

Leveraging Cloud Storage with SharePoint and StoragePoint

By Chris Geier
Published: June 2010



CONTENTS

OVERVIEW AND SCOPE	2	STORAGEPOINT CLOUD ADAPTERS.....	6
ABOUT STORAGEPOINT	2	COMPRESSION	6
WHAT IS CLOUD COMPUTING?	2	ENCRYPTION	7
SHAREPOINT IN THE CLOUD?	3	LOCAL CACHE.....	7
CLOUD STORAGE RISKS AND BENEFITS	3	DATA SHREDDING	7
RISKS	3	CONTENT EXTERNALIZATION AND RECALL TIMER JOBS	8
Security	3	USAGE SCENARIOS	8
Legality and Ownership of Data	4	Corporate Records Management	8
Availability and Performance	4	Static or Infrequently Changing Content	8
Viability of Cloud Storage Providers	4	Primary Data Storage Location	8
Bandwidth Costs	5	RISK MITIGATION	9
OFF-PREMISE STORAGE BENEFITS	5	STORAGEPOINT SCENARIOS	10
Significant Cost Savings	5	TIERED STORAGE WITH STORAGEPOINT	10
Reduced Maintenance	5	TRADITIONAL ARCHIVING VS. TIERED STORAGE.....	11
Location Independence	5	RECORDS RETENTION IN THE CLOUD	12
Scalability	5	STORAGEPOINT CONSIDERATIONS WITH CLOUD STORAGE....	13
EXTERNALIZING SHAREPOINT CONTENT WITH STORAGEPOINT	6	FULL-TEXT INDEXING AND CLOUD STORAGE	13
		CONCLUSION	13

OVERVIEW AND SCOPE

Microsoft SharePoint has quickly become the main location for content within many organizations. This leads to a significant increase in cost as all SharePoint content is stored within a Microsoft SQL database as a binary large object (BLOB). StoragePoint provides the means for an organization to externalize this content to a location outside of the SQL database such as a file system, SAN, or even a cloud storage provider.

Cloud storage platforms can provide significant reduction in cost for IT departments by offloading the acquisition cost and monthly operating cost to the cloud storage provider. Using StoragePoint to externalize content to a cloud storage platform can help an organization free content that would normally be constrained in a SQL database and leverage a cloud storage vendor to reduce their overall storage costs and footprint within the organization. This paper examines the benefits and barriers that an organization can expect to face when leveraging a cloud storage platform.

ABOUT STORAGEPOINT

Metalogix StoragePoint is an easy-to-install, remote BLOB storage and archive solution that allows you to consolidate and optimize SharePoint storage. StoragePoint improves performance, scalability and compliance requirements while decreasing overall storage and administrative costs of your growing SharePoint environment. Use StoragePoint and gain the flexibility to store unstructured SharePoint content on virtually any tier of storage device including SAN, NAS, and cloud storage platforms while improving your search, index and back up windows by 95%

WHAT IS CLOUD COMPUTING?

Cloud computing is Internet-based computing, whereby shared resources, software and information are provided to computers and other devices on-demand, like the electricity grid. Cloud computing describes a new supplement, consumption and delivery model for IT services based on the Internet, and it typically involves the provision of dynamically scalable and often virtualized resources as a service over the Internet.

http://en.wikipedia.org/wiki/Cloud_computing

SHAREPOINT IN THE CLOUD?

SharePoint in the cloud will mean different things to different people and conjure up a variety of imagery. There have been significant discussions about SharePoint in the cloud with the SharePoint 2010 launch. What it means to StoragePoint users is simple; we help you with more options to make SharePoint in the cloud easier. When you talk to different organizations about using the cloud, the first reactions are about control and security. With common implementations of SharePoint in the cloud you are presented with only one real option; a hosted model where everything is in the cloud. This can be a shared hosting model or dedicated. To many this feels like handing over everything to someone else and they are left feeling like they have little control over their environment.

At some level people still want to be able to see, and touch their environment even if the ability to do so is not a tangible benefit.

By using StoragePoint with SharePoint you can host your SharePoint servers locally, which can provide you with the grounded feeling of having ultimate control and being able to see and touch your environment. At the same time, you are taking advantage of cloud based storage and all it provides you in terms of scale, availability and cost. Using StoragePoint's BLOB externalization you can have the best of both worlds; local control with cloud flexibility and cost.

