While growing SharePoint environments can inhibit performance, rapid growth can also have an immediate impact on storage. To understand why, let's look at 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 up web content to clients directly. This is the tier that the clients connect directly to via HTTP or HTTPS and is often load-balanced for availability.

Figure 1: Three Tiers of SharePoint Architecture



The second tier is a service application tier, which is used to run the multiple services that are consumed by numerous systems within a SharePoint environment, both within the immediate farm and sometimes outside of the farm itself. These service applications include Search, Excel Services, PerformancePoint, the Managed Metadata Service, and the User Profile Synchronization Service.

The third SharePoint tier is the data tier, where all the information that is shared and managed within SharePoint is kept. With the exception of search indexes, this information is stored within Microsoft SQL Server databases. By default, this includes all content within SharePoint, both structured, such as the metadata and contextual information, and unstructured, such as the actual documents themselves. These unstructured objects, known as

Binary Large Objects (BLOBs), degrade performance and create SharePoint database sprawl, which is hard, time consuming and expensive to manage.

These content databases are simply SQL databases used for storage of the content that is created and consumed in the SharePoint environment.

Shredded Storage in SharePoint 2016 Attempts to Put Things Right

Prior to the release of SharePoint 2013, SharePoint stored every version of every file as a full-sized BLOB, which led to significant database size growth. SharePoint 2016, on the other hand, introduced a concept known as “Shredded Storage”, in which only the first version of a document is stored as a full-sized BLOB, and each subsequent version stores only the changes made between versions. This reduces the steepness of the total storage growth curve for SharePoint and also cuts down on data redundancy.

One disadvantage to shredded storage, however, is the fact that there is a performance cost associated with

reassembling the BLOB versions. Natively, SharePoint will always serve up BLOBs directly from SQL Server, a process that can be up to two-times slower than accessing data from other storage mechanisms. In addition, shredded storage cannot be applied retroactively. In other words, documents that are migrated from SharePoint maintain their full BLOB sizes for older versions.

Despite the introduction of shredded storage, environments that measure their SharePoint storage tier sizes in terabytes are still common. 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 must pay close attention to how they architect the data tier and how they manage their data. It is important to understand Microsoft's official and updated guidance on this topic as well which we'll look at next.

Microsoft Guidance on Storage Sizes and Limitations

The guidance from Microsoft on SharePoint storage restrictions and size limitations can be confusing and if complicated by the fact that there are two defined levels. Boundaries which are absolute limits that cannot be exceeded by design, or Thresholds which are those that have a default value that cannot be exceeded unless the value is modified.

Current Limitations and Restrictions

There are a few key limitations and restrictions of SharePoint components that factor into storage design for SharePoint 2016. Within a SharePoint farm content is housed within logical groupings known as web applications. Each SharePoint farm can house up to 500 content databases. While there is room for expansion at the content database tier, (each content database can house up to 250,000 non-personal site collections) a single site collection

can exist only in one content database - it cannot span multiple content databases.

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 complicates storage management planning and can lead to the performance issues

previously indicated. Indeed, best practices for SharePoint environments generally dictate that content should be distributed across a document management environment in multiple site collections that are stored in multiple content databases, as illustrated in Figure 2. This can help improve performance by decreasing the density of 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 concept is ideal, it is not always followed and comes with some challenges. For example, Microsoft makes the default navigation structure at the site collection level, so any type of menu-based navigation across site collections must be created by the custom development of the master pages in SharePoint. For this and other reasons related to natural organic growth of the SharePoint environment, many organizations may end up with very large content database sizes. This becomes problematic given Microsoft's recently updated guidance on database sizes and the role this plays in the architecture of the data tier.

