AvePoint

Best Practices for Microsoft® SharePoint® Backup and Restore

White Paper

Cesar Coba
SVP of Customer Success, AvePoint

Mary Leigh Mackie
CMO, AvePoint

John Hodges
SVP, Product Strategy, AvePoint

Last revised: May 2021

AvePoint DocAve

Table of Contents

Executive Summary 3

Evaluating Your Disaster Recovery Strategy 4

- Anchoring a Successful SharePoint Backup Plan 4
  - Reducing Backup Windows 4
  - Satisfying Recovery Time, Recovery Point, and Recovery Level Objectives 5
  - Providing Comprehensive Support of Farm-Level Components 7
- Determine What Needs to be Protected 7
  - Core Content 7
  - Platform Components 8
  - Combine Assets into an Appropriate Backup Strategy 8
- Analyzing SharePoint's Taxonomy 8
  - Determine the Characteristics of Core Content 9
- Defining Your Service Level Agreements 10
  - Approach 1: Not All Data is Created Equal 11
  - Approach 2: Aggressive RPOs for all 12
  - Stakeholders 13
  - Recovery Scenarios 13
  - Data Integrity 13
  - Measurability 14
  - Performance Best Practices for SharePoint Backup and Recovery 14

Deploying DocAve Backup and Restore 18

- Overview of the DocAve System Architecture 18
- Platform and Granular Backup and Restore 19
  - Platform Backup and Restore 19
  - Granular Backup and Restore 19
  - DocAve Manager 20
  - DocAve Media Service 20
  - DocAve Agent 21
- Considerations for Large-Scale Deployments 21
  - Plan for SharePoint Boundaries and Performance Guidelines 21
  - Limiting Backup Sizes per Backup Plan 22
  - Utilize Hardware Snapshot Technology 22
  - Utilize Multiple Media Services 22
  - Use Load-Balanced Agent Groups in Granular and Platform Backups 23

Optimizing Your Backup Plans 24

- Evaluating Your Storage Requirements 24

Estimating Storage Space for DocAve Backups 25
Setting up Business-Aware Backup Plans 26
Granularity of Backups 26
Security Trimmed Access to Backup Data 27
Plan for Probable Recovery Scenarios 27
Protecting Farm-Level Components 29
Performing a Platform Backup in DocAve 29
Selecting DocAve InstaMount 29
Coexisting with SQL Server Backups 30
Identifying and Selecting Components to Protect 30
Determining Restore Granularity Level 31
Platform Restore Processes 31
About AvePoint 33
Notices and Copyright Information 34

Executive Summary

Organizations of all sizes are rapidly adopting Microsoft® SharePoint® and Office 365 as a standard platform for their online collaboration, portal, enterprise content management, and other mission-critical initiatives. As end-users increasingly utilize SharePoint and Office 365 for their day-to-day business activities, organizations are confronted with the consequences of exponential growth in the volume of business-critical data residing in the platform amidst the overall data footprint. As a result, it has become crucial for today's competitive organizations to vigilantly prepare for unplanned disasters with robust data protection and recovery solutions.

With a complex and distributed SharePoint hybrid deployment model and increasingly constrained IT resources, organizations need an efficient and comprehensive disaster recovery solution to protect the wide variety of data and components that constitute their SharePoint farms and Office 365. This solution must satisfy the most stringent business requirements and service level agreements (SLAs), and have the capacity to scale effectively to maintain performance as the SharePoint footprint expands.

DocAve® 6 Backup and Restore for SharePoint and Cloud Backup for Microsoft Office 365 addresses these challenges by providing a granular, full-fidelity, item-level through platform-level backup and disaster recovery solution for Office 365, SharePoint 2010, SharePoint 2013, SharePoint 2016, SharePoint 2019, and SharePoint Online environments. In the coming pages, we will outline the planning, guidelines, and implementation considerations for SharePoint data protection and disaster recovery, and briefly review how AvePoint backup solutions are able to execute these proven practices. This document is intended to aid IT administrators responsible for managing SharePoint hybrid deployments in planning and implementing a comprehensive, reliable, and efficient data protection strategy appropriate to their organizational needs.

Evaluating Your Disaster Recovery Strategy

Successful SharePoint hybrid deployments require effective management, vigorous protection, diligent compliance protocols, and continuous availability. As companies increasingly rely on SharePoint to act as the sole point of access for enterprise-wide data and content, these requirements become crucial catalysts for optimal production and minimized exposure to costly downtime and data loss. The fundamental question for every organization operating in today's 24/7 economy is: How prepared are you for interruptions to your mission-critical SharePoint environment? Success in today's demanding business environment requires organizations to be properly equipped to maintain uninterrupted access to critical digital assets.

Anchoring a Successful SharePoint Backup Plan

Backup and recovery operations can be very consuming to server resources, causing lower performance levels while operations are still in play. With that said, how confident are you that your current disaster recovery strategy can protect your SharePoint environment? How much does your strategy affect the current performance of your environment? A comprehensive SharePoint disaster recovery strategy for optimized platform performance must seek to achieve the following goals: reducing backup windows, satisfying recovery time and recovery point objectives, and providing comprehensive support of farm-level components.

> *How prepared are you for interruptions to your mission-critical SharePoint*

Reducing Backup Windows

With an ever-growing amount of data stored within the SharePoint SQL Server content databases and Microsoft data centers, it is an ongoing challenge to ensure that backups complete quickly. Long-running backups often have serious ramifications, including performance loss, inefficient use of system resources, and increased likelihood of data loss as fewer recovery points are taken. Additionally, because open files residing in system memory will likely be excluded from backups, elevated levels of data inconsistency are a consequence of long-running backup processes. In order to reduce backup windows, there are several tactics that can be implemented:

- Increase backup speed with hardware-based Microsoft Volume Shadow Copy Services (VSS) snapshots
- Reduce the volume of data being backed up
- Adjust backup frequency based on data criticality and usage
- Perform backups incrementally
- Execute backup jobs based on a more business-appropriate schedule

The most business-appropriate strategy for a given organization generally includes some combination of all these tactics.

One easy way to reduce backup windows is to use faster backup technology provided by the hardware infrastructure supporting the SharePoint environment. Several advanced hardware solutions are compatible with VSS, which can create an application-consistent backup of entire databases, servers, and other data stored on a given volume. The upside of VSS is that it creates shadow copies very quickly because it is only writing the changes to disk. The disadvantage is that in order to fully restore the information, the original data must still be available, which is a challenge when facing disaster recovery situations where the disk may not be usable. Without the original data, the shadow copy is incomplete and cannot be used. Any backup strategy should incorporate tools that integrate or utilize hardware-based VSS technology to ensure farm-level backups are performed as quickly as possible.

One of the most effective ways to reduce backup windows is to reduce the amount of SharePoint data included in a given backup. While ensuring that digital assets are protected, organizations must determine the business-criticality of each specific site, list, and/or document library within the deployment. When the criticality of each of these datasets has been determined, appropriate levels of protection (e.g. frequency of backup executions) can be prescribed. This tiering process empowers organizations to appropriately differentiate their data so critical content is backed up more frequently and data of relatively less business importance will be backed up less frequently.

Another tactic for minimizing backup windows is to perform incremental backups where only the content that has been updated between backup intervals will be selected for backup, drastically reducing the size of the backup file created. This tactic ensures data protection while reducing execution time, consequently allowing for shorter backup windows. If organizations are using this tactic in conjunction with SQL log shipping, it is important to consider the impact incremental backups will have on the overall disaster recovery scenario.

Another option is to use appropriate scheduling. When backups are executed during off-production hours, network resources can be exclusively dedicated to the task. In contrast, backups performed during production must share resources with end-users and other infrastructure processes. It is also possible to schedule backup processes during production hours with the appropriate choice of a hardware snapshot solution, or one that does not lock content during the backup process. A necessary component of this tactic is the ability to schedule backups; therefore, a robust data protection strategy should include tools with appropriate scheduling capabilities.