You can even mix and match, where some of your data will be hosted locally while you can have some data pushed out to the cloud. The concept of mix and match allows your organization to start slowly with cloud-based storage and expand as your implementation grows and your comfort level increases.

You can also reverse the flow, where the servers are in the cloud and the storage is local. This gives your organization control of the data but gains in cost reduction by taking advantage of a dedicated hosting model for all of your SharePoint needs. The key to this is the provider that allows for a dedicated model where you can install the StoragePoint solutions.

CLOUD STORAGE RISKS AND BENEFITS

Cloud computing provides excellent opportunities for enterprise IT departments to gain operational efficiencies and significantly reduce overall costs. While offering significant benefits in the form of cost

savings and scalability, there are additional risks to consider as well. The decision to leverage cloud storage must be based on solid information. Reading through the next sections will help you better understand these decisions and provide information you should know.

RISKS

SECURITY

Leveraging a cloud platform to store corporate information introduces a new surface to defend from attackers. Enterprise data is now being transmitted across the internet to a third party for storage. Enterprises have traditionally been wholly responsible for the creation, retention, and disposition of content. Backup copies were typically shipped offsite via physical media or transferred through a private network. Using a cloud storage provider expands the attack points available to an attacker. By removing content from the control of IT staff, a company is increasing the risk that it could be compromised. If an organization decides to leverage a cloud

storage provider they will need to take steps to mitigate the risk during both the transmission and retention of data. Typically mitigation of this type involves encryption of data prior to transmission as well as on disk when stored.

LEGALITY AND OWNERSHIP OF DATA

Legal challenges exist when externalizing content from within the confines of your data center to a cloud provider as data located in different jurisdictions is subject to different laws. Additionally using a cloud provider has the potential to cross international borders and may be subjected to different international laws. Understanding where your data originates, where it will be stored and under what jurisdiction it may fall are critical factors to consider when making an informed decision to use cloud computing and what cloud computer provider to use.

AVAILABILITY AND PERFORMANCE

While the internet has become a highly stable environment, there is still reason for concern when relying on it and storing content off premise. Latency issues as well as outages can cause issues when attempting to access any off premise data. Companies need to account for this when determining what content should be stored in the cloud and which should be hosted on premise. Companies also need to take this into account when planning for network access, ensuring connections to the internet and their associated redundancy match those that they need to provide to the business for SharePoint availability. Although it is rare for a storage provider to be unavailable, the impact of this must be properly assessed. When customers implement cloud based storage and/or computing they are often provided with an availability service level agreement (SLA) which defines what the up time and accessibility will be of their application and/or data. They will also guarantee that the system or content is not lost due to any action or failure on their part. This does not however guarantee that they will be able to recover any data that the customer deletes either by accident or by design. The ability to delete information that may need to be kept necessitates the need for backup of data in some manner, either through replication or conventional means. Backing up data from the cloud has inherent issues, the primary of which are bandwidth in terms of time to backup, and additionally cost of bandwidth use. Through the use of the recycle bin and StoragePoint through the use of the BLOB retention policy have mechanisms that help mitigate these concerns.

VIABILITY OF CLOUD STORAGE PROVIDERS

When entrusting enterprise content to a third party you must ensure you are comfortable with the viability of the provider. Storing terabytes of content with a third party can be a challenge to move from one vendor to another or even back into the internal environment thus you need to ensure that your organization has a clear exit strategy when leveraging a third party to store its corporate records. With StoragePoint you can selectively store SharePoint content with a cloud based storage provider. Additionally with StoragePoint's shallow migration you can easily migrate content from on-premises to cloud, or from one cloud provider to another, easily and seamlessly to the users without an interruption in service.

BANDWIDTH COSTS

Most cloud storage vendors have costs associated with both incoming and outgoing transfer of data. Careful planning is needed to ensure you understand your projected traffic patterns to and from the storage provider. If a large user base is consistently accessing content stored off-premise it could negate many of the cost savings benefits. This makes tiered storage models for use with SharePoint an appealing option where you would only put rarely accessed materials or those that are for retention reasons only up on Cloud. We will discuss examples a bit later.

