

# Protect Your SharePoint Data

February 2014

## CONTENTS

Introduction.....	3
SharePoint Disaster Recovery (DR) Planning Concepts.....	3
DR Plan Benchmark Targets.....	4
Prerequisites and Assumptions.....	6
A Few Words of Guidance.....	9
Conclusion .....	10
About Metalogix .....	11

## INTRODUCTION

SharePoint is one of the most popular server platforms Microsoft has ever released. This success, coupled with SharePoint's natural strengths in document storage and information management, means that organizations are utilizing SharePoint on a daily basis for the storage of many forms of business critical data.

The centralized storage of any data is an inherently risky business. Disasters take many shapes and forms, and they can range from something as small as the deletion of a single document all the way up to the destruction of a key facility by a tornado. Because disasters come in so many shapes and sizes, and because they are inherently unexpected, it is an exercise in futility to attempt to prevent them. The most prudent course of action is to simply accept that disasters occur and to plan for a time when your environment is impacted by one.

This whitepaper introduces the broad topic of disaster recovery (DR) and how its principles can be specifically applied to the Microsoft SharePoint 2010 and 2013 platforms. Though the technical mechanics of how to backup, restore, and generally preserve SharePoint content are not discussed, backup and restoration are investigated with the intent of explaining how these approaches fit into a more far-reaching DR strategy.

## SHAREPOINT DISASTER RECOVERY (DR) PLANNING CONCEPTS

It is all too common that the phrases "SharePoint disaster recovery" and "SharePoint backup/recovery" are used interchangeably. While the latter can be (and oftentimes is) an aspect of the former, it cannot be said that the reverse is true. An effective SharePoint DR strategy needs to be more than just backups and restores, extra hardware, and an off-site data center. Focusing solely on these technological aspects without concern for the underlying business drivers and targets results in an untargeted (and perhaps inappropriate) DR plan more often than not.

One of the most common misconceptions surrounding disaster recovery is that DR planning is a technological function that should be driven by IT personnel. While information technology (IT) personnel typically care about the viability and stability of the systems they manage, they may have little stake in the content and business data contained within those systems. While DR planning should involve IT personnel, it is critical that business stakeholders drive the process so that the true needs and goals of the people using SharePoint are considered and met by an effective DR approach.

Properly executed, the process of formulating a SharePoint DR strategy commonly begins with discussions and investigations that seek to build an understanding of an organization's business processes, data, and how the two combine within the SharePoint environment. Through risk assessments and associated business impact analysis (BIA) exercises, an organization gains the insight needed to prioritize and rank the relative importance of its SharePoint activities and information—oftentimes with specific dollar amounts. These analyses and rankings then serve as the basis for the creation of one or more business continuity plans (BCPs) that seek to minimize risk exposure, prevent loss (monetary, material, etc.), and sustain an organization as effectively as possible in the event of a disaster.

## DR PLAN BENCHMARK TARGETS

Though IT professionals and those tasked with responsibilities for a SharePoint environment may have some visibility and input into the process just described, quite often they do not. Involvement in the process is not a prerequisite when creating an effective SharePoint DR plan, but comprehension of the process is important from the perspective of knowing what a BCP is and understanding that it defines many targets, processes, procedures, and system-level recovery targets—including those that dictate what constitutes an acceptable SharePoint recovery plan. In particular, BCPs that cover a SharePoint environment commonly dictate two specific types of target benchmarks that any DR strategy needs to acknowledge and accommodate in order to be considered successful: recovery point objectives (RPOs) and recovery time objectives (RTOs). Without an understanding of RPO and RTO targets for a SharePoint environment, attempts to formulate a DR plan are nothing more than uninformed “best guesses.”