Satisfying Recovery Time, Recovery Point, and Recovery Level Objectives

One of the primary objectives of data protection and disaster recovery planning is to mitigate disruptions to business continuity. When analyzing tactics to ensure business continuity,

recovery of lost or corrupted data is the fundamental imperative. The key concern is to get the data back to avoid disruptions to business productivity. Interrupted access to content residing on SharePoint and Office 365 can–and usually will–have serious implications for revenue generation, end-user productivity, and other bottom-line business objectives.

Three recognized measures with regard to data recovery are:

- **Recovery Time Objective (RTO)** – Duration of time that a business process must be restored after a disaster or data loss in order to avoid unacceptable consequences associated with a break in business continuity.
- **Recovery Point Objective (RPO)** – Acceptable amount of data loss measured in time.
- **Recovery Level Objective (RLO)** – Granularity level that you need to recover data from (e.g. the whole farm, Web application, site collection, OneDrive for Business, site, list/library, or item level).

In terms of the scope of recovery, not all data needs to be recovered in the event of a failure in order to meet business service level agreements (SLAs).

Occasionally, data loss is the result of catastrophic failure, such as a fire, flood, or hardware malfunction. For these occurrences, it is critical for administrators to maintain a platform-level recovery strategy. More often, content loss is a result of accidental deletion or discrete corruption and can be resolved without a full-system recovery. In these instances, SharePoint administrators can minimize recovery time and meet a business-appropriate RTO by implementing recovery strategies at the granular level. If only select data needs to be recovered, full-system recoveries derived from a complete backup are inefficient because they take an unnecessarily long time to perform and do not utilize system resources economically. An optimal recovery strategy would allow administrators to target and recover only those lost or corrupted objects directly to production. Establishing the RLO with regard to data recovery is essential to meeting aggressive RTOs.

Alternately, satisfaction of a business-appropriate RPO is a function of backup frequency–including a focus on–the intervals between backups. Whether the situation involves a catastrophic failure or the targeted loss of content, satisfaction of RPO is dependent on the backup interval. Administrators must implement data protection routines that allow for backups at intervals complying with the organization's acceptable data loss thresholds. Depending on business needs, an organization may need to run hourly backups and retrieve them on demand to establish an RPO of one hour.

True data protection includes not only recovery of primary content, but also all associated metadata, version histories, and access permissions. Because SharePoint is largely a collaborative platform, established RPOs, RTOs, and RLOs must include full-fidelity data recovery in order to

be truly effective. Administrators must ensure their data protection strategies provide complete preservation of all content metadata.

Providing Comprehensive Support of Farm-Level Components

A SharePoint farm consists of not only the sites and content that reside in the SQL Server databases, but also additional farm-level elements that must be backed up accordingly to provide comprehensive protection against disasters. These elements include the configuration and central administration databases, Web applications, Shared Services providers, the Search Index, IIS configurations, and other customizations likely residing in locations on the front-end Web servers commonly referred to as the "Hive" for SharePoint. Without properly protecting these components, a complete recovery of a SharePoint environment would require extended down time, complex reconfiguration, and countless manual tasks prone to human error. It is therefore recommended that disaster recovery strategies provide adequate protection for the full spectrum of SharePoint farm-level components.

It is also very important to the overall recovery of SharePoint that the ability to restore these items individually allows organizations to recover very quickly from component failure, reducing the impact of those failures on the overall availability of the service.

Determine What Needs to be Protected

Everything within a SharePoint environment needs to be protected, but no single strategy will serve as an appropriately efficient solution for the entire deployment. Rather, strategies should be applied based on roles, business importance, and means of restoration. The first step in taking these measures is to break down your SharePoint assets into core content and platform components.

	What is it?	SharePoint Assets
Core Content	User-facing content in SharePoint.	Site collections, OneDrive for Business, sites, lists, document libraries, list items, documents, and discussion threads.
Platform Components	Components focused on services offered to SharePoint users.	Configuration and central administration databases, Web applications, Shared Services Applications for SharePoint, Search Index, IIS configurations, deployed solutions, and other customizations located on the front-end Web servers.

Core Content

Core content is stored within a SQL Server content database and is accessible and viewable by the end-user regardless of whether that end-user is an administrator, a site owner, content owner, or a limited-rights user. From the organizational/business perspective, core content

represents the most critical data within SharePoint. It is an organization's sales leads, contact information, financial statements, and other mission-critical data that is stored in the content database. Any loss of core content would likely result in significant business disruption for the organization. It is also pertinent for administrators to consider any binary large object (BLOB) content that has been externalized onto file- or cloud-based storage devices through the use of Microsoft's RBS BLOB APIs, as this content is still used by end-users and must be incorporated into any comprehensive data protection strategy. For more information, please refer to AvePoint's Optimized SharePoint Storage with BLOB Externalization white paper.

Platform Components

Platform components consist of the infrastructure elements required for SharePoint to function properly. These components represent the functional architecture of SharePoint, and are generally not accessible by end-users but support the services you offer to support end-users (including profiles, search, Microsoft Excel calculation services, features, and solutions). The SharePoint administrator is typically responsible for managing and configuring these elements. Loss of any of these components would likely result in platform access disruption.

Combine Assets into an Appropriate Backup Strategy

In developing an appropriate backup strategy, it is recommended to approach these two classes of SharePoint assets as a whole as part of your overall backup strategy for your SharePoint farm and Office 365, but take into consideration that one may be better suited than the other for your specific use case. To put it simply, the protection strategies appropriate for one class might not be so for the other.

In light of the concerns we addressed in the previous section, we can see how relying on only farm-wide snapshot backups–though simple to execute–would not deliver the protection required (as measured in RTO and RPO) for core content at a typical organization. Alternatively, a strategy that can recover core content efficiently might not enable a quick restoration of the platform's complex components and related configurations. Any robust disaster recovery plan should provide optimal support for protecting both the core content as well as the complex platform-level components.

Analyzing SharePoint's Taxonomy

Efficient and successful SharePoint deployments must have a properly designed taxonomy. By implementing an effective content taxonomy, knowledge workers will be better equipped to harness the value of SharePoint to organize an unstructured web of information. An appropriately designed taxonomy is straightforward and relatively easy to implement with the proper planning that includes buy-ins from both the business and IT departments in an organization. Before content sprawl is allowed to occur, the following simple tasks can frame an effective taxonomy:

- Sites and sub-sites should be logically grouped
- Document libraries and lists should be properly categorized
- Information should be targeted to the appropriate group of users

Change is constant in today's business climate, and a well-defined taxonomy in SharePoint serves as an effective framework for organizing the growing volume and changing nature of information within an organization. Not only can the many dimensions of a well-formed taxonomy facilitate information access for end-users, it can also greatly simplify the backup strategy with respect to core content.

Administrators can build granular backup plans more efficiently when they are targeted as subsets of sites, sub-sites, or libraries with content that is logically grouped by taxonomy. By doing so, administrators can appropriately allocate system and storage resources needed to protect different categories of their functional content. SharePoint content that is deemed of high business importance (e.g. Sales Productivity Applications) can be backed up more aggressively, while content of relatively less business importance (e.g. Facilities Manuals) can be backed up less frequently.

We will discuss this concept in further detail later in this paper when we review concrete procedures and tools available to target and execute backups based on criticality.

Determine the Characteristics of Core Content

Constant content growth places ever-evolving burdens on existing backup strategies. Because of this burden, there is no ideal backup strategy. It is therefore important for organizations to perform regular detailed analysis of their SharePoint environments to maintain an accurate understanding of how much and what type of content needs to be protected. In this process, it is important to consider a variety of factors, including:

- **Nature of each content database – Its overall size and the type of content that is stored (e.g. list items, documents, discussion threads, and audio files).**
- **Nature of content externalized onto cloud- and file-based storage – Frequency of access or modification of data, whether the cloud provider offers its own SLA for content availability, and whether the hardware is under its own separate backup plan or schedule (e.g. taking snapshots).**
- **Rate of change of content – Usage/access frequency and the volume of content that gets** generated **weekly or monthly.**

While considering the factors above, it is important to keep in mind that content from a given database can be backed up to different tier levels of storage (most likely based on the age of backup data) depending on the length and cycle for retention of specific content. To implement an optimized backup strategy, content objects must be appraised granularly and require backup tools that allow for granular precision.