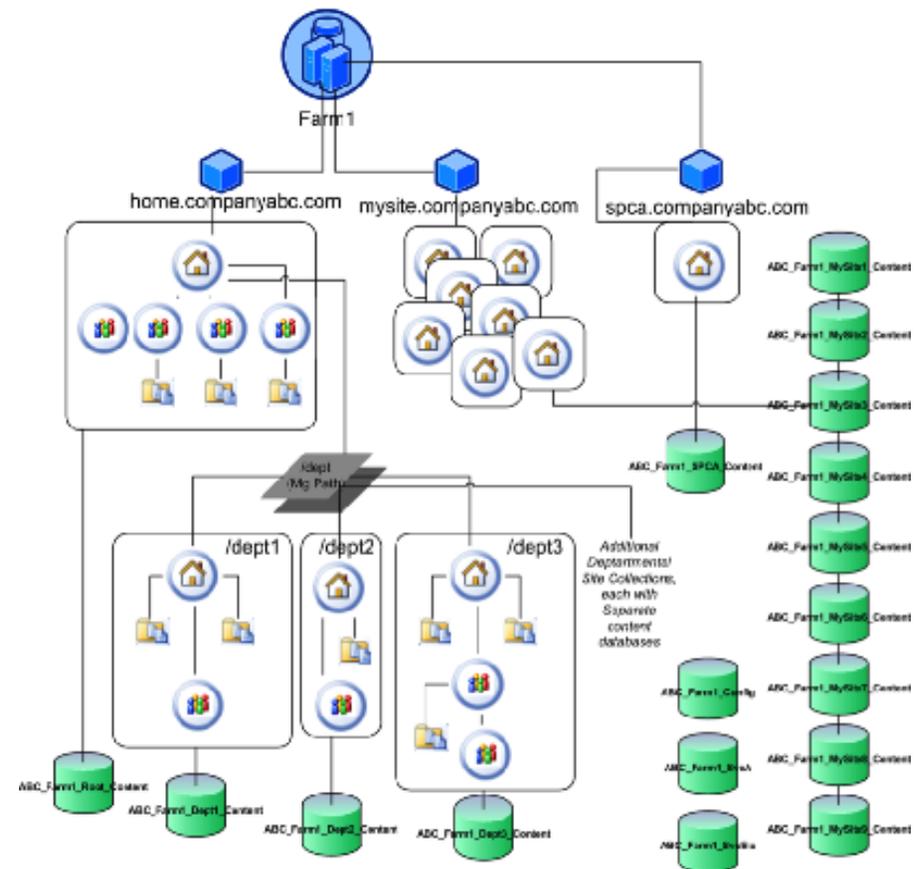


Figure 2: Distributing Content across Multiple Content Databases

Microsoft Guidance on Content Database Sizes

With the release of Service Pack 1 for SharePoint Server 2010 and SharePoint Foundation 2010, Microsoft updated its recommendations for content database sizes. These recommendations also carry over to SharePoint 2016 environments. 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 database sizes for records management or archive scenarios. The only hard limitation then becomes the restriction to a maximum of 60 million objects in any single content database. This represents a large amount of items, but is not an unprecedented number in some larger SharePoint deployments.

The change in Microsoft's upper recommendation limits for content database sizes in Service Pack 1 was a result of usage reassessment and the addition of 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 that administrators restore databases is to recover deleted sites. Since that is no longer a major issue post-2010 SP1, Microsoft updated their guidance and started to allow for massive database sizes. It's important to note that the previous 200GB size "limitation" was not an actual limitation at all, and SQL would allow databases that were much larger than that. Organizations with massive pre-2010 SP1 databases, however, found themselves in a tricky spot and were

forced to recover an entire database simply to restore a single deleted site.

This doesn't mean that we advocate that organizations should immediately go out and create massive content databases. In fact, very large content databases can still have a negative impact on the overall functionality and the design possibilities available. For example, high availability and disaster recovery can be complicated by massive database sizes, since technologies such as SQL Mirroring and Log Shipping do not work well on larger databases. Traditional restore techniques can also take too long in these scenarios. It can also be more difficult to move data around to rebalance the load on SQL Servers when databases are extremely large.

A More Efficient Route to Improving SharePoint Storage Performance

The architectural limitations of SharePoint's data tier coupled with its high disk IOPS requirements is driving organizations to seek new ways of improving SharePoint performance without investing significant funds in their storage infrastructure.

SharePoint storage can be expensive business. Organizations must factor in the cost of high-performance SAN and NAS infrastructures combined with the IOPS requirements of dealing with a large document management repository. These costs alone can dwarf the other budget items required to implement SharePoint. Clearly, organizations need a better way to manage SharePoint performance without the massive investment at the storage tier.

Improving Storage Performance by Dividing Content Databases Into Multiple Files

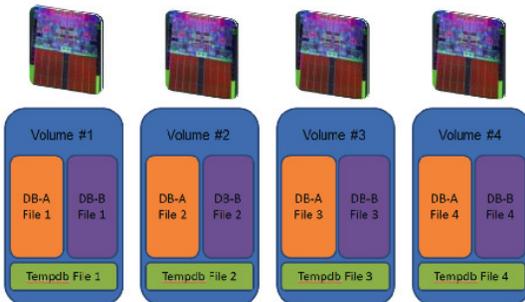


Figure 3: Distribution of Content Database Files across Volumes

One simple way to improve the performance of the data tier is to break a content database into multiple pieces. This can be done by creating multiple distinct files for each content database within SQL Management Studio. These files can then be distributed across different storage volumes, as illustrated in Figure 3. If the database files are distributed across separate disk aggregates, better performance is achieved at the data tier resulting in faster page loads.

Administrators don't need to span multiple disk aggregates to get better performance, since multiple files for each database can result in better parallelization of the database traffic.