### RECOVERY POINT OBJECTIVES

Put bluntly, an RPO defines the maximum amount of data loss that is deemed acceptable in the event of a disaster. An RPO is commonly measured in hours or minutes, but intervals can be substantially longer or shorter depending on the function of the SharePoint farm (or specific aspect of the SharePoint farm) being described. Systems or data that are of minimal relative importance to an organization may have RPOs of a week or more, while highly critical systems could have an RPO of zero (that is, no acceptable data loss) at the other extreme.

### RECOVERY TIME OBJECTIVES

An RTO target describes the maximum amount of time that recovery personnel have to return a SharePoint environment, or a specific piece of an environment, to a pre-determined state of operational readiness following a disaster. RTO targets, much like RPO targets, are typically measured in hours or minutes. As with RPOs, RTO intervals can vary significantly. Less critical SharePoint environments (basic FAQ or non-revenue-generating marketing sites, for example) may have RTO targets measured in weeks, while business-critical SharePoint environments could have RTO targets that are measured in seconds or even zero (that is, zero loss of functionality in the event of a disaster) for real-time failover.

### PERFORMANCE OFTEN COMES AT A COST

If money were not a factor, most SharePoint personnel would opt for a DR solution or strategy that negates data loss and downtime completely. Unfortunately, highly performing SharePoint DR strategies and tools typically come with a greater price tag (often significantly so) than those that perform less quickly or effectively. As a rule of thumb, you can expect the cost of a SharePoint DR solution to grow as your RPO and RTO benchmark targets decrease. As these targets approach zero (that is, zero data loss and zero recovery time), the cost of an associated DR solution can rise quite dramatically.

Generally speaking, the total cost of any DR strategy is going to include more than just the cost of a tool

or license, as well. While low-end DR solutions may involve little automation and require constant human intervention, the overwhelming majority of SharePoint DR strategies include multiple pieces of software. A well-formed strategy may include not only tools and software, but additional hardware, one or more remote data centers, additional personnel, testing and maintenance cycles, and more. Once again, as RPO and RTO targets approach zero, the number and prevalence of these additional components increases. These increases, in turn, result in increased complexity and cost.

## SHAREPOINT DR STRATEGIES AND THE RPO/RTO CONTINUUM

All SharePoint DR strategies implement one or more approaches that fall somewhere on the RPO/RTO continuum. For purposes of further discussion, SharePoint DR solutions are broken into three segments or categories: “Adequate,” “Better,” and “Best.”

“Adequate” solutions are those that meet a disaster recovery requirement to preserve the entire SharePoint farm but do not stress fast backup speeds and rapid recovery times. These solutions are typically free and come using what is provided by SharePoint and/or SQL, and their performance (or lack thereof) reflects that fact. “Adequate” DR strategies are oftentimes used only in specific scenarios where RPO and RTO target intervals are quite long (weeks or more, for example) and the underlying SharePoint environment configuration is not considered business-critical in any way.

SharePoint DR solutions that fall into the “Better” category typically incorporate the use of tools to regularly capture changes in content and business data within a SharePoint farm, as well as dependent system and environment configuration. This information is then made available (commonly at a remote location) for restore operations in the event of a disaster. Fast and simple backups and restores separate these solutions from what is available out of the box. Of the three solution categories, approaches that fall within the “Better” category typically offer the best performance per dollar.

Finally, DR solutions that are considered “Best” provide superior performance and oftentimes come with price tags that are dramatically larger than those found in the previous category. Solutions in the “Best” category often leverage high availability (HA) hardware and software to address RPO and RTO targets that may be measured in minutes or seconds. DR solutions in this category are typically only employed in scenarios where downtime and data loss must be avoided at all costs.

## PROTECTING YOUR SYSTEMS

When considering how to protect your SharePoint environment with a DR solution, it is important to realize that no one tool or process is going to address all requirements and recovery targets. You should be prepared to implement a tool or strategy to back up your critical SharePoint content, but the approach selected needs to fit into an overall DR plan, not take the place of it. Though the SharePoint platform comes with a set of backup and restore tools, these tools address only a subset of the full range of DR concerns. These tools also come with their own unique set of idiosyncrasies, limitations, and problems that can directly impact when and how they are used in the event of a disaster. Remember, it is just as important to know what your tool or strategy cannot do as what it can do to when creating a successful DR strategy in your environment.