Defining Your Service Level Agreements

As we briefly discussed earlier, a viable SharePoint content and platform protection strategy requires the generation of a business-appropriate SLA. The SLA is often the most difficult aspect of protection strategy development, since the organization must solidify terms and conditions–in terms of RTO and RPO–in coordination with management and end-users. A formal document should be prepared that communicates agreed expectations on SharePoint performance. For the administrator, these agreements will **serve as the principles that frame the details of a backup plan and justify the budget and resource allocation requirements for implementation.**

Throughout the process of developing a formal SLA, the administrator must also serve as the reality checker. Management and end-users often have unrealistic expectations of achievable service levels. Under these circumstances, the level of service desired does not match the level of financial and time commitment these parties are willing to invest to achieve the desired level of service. Managers and end-users often expect a guarantee of near-zero data loss, or continuous uptime, without committing the energy and resources needed to realize such guarantees. It is often the role of the administrator to inform management and end-users of the true costs related to a given SLA.

Supporting robust SLAs for business-critical content typically requires more frequent backup schedules, and possibly unique requirements for storage. For example, you need to back up to higher tiered or secure storage locations to support the separation of sensitive or regulated data. Content with more robust SLAs typically taxes IT resources and personnel the most. In order to address costs associated with supporting robust SLAs for secure, highly regulated, or business-critical content, IT departments can begin to align SLAs for backups and restores with additional site-based options and configurations (such as permission levels) and package these combinations as services (e.g. site provisioning) that can be delivered back to the business. These combinations of site configurations and data protection SLAs can be combined into service tiers such as Tier 1 for most critical content, followed by Tier 2 and Tier 3 for content that may require less stringent controls. These service tiers each align to their own set of SLAs, permissions, and site configurations. Organizations can start to associate costs with each tier, ultimately facilitating charge-back.

In addition to explicit RTO, RPO, and RLO measurements, a well-defined SLA for SharePoint should identify stakeholders, recovery scenarios, data integrity, and measurability.

The SLA is often the most difficult aspect of protection strategy development

For more information on how AvePoint helps automate the presentation of service requests to business users for common requests such as site provisioning, associating appropriate SLAs, permissions, and configurations upon approval, please see DocAve Governance Automation.

Approach 1: Not All Data is Created Equal

The most challenging business decisions made by a SharePoint administrator during the design of an appropriate backup architecture relate to the scheduling of backups for various data types. To achieve this, three criterias need to be evaluated:

- **Business criticality of the content**
- **Frequency of use**
- **Age of the content**

A useful means for visualizing these criteria, and how they relate to one another, is to map them out on a two-dimensional matrix. This matrix–known as the SharePoint Backup Planning Criticality Matrix™–assigns business criticality on the *x*-axis and usage activity on the y-axis. Figure 1 demonstrates how a sample organization might construct a Criticality Matrix. As displayed in this example, sales leads and customer records stored in SharePoint are extremely important to this organization. They are the lifeblood of the company, and any downtime or loss of such information would have drastic consequences. This type of content is also very frequently updated, perhaps several times each hour. A more aggressive backup plan will be required to limit any amount of content loss due to the frequency of content updates. In the matrix, this content will be associated with the upper-right cell.

Conversely, at the other end of the spectrum there resides content updated relatively infrequently–probably once every six months–that is of low relative importance to business processes. This content would reside in the bottom-left cell. In this example, employee guides and vacation policies are represented, as they are generally not critical to business and are modified less frequently than sales leads or company financial information.

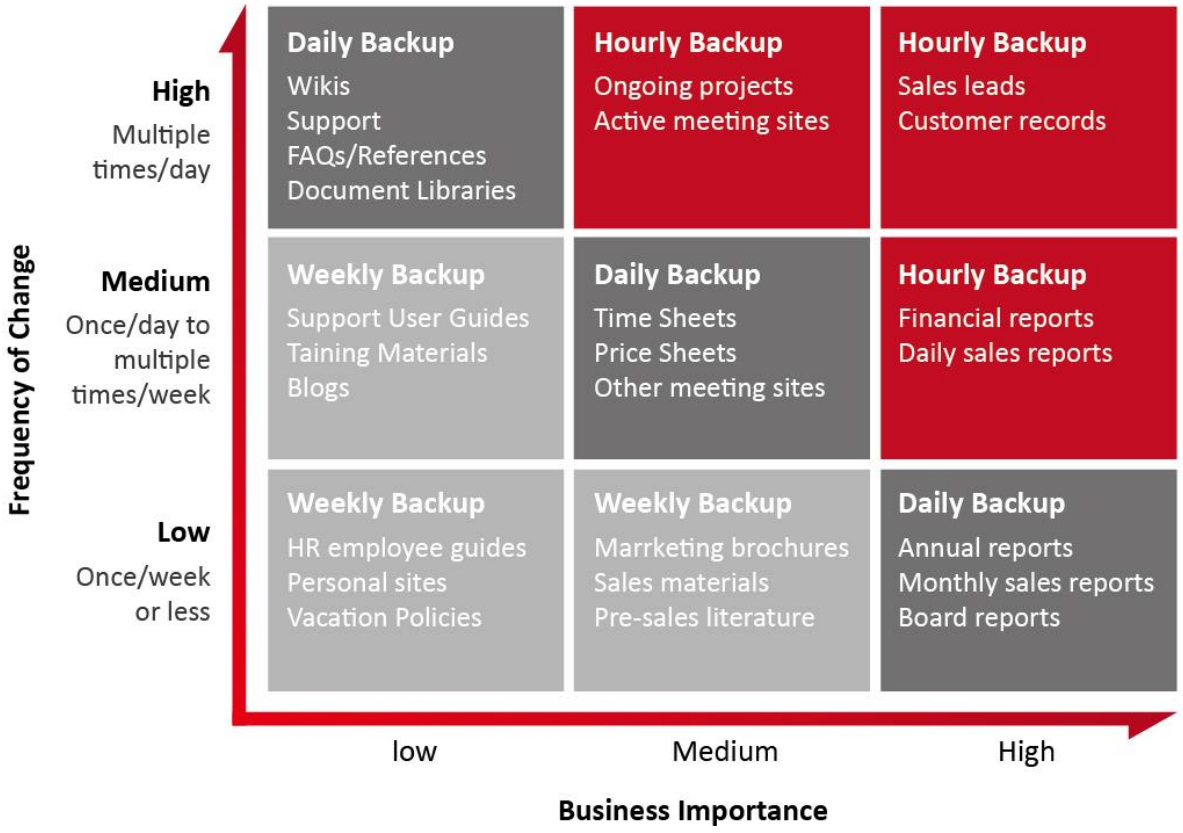

Figure 1: Sample SharePoint Backup Planning Criticality Matrix.

Once the criticality of all content has been identified, constructing a backup plan becomes relatively simple. DocAve allows administrators to build an unlimited number of backup plans. Backup jobs can be full, differential, or incremental (see Figure 3 and Figure 4 for details) and can be scheduled at specified hourly, daily, weekly, or monthly intervals. Furthermore, DocAve supports the protection of externalized content, providing administrators with the utmost flexibility in generating the perfect combination of backup schedules in order to truly cover the entire SharePoint environment.

By extending this approach to all of the other sites within the SharePoint environment, a modular and flexible backup system can be created that runs according to organizational timetables, mitigates duplication of backups, and ensures the security and integrity of all SharePoint core content.