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. In addition, it is 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 used 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 created 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. There is no perfect equation for this process and, since 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 allow them to handle the increased number of write operations that are performed. The ideal RAID level for SharePoint content is RAID 0+1, which ensures the highest performance and availability options. RAID 5 can be a cost-effective option that results in larger drive sizes, but administrators should 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)

One possible option for improving SharePoint storage is to take advantage of technologies that allow the actual documents and files to be stored outside of the content database, thereby keeping the overall size much lower. This is known as Remote BLOB Storage (RBS). As discussed earlier, a BLOB is the format in which all documents and files are stored within the content databases. This type of data is known as unstructured data and SQL Server has not traditionally been the best place to store this type of content. SQL Server works better if it is given structured data to work with, such as metadata and the context of files. Using a process such as RBS,

organizations can extract BLOBs from the SQL databases and store them on alternative storage, as shown in Figure 4. This approach provides the benefit of managing content from

within SharePoint, without incurring the performance and storage hit associated with physically storing BLOBs in a SQL Server database.

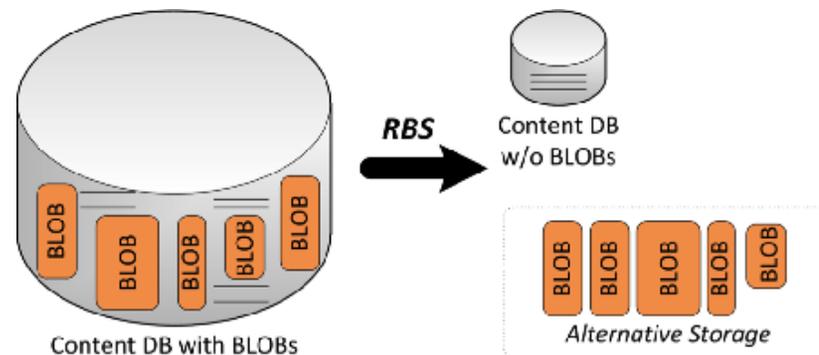


Figure 4: Understanding Remote BLOB Storage (RBS)

Since the majority of content in the average SharePoint content database is typically composed of BLOB files, by storing content outside of the SQL content database, the size of the database itself can be reduced up to 95 percent. This opens up a myriad of different architectural options for SharePoint, such as the ability to tier or segment the storage, or the ability to create complex archiving policies.

For this reason, it is recommended to use a third-party tool that can take advantage of the RBS features. These tools provide additional functionality well beyond that provided by the FILESTREAM provider, which was released by Microsoft as a sample provider, not as a fully baked solution. Additional functionality provided by third-party tools include using RBS to take advantage of storage tiers that may be remote, such as cloud storage, slower and cheaper SATA disk volumes, and even keeping the data on file servers while accessing it in SharePoint via RBS, a concept known as a Shallow File Copy.

Out-of-the-box, Microsoft provides a tool – “FILESTREAM provider” – that allows for RBS to be implemented with SharePoint, but there are some significant limitations to FILESTREAM provider:

- It lacks basic features for enterprise deployments, including:
 - No user interface
 - Lack of support for remote storage
 - No multi-threading for garbage collection
- RBS 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

An additional advantage of using RBS to store the BLOBs on an alternate storage location is that it opens up the environment for the use of data de-duplication options that were previously not accessible when the BLOBs were stored within the content databases. Many SAN vendors offer de-duplication at the block level for content within their flat file volumes, so if data is repeated multiple times, the repetitive part of

that data is only stored once. Since one version of a document may be 90 percent similar to the last version in SharePoint, by storing the nearly identical BLOBs on the de-duplicated SAN volume, the overall space taken up by the files themselves can be reduced. Check first with your SAN vendor to see if this is supported, but it is an option that provides a way to deal with large document management environments.

Conclusion

Because of SharePoint's robust disk requirements of up to 0.5 IOPS per GB, organizations with rapidly growing SharePoint environments are trapped between limiting SharePoint adoption and use or needing a significant investment in SharePoint specific storage. Fortunately for these organizations, there are methods that can rein-in storage growth while maintaining performance.

These methods include simple techniques such as creating multiple files for each SharePoint content database, but they can also include advanced options such as RBS, which drastically reduces the size of SharePoint content databases, allowing them to maintain their speed requirements while storing the BLOBs in an alternative, cheaper location.

Organizations planning for growth within their SharePoint environment should also consider these options—or run the risk of storage performance issues eventually sabotaging their plans for SharePoint.

About Metalogix

Metalogix is the premier provider of unified management software to migrate, manage and secure content across enterprise collaboration platforms in the cloud and on-premises.

Over 20,000 clients trust Metalogix to optimize the availability, performance, and security of their content across the collaboration lifecycle.



www.quest.com

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats and regulatory requirements. Our portfolio includes solutions for database management, data protection, unified endpoint management, identity and access management and Microsoft platform management. For more information about the Metalogix integration, please visit www.quest.com/quest-acquires-metalogix.