## PREREQUISITES AND ASSUMPTIONS

It is important to identify a few prerequisites and assumptions that should be taken into account when designing a DR strategy.

First, your plans should include and target an offsite facility or datacenter that can be brought online in the event of a disaster. By including such a location in your strategy, you reduce your exposure to risks that could arise from either the destruction of your primary facility or an inability to access it. Think of it like the investment strategy of diversification; resources should be spread beyond your primary facilities so that business operations can continue in the event of a cataclysmic event at your primary location. Without this type of redundancy, you are highly susceptible to geographically natured risks, power grid failures, and other large-scale disasters.

High availability, or HA, is the concept of introducing redundancy for key components, such as hard drives, SharePoint web front-end servers, and SQL Server hosts. The use of VM snapshots in SharePoint DR is not discussed, because Microsoft does not support the use of virtual server snapshots as a DR solution because of performance degradation. While HA may be leveraged in many environments for reasons of performance and responsiveness, it is discussed here only for the purposes of continuity and redundancy that are important in SharePoint disaster recovery scenarios.

## THE OPTIONS: IS GOOD, BETTER, OR BEST POSSIBLE?

Like most IT solutions, designing and implementing a SharePoint DR strategy is often a series of trade-offs. It can come down to weighing the benefits and disadvantages of several different approaches to find a solution that meets the functional needs of the business while best leveraging the available resources of your IT department. In a perfect world, each SharePoint DR solution could be described as “Good,” “Better,” and “Best.”

The reality, though, is that available options are better described as falling into the categories of “Adequate,” “Better,” and “Best.”

## THE “ADEQUATE” SHAREPOINT DR SOLUTION

At a minimum, every SharePoint DR solution needs to focus on the content and business data that users supply and trust SharePoint to store. An “Adequate” SharePoint DR solution is one that focuses on preserving this data in a repeatable, consistent, and reliable fashion to safeguard it against damage, corruption, or even destruction.

In the absence of a third-party product, three tools exist within the SharePoint environment to address this requirement: Windows PowerShell, Central Administration, or SQL Server tools. Thus, these tools are not mutually exclusive, and they can be employed in tandem to create complete farm backup or redundant backups for greater peace of mind at no additional investment in software. Note that SQL Server’s tools do not backup the farm’s configuration.

These tools are considered “Adequate” because they give you the ability to protect the most critical pieces of your SharePoint environment, but do not provide full coverage for your entire farm and some of its dependencies. Farm backups will not include Internet Information Services (IIS) configuration, remote BLOBs that are not using the FILESTREAM provider, trust certificates, registry keys, and some key files on front-end web servers like manual changes made to a web.config.

Central Administration and PowerShell also provide granular site collection backups. These can be more resource intensive depending on the amount of content in the site collection. Consult with the business to ensure you granularly backup only the most critical data necessary. Microsoft recommends using granular backup for site collections under 85 GB.

Planning for different types of recovery scenarios from these backups is just as critical to your backup schedule in order to meet your desired RPO and RTO targets. Recovering from a disaster using an “Adequate” DR solution entails more manual effort to bring the farm completely up. For example, the configuration and Central Administration databases cannot be restored using out of the box SharePoint tools and service applications must be restarted. Simple item level recovery is also not possible using an “Adequate” solution. These items are what commonly lead to long RTO target intervals for this category of DR solution.

## “BETTER” SHAREPOINT DR SOLUTIONS

A “Better” SharePoint DR solution addresses the short comings of the “Adequate” solution such as the lack of full farm coverage by backups, built in scheduling, alerting on issue, backup grooming, and an inability to meet shorter RTO targets due to the amount of time it takes to restore individual pieces of a farm. The options that can turn an “Adequate” SharePoint DR solution into a “Better” solution can go a long way towards extending DR coverage for your environment and reducing your organization’s exposure to the risks associated with a disaster.