Approach 2: Aggressive RPOs for all

As IT is often unable to fully understand the criticality of each site or piece of content without a well-organized corporate taxonomy, many may prefer to backup everything on the most aggressive frequency possible. To avoid missing RPOs for your most critical content, organizations then have to backup all data at the most aggressive desired RPO.

Historically, this may have resulted in an explosion of backup data, impractical for most organizations to attempt to maintain. For those who tried, this huge backup data volume would have to be periodically pruned or continuously moved to lower tiered or tape storage. However, with the advancement of backup technology and storage technology in Microsoft SharePoint and especially Microsoft Office 365, this options becomes the more pragmatic.

Backup technology: Using incremental backups, organizations can choose to capture only changes since the last backup. The ability to quickly snapshot a point in time, with the Azure AD Apps online or VSS on prem, is increasing the speed at which a full environment or farm can be captured.

Storage technology: Now, in the Microsoft Cloud, versions of the same file are stored as BLOBs and block changes. Prior to SharePoint 2016, for each additional file version, the total storage impact of that file would be multipled by the total version number. Now, only the file's changes are stored in subsequent versions, resulting in a drastically reduced storage footprint, and therefore backup storage footprint.

When combined: Incremental backup technology, paird with the advancements in file storage in Office 365, combine to ensure only the small, incrementally changed blocks, are included in incremental backups. Speed is once again reduced along with the backup storage footprint.

For SharePoint Online and on-prem, these combined technology enhancements enable a set-it-and-forget-it approach where your entire environment can meet your most aggressive RPOs.

Stakeholders

Stakeholders include not only management and end-users, but also the personnel that will play key roles in the event of a disaster. The roles of IT personnel should be clearly defined and appropriate communication channels should be established by and between all affected business groups (e.g. executive staff, end-users, and IT personnel).

Recovery Scenarios

Planning for recovery needs a comprehensive look at all scenarios possible, including:

- What will be done if the entire production environment is down?
- If only a single site is lost, how will a single site restore be performed?
- Can lost documents be recovered without performing full farm recoveries?
- If extended downtime is expected, can critical documents be restored to a file system for temporary access?

Data Integrity

End-users must know what to expect when content is recovered to the SharePoint environment. The SLA document must explicitly define what users should expect with regard to object

metadata preservation, access permissions, and versions. For business-productivity and compliance-related purposes, management should have an explicit understanding of how both security and content (including metadata) will be affected.

Measurability

All services described in an SLA document should be described in measurable terms. With regard to RPO and RTO measurements, service availability should be defined as a percentage of uptime, while recovery time should represent the average time end-users can expect content to be restored. A helpful measurement with regard to platform availability is presented as a function of Mean Time Between Failure (MTBF) and Mean Time to Recover (MTTR) using the following formula:

$$Availability = \frac{MTBF}{(MTBF + MTTR)}$$

Minimizing MTTR maximizes availability. This relationship helps illuminate the importance of strategies that include the ability for granular content recovery. In addition, how much content will end-users expect to lose during downtime? This amount of content can be determined by the backup frequencies covered by each plan. For example, if a document library is protected by an incremental backup plan that runs every hour, end-users can expect to lose a maximum of an hour's worth of data if a disaster occurs.

With these factors addressed, an SLA document will serve as a valuable blueprint for all affected parties in the case of data loss or platform failure. But an SLA is similar to any well-crafted governance or compliance plan in that it is a living document. To be effective, this SLA document must never remain static for too long. Appropriate and regular revision of an SLA–taking into account changing volumes of content and the business-criticality of each dataset–is vital for maintaining optimal data protection and disaster recovery strategies in an environment of constant deployment and organizational evolution.

It is also very important to conduct regular platform disaster recovery tests. Just as it is imperative to have a well-documented and precise disaster recovery plan, it is equally as critical (if not more so) to be able to execute it properly. Organizations that are successful in the event of a real disaster are those that have taken the time and attention to practice their disaster recovery plans in advance.

Performance Best Practices for SharePoint Backup and Recovery

In order to meet more aggressive SLAs for SharePoint content, meeting lower RTOs and improved RPOs, organizations have a variety of solutions and options for backup and restore methodology. Each available methodology, whether running out of the box backups or utilizing a third party, can be optimized to ensure best possible backup and restore performance for the

chosen method. Let's take a look at several best practices that can help optimize the speed of a variety of backup and restore solutions, allowing organizations to meet more aggressive SLAs.

For additional recommendations from Microsoft, please refer to TechNet for Backup and Restore Best Practices in SharePoint Server 2016.

Minimize latency between SQL Server and the backup location

In general, it is efficient to back up to a local disk on the database server instead of a network drive. You can then copy the data later to a shared folder on the network. Network drives with 1 millisecond or less latency between them and the database server perform well.

Avoid Processing Conflicts

Do not run backup jobs during times when users require access to the system. Instead, use maintenance periods or down-time in your environment by using reporting or analytics tools to identify periods of low SharePoint usage and then run your backups during those windows of time. A best practice is to always run incremental backups to safeguard against server failure. Also, consider staggering backups so that not all databases are backed up at the same time. For organizations that have global SharePoint environments that are active around the clock, look for backup tools that do not lock databases during backup operations. To best support this recommended practice, utilize backup tools with comprehensive scheduling capabilities.

Monitoring SharePoint Size for Faster Recovery Times

Keep databases small to speed up both backup and recovery times for full-farm or database-level operations. You can use multiple content databases for a Web application instead of one large content database and stick to the Microsoft recommended database size of 200GB. For Office 365 assets, there are also software boundaries and limits that you should know about. Organizations with larger amounts of critical data should meet suggested defined conditions. Refer to Microsoft's Software Boundaries and Limits for SharePoint 2016 and 2019 and SharePoint Online Software Boundaries and Limits for more recommendations.

Microsoft has revised the best practices boundaries to include the ability to extend databases beyond 200GB (in some cases up to 4TB) when RBS is used. While RBS does imply that databases using RBS will be smaller in size and therefore faster to backup, it is important to note that data protection will also need to include periodic backups of content created on the BLOB store to protect against corruption, keeping the backup data set relatively large.

Separating content into multiple, smaller databases allows you to perform parallel backups of content and greater flexibility in recovering databases.

You should also consider the site collection or OneDrive for Business size when determining the tools to use. If the business requires site collection backups in addition to farm-level or database-level backups, select the tools that you will use based on the site collection size.

- **15-100 GB** – Use the **Backup-SPSite,** a SharePoint Server 2016 tool, a SQL Server tool, or other database backup tool to protect the content database that contains the site collection.
- **Larger than 100 GB** – Use a differential backup solution, such as SQL Server 2014 or System Center 2012 - Data Protection Manager (DPM) R2, instead of the built-in backup and recovery tools. See Figure 4 for more on differential backups.

For more information, review the Microsoft TechNet article: Back Up Site Collections in SharePoint Server.

Avoid Getting Throttled or Blocked in SharePoint Online

SharePoint Online uses throttling to maintain optimal performance and reliability of the SharePoint Online service. Throttling limits the number of user actions or concurrent calls (by script or code) to prevent overuse of resources.

Here are three best practices to assist handling throttling:

- Create an Azure AD App Profile to perform backup operations
- Avoid using Service Accounts to perform backup operations
- Register your backup application or functions with an AppID and AppTitle to fast-track any issue resolution
- Reduce the number of operations per request
- Reduce the frequency of calls
- Use incremental backup to reduce the number and frequency of calls until no more throttling occurs

Incremental back off uses progressively longer waits between retries before trying again to run the code that was throttled.

Backing off is the fastest way to handle being throttled because SharePoint Online continues to log resource usage while a user is being throttled. In other words, aggressive retries work against you because even though the calls fail, they still accrue against your usage limits. The faster you back off, the faster you'll stop exceeding usage limits.

For more information, please refer to Microsoft's article: How to Avoid Getting Throttled or Blocked in SharePoint Online.

If your usage in Microsoft Office 365 extends beyond SharePoint Online, to Microsoft Teams and other applications, it's important to note that each service comes with it's own SLAs and recycle bin windows. You may have to plan for protecting each service. However, as SharePoint Online is the underlying content service behind OneDrive for Business, Microsoft Team Sites, and Group Team Sites, it's also critical that any backup approach avoid disrupting interdependencies among these services.

Use Incremental Backups for Large Databases

Use incremental backups for large databases. Incremental backups can be restored faster and more efficiently than full backups for larger databases. For more information on incremental backups, see Figure 3.

Use Compression During Backup

In some circumstances, you can use compression to decrease the size of backups and the time to complete each backup. Because a compressed backup is smaller than an uncompressed backup of the same data, compressing a backup typically requires less device I/O and therefore usually increases backup speed significantly.

It's to be noted that SharePoint Server 2016 supports SQL Server backup compression. SQL Server data compression is not supported for SharePoint Server 2016 databases. For more information about how backup compression affects performance in SQL Server, see Backup Compression (SQL Server).

Follow SQL Server Backup and Restore Optimization Recommendations

If you are using SQL Server backups, you can either choosing to use multiple backup devices or using a combination of full, differential, and transaction log backups (for the full or bulk logged recovery model) to minimize recovery time. Using multiple devices can increase throughput in proportion to the number of devices used. Similarly, the backup can be restored from multiple devices in parallel. Differential database backups are usually faster to create than full backups and reduce the amount of transaction logs required to recover the database.

For detailed recommendations about how to optimize SQL Server backup and restore performance, see Microsoft's article: Optimizing Backup and Restore Performance in SQL Server.

Now that we have reviewed the key considerations an organization must take into account during its SharePoint data protection planning, we can now examine how DocAve Backup and Restore can be utilized to support your data protection strategies.

Deploying DocAve Backup and Restore

DocAve Backup and Restore provides full-spectrum, item-through-platform level data protection for Microsoft SharePoint assets whether residing on premises or on a hybrid environment. DocAve protects entire SharePoint environments, from individual items and sites to Web applications, SQL Server databases, index servers, external BLOB stores, and IIS settings. Importantly, DocAve Backup and Restore delivers full-fidelity recovery of all metadata, security settings, version histories, and customized layouts. Uniquely, granular backup scheduling enables organizations to differentiate data and execute backup processes according to their specific, unique business needs. Additionally, DocAve Backup and Restore can be configured to automatically detect new SharePoint content and perform scheduled or on-demand full, incremental, and differential backups on this content. With these robust capabilities of DocAve Backup and Restore, organizations can be confident their SharePoint environment is optimally protected and their SLAs are met.

If you're currently living in a hybrid state, or have begun to transition to Microsoft Office 365, we recommend taking a Cloud First approach to protecting your cloud content. Fundamentally, the way we authenticate against and "talk to" cloud data, in order to protect it, can no longer be effectively supported from an on-premises architecture. Use AvePoint Cloud Backup with its simple set up and approach to meet aggressive RPOs side-by-side with DocAve for your on-premises SharePoint servers. Cloud Backup is 100% hosted on Azure as a SaaS solution. Cloud Backup protects SharePoint Online, along all your modern workloads. Your Teams, Groups, even Projects and Planner, are all protected with optimized cloud performance. Cloud Backup is built on AvePoint Online Services, our secure and scalable SaaS platform that is hosted across more than a dozen global Azure data centers. Simply sign in, and begin. Leave the capacity planning to us! To support your on-premises SharePoint investment, keep reading to learn how to optimize DocAve to support 1 or 100 SharePoint servers.

Overview of the DocAve System Architecture

DocAve Backup and Restore employs redundant and fully distributed components to ensure continuous uptime and failover for any DocAve processes being executed. While employing an agent-server deployment model that allows it to scale across multiple SharePoint instances and versions, each of the individual components can be broken down into services for further workload distribution, while enabling centralized management and control from a single browser-based interface accessible from anywhere within the organization's internal and external networks. This architecture is depicted in Figure 2.

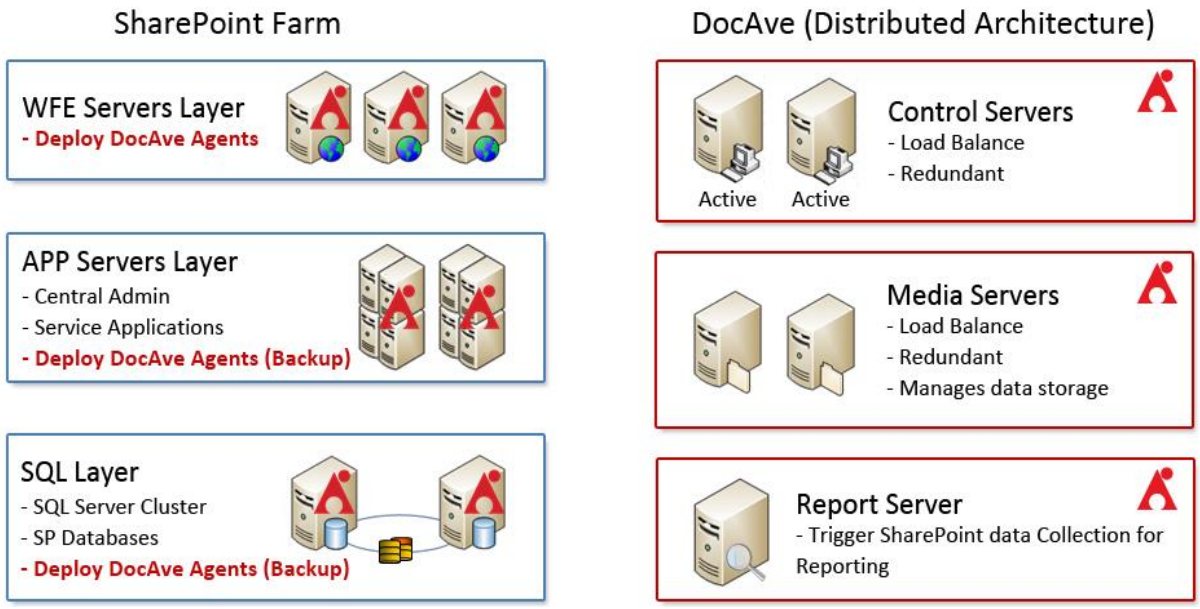

Figure 2: Fully Distributed DocAve System Architecture

Platform and Granular Backup and Restore

Operating together under the DocAve Backup and Restore umbrella are Platform Backup and Restore and Granular Backup and Restore. This two-part product features improved user experience, enhanced scalability, and even deeper integrated with leading storage hardware providers. Administrators can leverage wizard- or form-based backup and restore operations and granularly restore from hardware snapshots to simplify truly comprehensive SharePoint protection.

Platform Backup and Restore

Platform Backup and Restore protects your organization from disasters with a comprehensive toolset for backing up and restoring your entire SharePoint environment, including all content, customizations, solutions, and features, as well as back-end SQL databases, all configurations, index/job servers, front-end internet information services (IIS) settings, file system resources, BLOB stores that have been configured with Microsoft's Remote BLOB Store (RBS) API via the DocAve BLOB provider, custom application databases and the VMs hosting your SharePoint servers. With this product, you can restore an entire platform or individual SharePoint environment components.

Granular Backup and Restore

Granular Backup and Restore provides full fidelity backup and recovery, from an individual content item to an entire SharePoint site collection, maintaining all metadata, security settings,

and version histories. It offers full, incremental, and differential backup capabilities for select SharePoint content to build backup plans and schedules that focus on frequent backup of high priority data–improving backup operations and storage efficiency. Content that has been externalized by Microsoft's RBS API–via FILESTREAM, DocAve, or other third party BLOB providers–will be included by default in DocAve's granular backup and restore processes.