A “Better” solution focuses on providing one solution for creating a configurable farm and granular backup as well as providing flexible restore options. IT professionals cannot be the only people involved in creating

a restore plan. All stakeholders must approve and understand the RPO and RTO targets so that expectations can be met when an emergency occurs. A fast and repeatable restore plan is necessary in “Better” solution.

When considering the use of a third-party SharePoint backup/restore tool, it is important to select a tool that addresses your environment’s specific DR targets and platform requirements. Cost is certainly one important consideration, but there are also many other questions that should be asked.

Here are just a few:

- ▶ Can the tool meet your RPO and RTO targets? Can it back up your farm frequently enough? Is it able to restore your farm quickly enough?
- ▶ What pieces of a SharePoint farm can the tool cover? Can it back up your search index? Can it backup custom code?
- ▶ What pieces of a SharePoint farm is the tool not able to cover? Can you do it all with one tool, or will you still need to implement additional solutions to attain the desired RPO and RTO?
- ▶ What other resources are required to leverage the tool? Are additional servers or specialized hardware necessary to install the tool in your environment? Do you need to configure SharePoint, or another application or system, in a specific way in order to use the tool?
- ▶ How does the tool use your storage? Will you need to purchase additional (or perhaps specialized) storage in order to implement it?

As mentioned earlier, another target of a “Better” SharePoint DR solution is the desire to meet shorter RTO and RPO targets; that is, to restore SharePoint functionality in less time. One way this can be achieved is through the use of more effective backup and recovery tools that cover more of your SharePoint farm, as explained above. Additional opportunities for meeting reduced RTO and RPO targets exist, though, through improving existing infrastructure. Simply improving network and storage speeds, for instance, commonly reduces the amount of time it takes to both capture a backup and restore from it. Shorter restore times translate directly into shorter RTO intervals, while improved backup speeds can make it possible to satisfy more aggressive RPO targets.

## “BEST” SHAREPOINT DR SOLUTIONS

A “Best” SharePoint DR solution is one that can be expected to deliver near or actual real-time recovery capabilities and meet the smallest of RTO targets. Such a solution is capable of preventing loss to all but the most recent of data, if any loss is even tolerated, in order to meet the most stringent of RPO objectives. Just as a luxury car is built to exacting standards, a “Best” SharePoint solution provides extreme coverage for your entire environment and outstanding performance for DR situations that demand it. As you might expect, the luxury car analogy holds true when it comes time to budget for such a DR solution, as well.

Similar to a “Better” DR solution, there are multiple ways to implement a “Best” solution. Solutions may begin by leveraging highly available technologies and systems within a primary data center. Such highly available



approaches may include both SQL Server clustering for automatic storage failover and web front-end load balancing for redundancy.

To achieve geographical distribution and further redundancy, one or more “hot sites” containing their own live SharePoint environments may interact with the primary data center through SQL Server database mirroring or AlwaysOn capabilities, SQL Server transaction log shipping, or even a third-party real-time replication solution. If the primary data center were impacted by a disaster such that it could no longer service requests, users could be redirected to the hot site (or one of the hot sites) to resume their work with a minimum of interruption.

In addition to the use of highly available technologies, “Best” class SharePoint DR solutions commonly leverage one or more additional backup solutions that are pulled from the field of “Better” class solutions. These redundant backup and restore systems not only act as another layer of insurance against catastrophic loss, but they are often leveraged by administrators to service the more mundane, non-catastrophic restore requests that users commonly generate (to recover a lost or deleted document, for example).

In the majority of cases, “Best” class SharePoint DR solutions are most effectively implemented when they are designed into the SharePoint environment rather than added to it after the environment has “gone live.” The more a SharePoint environment matures, the more difficult it becomes to retrofit it for the type of enhancements and technologies a “Best” solution demands without incurring substantial outages or down time. After the-fact implementation costs are typically much higher, as well, since they often include substantial new hardware and licensing purchases on top of those that may have already been made.