The solution is composed of three main components (DocAve Manager, DocAve Media Service, and DocAve Agent) designed to maximize workload distribution while providing flexible, intelligent, and granular backup and restore functionality. As depicted in Figure 2, the DocAve Manager Service consists of load-balanced Web services that manage requests/responses and an active-passive pair of Control Services. This highly available architecture for the Control Service allows DocAve to failover to a standby job service to continually process any jobs in the queue.

DocAve Manager

DocAve Manager centrally administers all activities performed by the DocAve Software Platform, enabling administrators to control routines, schedule jobs, and view or act on any error conditions raised by the agent processes. DocAve Manager consists of three services: Control Service, Media Service and Report Service. They can either be run on the same server as your DocAve Agent or split across several servers.

The Manager can be installed on any Microsoft Windows-based hardware with appropriate processing power and access to its member agents throughout the network. It is recommended to install the DocAve Manager on a machine with high availability, although it does provide an active-passive control services cluster to ensure continuous availability. This highly available architecture for the Control Service allows DocAve to failover to a standby Control Service to continually process any queued DocAve jobs.

DocAve Media Service

DocAve Media Service provides dedicated computing resources to read and write to storage media. For largely distributed deployments, it is recommended to have the media services deployed within close proximity of the back-end Web servers and the physical storage, but not on the same hardware. It is also recommended to host the media services on hardware with high reliability in addition to high availability to prevent backups from being interrupted due to hardware failure and needing to re-run.

Multiple dedicated media services can be distributed across multiple SharePoint environments to handle additional process workload, thus improving the overall performance of backup and restore jobs. Additionally, media services also support the ability to persist data to multiple logical storage devices. Media services can be associated to a limited subset of virtual storage devices, so DocAve can automatically attribute the backup processes through the designated

media services. By doing so, DocAve can further optimize the storage pools according to both business activity and established SLAs.

DocAve Agent

SharePoint agents are responsible for running DocAve jobs and interacting with the SharePoint object model. At a minimum, DocAve agents enable DocAve Manager to communicate with the respective servers, allowing for Granular Backup and Restore commands to function properly. Each DocAve Client Agent is a lightweight, application-specific client component. In the case of DocAve Backup and Restore, the Client Agent resides on a SharePoint server, including application servers such as the FAST Search server.

The operations of the agent are controlled by the Control Service on the DocAve Manager via XML messages over a network. The client agent reads and streams selected data via the SharePoint server. Multiple DocAve client agents can be directed to communicate with one another, allowing the transfer of data between SharePoint environments. This communication is how Backup and Restore covers elements from one SharePoint farm to a different farm for restoration purposes.

Unique capabilities within the DocAve architecture ensure that communication between agents can endure very noisy or even intermittent data channels. Successful results from strenuous data transfer performance and stress tests have showcased the strength of DocAve's data packet level fault tolerance features. While it is possible to have the DocAve Manager and DocAve Agent on a single server, it is not recommended. For the best performance, install all the manager's services across multiple servers and activate only the necessary agents on the agent servers.

Considerations for Large-Scale Deployments

For larger SharePoint topology, especially the hybrid deployment that encompasses both on-premises SharePoint Server and SharePoint Online instances, the distributed architecture of DocAve Backup and Restore allows the solution to scale with the growing needs of the organization. Performance and management of large deployments can be maximized by addressing the considerations described in the following sections.

Plan for SharePoint Boundaries and Performance Guidelines

The first step to ensuring that data protection and disaster recovery plans can scale effectively is to deploy the SharePoint environment according to Microsoft's Software Boundaries and Limits for SharePoint 2016 and SharePoint Online Software Boundaries and Limits article on TechNet. As we mentioned in the previous section on Monitoring SharePoint Size for Faster Recovery Times, following Microsoft's recommendations will help administrators provide management and end users with appropriate solid recovery estimates.

Limiting Backup Sizes per Backup Plan

As we discussed earlier, the amount of data backed up inversely relates to the duration of the subsequent backup window. Longer backup windows can lead to performance loss, inefficient use of system resources, and increased data loss/inconsistency. To prevent such issues, a generally accepted rule of thumb is to make efficient plan groups and limit each backup plan according to the backup methodology utilized. For instance, limiting the scope of granular backup plans to less than 100 GB of data would be most efficient. But, in platform level backups where advanced VSS technology can be utilized, the scopes can be as large as your organization requires. The out-of-the-box granular capabilities and unlimited number of backup plans that can be created within DocAve Backup and Restore provide ample flexibility to support this deployment scenario. For more information on limiting backup sizes per plan, refer to the Microsoft TechNet article Backup and Restore Best Practices in SharePoint Server 2016.

Utilize Hardware Snapshot Technology

When discussing the importance of rapid backup and restore options for all of your content, it is crucial to address hardware snapshot technology. AvePoint partners with hardware vendors that provide and generic and custom hardware options to support VSS hardware snapshot technology to satisfy the backup requirements of growing SharePoint farms.

One key example that showcases the usefulness of hardware snapshots is the backup of production data during business hours. This scenario could potentially be disruptive to the network (in copying data to storage devices), to the SharePoint server load (in indexing of the backup data), or in locking SharePoint content during the backup of a given site. Being able to take a snapshot of production data limits the interruption to the services. DocAve is able to schedule the activities of indexing content and moving to storage devices at a later time so hourly backups can run without impact to the business. In addition, the backup window itself is much shorter due to the hardware snapshot process.

Utilize Multiple Media Services

DocAve Backup and Restore supports the deployment of multiple DocAve Media Services on separate physical hardware. As the amount of SharePoint content requiring backup grows, there is increased likelihood that multiple backup processes will be scheduled for execution at the same time. In order to optimize the media services' writing of backup data into storage, it is recommended to deploy additional media services to limit each media server's load processing to five concurrent jobs. Additionally, because media services should be deployed within close proximity to SharePoint infrastructure, deploying a minimum of one media service per geographic region is also strongly suggested.

Use Load-Balanced Agent Groups in Granular and Platform Backups

DocAve Backup and Restore provides built-in capabilities to handle agents within agent groups. Within heavily distributed SharePoint farms, one agent should be deployed and enabled within each load-balanced front-end Web server, so DocAve can automatically identify and assign a backup process to the least loaded server available, in order to accelerate any job already in process.

Optimizing Your Backup Plans

Once DocAve Backup and Restore has been deployed, there are several recommended approaches for implementing backup plans. By fully utilizing DocAve's built-in capabilities–such as granularity and storage pooling–administrators will be able to deliver more robust protection coverage and satisfy even the most stringent SLAs.

Evaluating Your Storage Requirements

After having determined all assets within the SharePoint environment needing protection, and calculating both the amount and the criticality of content, the next step is to establish specific backup plans in DocAve 6 Backup and Restore. The first step in this process is to ensure that the appropriate storage resources are in place to support anticipated data protection routines. To determine storage requirements, the following questions must be answered:

- **How much hardware will be needed to host the backup software, and how much disk space will full backups consume on the storage media?**
- **Will cloud storage devices be utilized for storing backups?**
- **Are there storage devices with externalized content that must also be incorporated into a backup and restore plan?**
- **What additional type of security options will be placed on the storage media?**
- **How should pruning rules be defined for the best possible use of valuable storage space when backups are no longer required?**

Answering these questions with accurate estimates is critical to optimizing backup plans. Your answers will influence how frequently full backup jobs can be executed, and how often incremental backup jobs can run to fill in the gaps. Similarly, by pooling available storage resources, multiple volumes can be managed as a single target backup destination to more effectively handle large data backups.

DocAve Platform Backup and Restore uses Microsoft Volume Shadow Copy Services (VSS) snapshot technology to generate snapshots (VSS shadow copies) and then protects your SharePoint farm. When configuring a backup schedule or running a backup plan, you can specify the type of backup to perform:

- **Full** – Backs up all of the selected content, including transaction log files, each time a backup is performed.
- **Incremental** – Backs up only the transaction log, drastically reducing the size of the backup file created (see Figure 3).
- **Differential** – Backs up all content that has been updated since the last Full backup (see Figure 4).

Figure 3: One Full backup followed by three Incremental backups.

In Figure 3, the first incremental backup backs up the newly-added data in the blue period. The second incremental backup backs up the newly-added data in the green period. The third Incremental Backup backs up the newly-added data in the orange period.

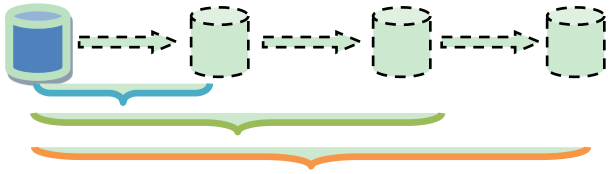

Figure 4: One Full backup followed by three Differential backups.

In Figure 4, the first differential Backup backs up the newly-added data in the blue period. The second Differential Backup backs up the newly-added data in the green period. The third differential backup backs up the newly-added data in the orange period.

Estimating Storage Space for DocAve Backups

The actual amount of physical storage space required for SharePoint backups is dependent on the v**olume of data (core content plus platform components) and the pruning rules in DocAve that define the number of backup cycles to be retained.** SLAs should specify the length of time and number of backups that should be maintained to adequately support recovering the variety of content being protected. These time periods will directly affect–and be influenced by–the storage resource requirements described in the aforementioned section. A good starting point for estimating the storage space requirements is the following equation (applies to both platform- and granular-level backups):

Required Storage Space =
(1.5 × Volume of Content) × Number of Backup Cycles + Volume of Content

For example, if you have 500GB of content and 2 backup cycles, your storage space requirement is 2000GB:

Required Storage Space = (1.5 × 500GB) × 2 + 500GB
Required Storage Space = 2000GB

For each backup cycle retained, ensure storage space is allocated for one full backup cycle of all data, plus an additional half-volume to compensate for any data updated during the cycle. Using DocAve Backup and Restore, pruning jobs can be executed after a backup job. Therefore, additional storage should be resourced for consideration of additional full backup cycles that complete before a pruning is executed.

Another factor affecting storage allocation is compression. When used efficiently, compressing backup files can effectively reduce the storage requirements by up to 75 percent. Of course, how much data can be compressed depends on the type of content. Compression also reduces disk space but increases the backup window and affects the impact of performing backups on the production infrastructure.

A final factor in considering storage allocation requirements is the level of optimization among the various backup plans. Often, content selected for backup can also be present in other backup plans depending on the type of backup plan you have. Therefore, the content will be duplicated in storage. Eliminating unnecessary backup jobs, such as optimizing backup plans to reduce redundancy, can reclaim storage space and reduce further outlays on wasted storage and reduce backup windows and the impact of taking backups on production operations.

Setting up Business-Aware Backup Plans

Because SharePoint is a business-oriented technology platform, data protection and disaster recovery solutions should always take into consideration the needs of the business. Many organizations may think they have a comprehensive disaster recovery plan in place, but simply having a plan is not enough. These plans need to be applicable to the organization's real-world situation and needs, and appropriately validated. Without such validation, satisfying applicable SLAs and maintaining critical business-processes becomes difficult.

Granularity of Backups

Backup plans can be created at the site collection, site, or item levels. Selecting the appropriate level depends on a number of different factors (see the previous section Not All Data is Created Equal).

The first factor concerns the granularity to differentiate discrete data for backups. There are occasions when sites are already logically segregated, and associating a backup plan for each individual site collection or site is simple and straightforward. This would reduce the number of backup plans that need to be managed, and will also more effectively limit the possibility of duplication.

The second factor concerns the restore options appropriate for the content. DocAve Backup and Restore indexes backup files for restoration according to the level of backup chosen. For example, a site-level backup will only allow entire sites to be restored, while individual content items as well as item versions can be restored from item-level backups. Bear in mind that item-

level backup plans also cover entire sites or site collections, as well as the nested sites that get created. So selecting an appropriate level of backup is really dependent on how much granularity is required for restoration and the time allotted (as defined in the SLA) to execute the backup job.[1] We recommend always beginning with item-level backups, as this provides the most flexibility in not just data differentiation, but also by providing item-level and item version restore capability.

One final consideration in determining backup levels is exactly what data is being restored at each level. Backup files generated from site-level backups are typically larger than similarly targeted item-level backups, as these backups encompass *all* of the assets within that site. This can affect storage resource allocations, as the execution of incremental backup jobs of entire sites will be backed up again, even if only individual content items were updated since the last full backup was executed.

Security Trimmed Access to Backup Data

Security trimmed access to the restoration of SharePoint backup data helps organizations distribute the administrative workload previously held by farm-level administrators and helps to further optimize IT resources and reduce recovery time and recovery point objectives. DocAve provides security-trimmed access for site collection owners and a security trimmed restore Web part.

Plan for Probable Recovery Scenarios

With comprehensive backup plans in place to provide adequate protection of the SharePoint environment, administrators should also plan for possible recovery scenarios. Again, granularity is crucial, as any type of recovery activity should not only satisfy SLAs, but also minimize any negative impact to environment activity.

Keep in mind that recovery procedures can also be delegated to other users with appropriate access permissions to DocAve. For example, IT helpdesk staff–or even the site/content owners themselves–can optionally perform simple recovery of lost or corrupted items, while full platform recovery might remain the sole responsibility of the SharePoint administrators.

Some common events that trigger recovery and require planning:

- **Hardware failure** – Hardware failure typically means environment downtime. In this case, it is important to be prepared to bring a standby SharePoint farm online as quickly as possible, in accordance with SLAs. If the production environment has been configured for high availability, with technologies to update content on the standby environment,

[1]Due to the indexing required during item-level backups, time necessary for such execution is longer than a site-level backup of the same site. However, in the case of incremental backups, site-level backups might possibly require longer time periods, as the job will attempt to backup the entire site even if deltas are present in only a single document within that site.

content loss can be minimized (DocAve High Availability provides this functionality). Otherwise, performing an out-of-place content restore to a separate SharePoint environment can serve as a temporary solution.

- **User error** – Users, developers, and administrators may accidentally delete content, remove entire site collections, or make incorrect configuration changes. User errors are the primary reason for site or content recovery.
- **Corrupted content (e.g. viruses, data corruption)** – SharePoint's native Recycle Bin feature will not be able to safeguard end-users from corrupted content. Corrupted content must be addressed by overwriting the current file with uncorrupt backups.
- **Data center destruction** – Destruction of a data center requires a full platform recovery, likely onto a separate SharePoint farm that can be brought online in a separate location. Alternatively, in order to better prepare for such a disaster, a high availability strategy using technologies such as SQL mirroring, log shipping or AlwaysOn Avaliability, can be established to create a warm standby SharePoint environment. In this circumstance, content loss would be determined by the frequency of content replication.
- **Unsuccessful implementation of changes** – A common cause of failure is planned changes or releases which don't complete successfully and require rollback. For more information on deploying changes across SharePoint environments, refer to DocAve's From Development to Production: Streamlining SharePoint Deployment with DocAve Deployment Manager white paper.

With enhancements such as automated, business-centric SharePoint backup and recovery as well as a point-in-time restore controller interface, DocAve is poised to again revolutionize the way administrators manage and protect SharePoint and Office 365.

Protecting Farm-Level Components

As we discussed previously, core content is the most significant asset within a SharePoint deployment, as it represents intellectual property that is critical to operations. But the underlying infrastructure of the SharePoint deployment is significantly important because it consists of many platform-level components coupled tightly together to create a working environment. Since there are so many moving parts in a SharePoint platform, a failure in any one component could adversely affect the entire environment.

DocAve Backup and Restore provides comprehensive protection for not only the core content residing in SharePoint's content databases, but also for the platform's full spectrum of farm-level components, their configurations, as well as content that has been externalized using Remote BLOB Store (RBS). All farm elements –including the configuration database, Web applications, search and index servers, IIS settings, and other file system artifacts on the front-end Web server–are protected.