## A FEW WORDS OF GUIDANCE

The array of tools and techniques that can be leveraged during the construction of a SharePoint DR strategy can seem somewhat dizzying, but there are a few general principles you should consider to help steer your efforts towards a solution that’s best for you and your environment.

First, remember that any hardware and software you select for your SharePoint DR solution should be chosen on the basis of quantifiable RPO and RTO targets. Many tools and systems make substantial promises with regard to backup, recovery, and DR, but the only true measure of success is your solution’s ability to meet established and agreed-upon DR targets and service level agreements (SLAs).

Remember that most SharePoint environments depend upon other services and systems for proper operation, such as AD, SQL Server, networking hardware, and more. The RPO and RTO targets of these systems directly impact your ability to meet recovery goals that are established for a SharePoint environment. A SharePoint environment that is restored without these systems is not considered operational from the perspective of those who depend on it. It is imperative that any SharePoint DR strategy that is proposed acknowledge and factor-in the recovery strategies of dependent systems appropriately.

Don’t lose sight of that which is most important to your users and business stakeholders: their data. A SharePoint farm contains many moving parts, and the selection of a DR strategy must address servers, services,

databases, supplemental software components, and more. At the end of the day, though, the most valuable commodity in any SharePoint environment is its content. Simply being able to restore a list item that has been deleted from the recycle bins is invaluable to your end users. It is only by ensuring that business data recovery is a primary target in your DR solution that you satisfy users and business stakeholders.

Finally, your solution's ability to meet established targets and SLAs isn't something that can be determined without testing, and this is why robust DR strategies are considered "living entities" that are regularly reviewed, updated, and exercised. A DR strategy should be treated as a full-fledged component of the overall SharePoint ecosystem and supported by your entire organization, not just the IT department. Just as your SharePoint installation and configuration mature over time, grow, or even contract, so may your approach to DR. Regularly test your DR strategy to make sure it works when you need it, and ensure that it remains appropriate for your environment. Through regular formal testing and validation exercises, you acquire the proof and confidence needed to assure stakeholders that your DR strategy meets stated recovery targets in the event of a disaster.

## CONCLUSION

SharePoint DR is more than just backing up your data and building extra hardware. At a minimum, you need to understand your user's needs and expectations for the availability of your SharePoint environment so that you can design an appropriate solution. Many tools and techniques are available for use when constructing a SharePoint DR plan, but it should be relatively clear by now that any DR solution you build or select will be as unique as your SharePoint environment itself. As stated earlier, there is no "one size fits all" DR strategy. Choosing the right solution for your SharePoint environment ultimately entails balancing requirements, performance, and cost.

## A FEW WORDS OF GUIDANCE

SharePoint Backup is the powerful full-farm protection product for Microsoft SharePoint that combines lightning-fast file operations with dynamic compression algorithms to create highly compact backup sets with blazing speed. With SharePoint Backup, administrators have a centralized, intuitive console to schedule, monitor, and manage backups, and users have the ability to recover their own documents.

Try Metalogix SharePoint backup for free!

<https://www.metalogix.com/Downloads/SharePoint-Backup-Download/SharePoint-Backup-Trial-Download.aspx>

## ABOUT METALOGIX

Metalogix provides content infrastructure software to improve the use and performance of enterprise content. For over a decade, Metalogix has transformed the way commercial and government organizations manage terabytes of content to improve knowledge sharing and collaboration. Today, more than 14,000 customers rely on the company's products to upgrade, migrate, organize, store, archive and replicate content on Microsoft SharePoint, Exchange and Cloud platforms. Metalogix has moved more than 50,000 terabytes of content.

### METALOGIX

5335 Wisconsin Ave NW, Suite 510, Washington DC 20015  
sales@metalogix.com | www.metalogix.com | 1.202.609.9100