Performing a Platform Backup in DocAve

The user experience during platform-level backup creation in DocAve Backup and Restore is similar to that of site collection, site, and item-level backups. Administrators can schedule backup plans to execute full, incremental, or differential backups on the servers where DocAve agents are installed. Since this is a platform-level backup intended to safeguard the environment against catastrophic events, the resulting backup files should be stored in a location different from the primary SharePoint farm to ensure that even if the full farm is not functional, access to the SharePoint backups is maintained. Options such as data retention (pruning), encryption, and compression are the same as those described in our discussion of core content backup options. There are, however, a few options specific to platform-level backups including: selecting DocAve InstaMount, coexisting with existing SQL Server backups, identifying and selecting components to protect, and determining restore granularity level.

Selecting DocAve InstaMount

In typical farm-component backup solutions, performing an item-level recovery from a database-level backup requires organizations to create a staging farm or mount a database back to a SQL server, a typically time-consuming process that is directly related to the size of the database.

Selecting DocAve InstaMount allows individual items to be recovered quickly from a full backup to the production environment without needing to first stage the database. InstaMount generates a virtual database while restoring data and uses a mapping file to record the relationship between the virtual database and the backup data. The data included in the virtual database depends on the node you selected from the farm tree. For example, if you select 10 documents from the farm tree to restore, the virtual database only contains the data of these 10

documents, not the whole database. InstaMount can speed up the recovery dramatically. Having this capability is critical in order to meet stringent SLAs and growing database sizes.

Coexisting with SQL Server Backups

Database administrators typically have their own SQL Server backup plans in place to protect the databases in the case of any unplanned disaster. These database servers include those used for custom applications or other operational data repositories, as well as all of the SharePoint content and configuration databases. DocAve Backup and Restore platform-level backups can coexist with institutional database backup utilities and automatically prevent any conflicts from taking place. Therefore, a recommended approach is to keep SQL Server backups as they are currently staged, and protect only the SharePoint content databases via DocAve.

Full SharePoint platform-level backups use DocAve Platform Recovery in addition to other existing SQL backup methods for SharePoint databases by taking a copy of the database and associated transaction log rather than committing and truncating the transaction log in the existing database. By doing so, SharePoint administrators have the flexibility to holistically capture a full SharePoint environment via a DocAve backup, as opposed to having to manage SQL Server backups/recovery with a separate utility.

Identifying and Selecting Components to Protect

Similar to the site collection, site, and item-level backup functionality in DocAve, the platform-level backup builder provides a data tree populated by the various components available for granular selection for each backup plan. These platform level components include:

- SharePoint farm configuration database
- Web applications/content databases
- Central administration content database
- Global search setting database
- State Service database
- Session State Service database
- Usage and Health Service database
- Shared Services Applications for SharePoint related components, including Managed Metadata Services, the associated databases and indices
- SharePoint solutions installation files
- Nintex databases, including Nintex workflow configuration database, content database, and solutions
- InfoPath Form Services and all the form templates installed on the front-end web server and Form Services configuration

- Single Sign-On configuration and database (Replaced by Secure Store Service in SharePoint 2013)
- Front-end Web server components, including IIS settings, SharePoint templates under 16 Hive [2]for SharePoint 2016, custom features, custom site definitions, and other file system folders
- SharePoint Learning Kit

With respect to the front-end Web server components, there are a number of SharePoint configurations that are not stored within the configuration database but reside on the front-end Web server file system itself. Configurations such as SSL configuration, forms-based authentication, and Web Part configuration are located under the IIS settings web.config file. Additionally, custom templates, site definitions, and features are deployed in the "SharePoint Hive". These customizations should be included in backups; administrators should not need to perform the customizations again once the environment is restored. As an added capability, the data tree provided in DocAve expands further into the front-end Web server file system, so any additional dependencies external to the SharePoint configuration can likewise be protected within the same backup plan.

Determining Restore Granularity Level

The platform-level data tree in DocAve lets administrators drill down and granularly select the farm-level components to be protected per backup plan. Administrators have the flexibility to group certain Web applications, content databases, and other individual components into their own backup plans. Furthermore, DocAve lets administrators simply select an entire SharePoint farm, and any child components it includes will be automatically protected within a single backup plan. Even new components that are created will be automatically protected.

Regardless of the granularity of backup plans, DocAve Backup and Restore can index the platform-level backups for item-level or even item version-level recovery, by allowing for selection of the appropriate restore granularity level when the plan is defined. If the restore granularity level is set to site level, the backup process will index the backup file to allow restoration at the site level. If the restore granularity level is set to the item level, then even individual items can be restored from a full farm backup.

Platform Restore Processes

In the event of a catastrophic SharePoint farm failure, it may be necessary to restore the entire farm using a platform-level restore. There are two platform-level restore processes to be considered–an in-place restore (put the farm back to the way it was) and an out-of-place restore

[2] The "16 Hive" for SharePoint 2016 is located on the file system at C:\Program Files\Common Files\Microsoft Shared\web server extensions\16\TEMPLATE.

(recover and restore the farm somewhere else). Depending on the type of hardware being used, the type of disaster recovery plan may call for reduced SharePoint topology or identical hardware to perform the plan that may not be available. The following prerequisites must be met in order to proceed:

- Windows Server 2008, Windows Server 2008 with Service Pack 2, Windows Server 2012, or Windows Server 2012 with Service Pack 2
- IIS with ASP.NET enabled
- .NET Framework
- Microsoft SharePoint Server installed and configured (if an existing farm already has front-end Web servers deployed, they should all be disconnected through the SharePoint Products and Technologies Configuration Wizard)
- SharePoint patch level should remain the same
- Server names and topology should remain the same
- SQL Server disk layout should remain the same
- Same domain account for SharePoint administration should be used

Once these prerequisites have been satisfied, the farm can be recovered by simply loading the appropriate platform backup through the DocAve Restore Controller. Due to dependencies between various elements of the SharePoint farm, this restoration process is usually a multi-step procedure. Before any front-end Web servers can be attached, the configuration database and the Central Administration databases must first be restored. Once restoration is complete, the front-end Web servers can be brought online and connected to the restored configuration database using the SharePoint configuration wizard. It is important to keep in mind that one of these front-end Web servers should host the central administration Web application.

After the Web front-end servers are online, additional farm-level components can be restored. These components include any customized IIS customizations (e.g. SSL or forms-based authentication, web.config), custom solutions, and Shared Services Applications.

Any restored front-end related components (e.g. IIS settings, custom site definitions, and features) should be restored on all of the Web front-end servers of the SharePoint farm. This process is drastically simplified and is presented to administrators in a step-by-step fashion with DocAve's Farm Rebuild Wizard.

About AvePoint

AvePoint is the Microsoft Cloud expert. Over 15,000 companies and 3 million cloud users worldwide trust AvePoint to migrate, manage, and protect their Office 365 and SharePoint data. AvePoint's integrated cloud, hybrid, and on-premises software solutions are enhanced by 24/7 support and award-winning services. Organizations across six continents and all industries rely on AvePoint to ease transition to the Microsoft Cloud, increase IT administrator productivity, and satisfy governance and compliance objectives.

A three-time Microsoft Partner of the Year Award winner, AvePoint has been named to the Inc. 500|5000 six times and the Deloitte Technology Fast 500™ five times. AvePoint is a Microsoft Global ISV Partner, Gold Application Development Partner, Gold Cloud Platform Partner, Gold Cloud Productivity Partner, Gold Collaboration and Content Partner, and US Government GSA provider via strategic partnerships. Founded in 2001 and headquartered in Jersey City, NJ, AvePoint is privately held and backed by Goldman Sachs.

Notices and Copyright Information

AvePoint, Inc.
525 Washington Blvd
Suite 1400
Jersey City, NJ 07310
USA