OFF-PREMISE STORAGE BENEFITS

SIGNIFICANT COST SAVINGS

Cloud storage does not carry with it the hefty acquisition cost that on-premises storage requires. Monthly operating costs are also dramatically reduced as a company does not have to staff a data center nor pay for electrical and additional cooling costs associated with the equipment. Equipment failure such as disk or other related hardware is the responsibility of the cloud storage provider. This reduces the need to warehouse additional hardware as standby equipment or carry on additional service contracts.

REDUCED MAINTENANCE

Miscellaneous maintenance tasks are not necessary as the cloud storage provider handles these items as part of their service offering. This allows an organization to maintain its current staffing levels without hiring new resources as SharePoint usage and content storage grows.

LOCATION INDEPENDENCE

Accessing data from anywhere there is an internet connection can be quite beneficial to an organization. Hosting information in the cloud allows an organization to consume that data from anywhere. Hosting data in the cloud can reduce the need for complex data replication scenarios to ensure important records are hosted offsite. Cloud based storage can also decrease the overall recovery time in the event of a system failure or disaster.

SCALABILITY

Cloud storage can provide virtually limitless storage to an organization. Ordering additional hardware and installing it into a local data center is not required when leveraging a cloud storage provider. More space is allocated as more content is created. Additionally the reliability and availability of the medium is built into the architecture, thus not requiring additional planning.

EXTERNALIZING SHAREPOINT CONTENT WITH STORAGEPOINT

SharePoint has become a service almost as essential as email within many organizations. It has changed the way users work within a business and allows them to easily share, search, and manage content with both internal and external entities. While this is helping users function more efficiently, it can impact many teams within the IT departments. SQL DBA's must maintain the SharePoint content databases. Storage teams need to allocate disk space to the DBA's. Backup teams must find ways to backup this data within the allotted maintenance windows. Externalizing that content from the database allows for more options to manage the content and can make everyone's lives easier.

STORAGEPOINT CLOUD ADAPTERS

StoragePoint allows organizations to move content that is confined to a SQL database server to virtually any storage location. If the organization chooses to externalize content to a cloud platform they can begin to realize many of the benefits of this architecture.

StoragePoint has adapters available that allow content to be externalized from SharePoint to many of the major cloud storage providers. These include:

- ▶ Amazon S3
- ▶ AT&T Synaptic Storage
- ▶ EMC Atmos Online
- ▶ Microsoft Azure
- ▶ Rackspace Cloud Files

The extensible architecture of StoragePoint allows additional adapters to be easily developed. StoragePoint also helps to mitigate many of the risks associated with cloud storage through its extensive feature set.

COMPRESSION

Most cloud providers charge for incoming and outgoing bandwidth. Compressing content prior to uploading it to the cloud can save a measurable amount of time and money in the form of bandwidth costs. By way of example and depending on the file type, compression ratios typically run anywhere from 25% to 80% when enabling compression from within a StoragePoint Storage Profile.

ENCRYPTION

Encrypting data is essential to address many of the risks of cloud storage. StoragePoint can encrypt the content before it leaves the corporate network. This protects the data from any vulnerability during transmission as well as when the content is “at rest” in the cloud. Even if records are subpoenaed from the cloud provider, they are still encrypted and virtually unusable without the ability to decrypt them.

Encryption Method for Content in BLOB Store

AES (256 bit) ▾

Encryption Key Passphrase

*Enter a passphrase to be used to generate a key or leave blank to generate a random key. The passphrase entered is **not** saved with the Endpoint.*

Encryption Key

LOCAL CACHE

StoragePoint can be configured to use a local, on-premises cache for files in the process of being externalized. This removes the latency associated with a write to the cloud over the Internet from the end user’s experience. Files that are uploaded by users are quickly written to the cache and control returned to the user. StoragePoint then moves the file out to the appropriate end point via a content migrator agent.

DATA SHREDDING

Where supported by the platform provider or where we have byte level control over the writes, StoragePoint is able to securely delete (shred) files. When a file is deleted from SharePoint and subsequently purged from the recycle bin, StoragePoint will overwrite the file multiple times with zeros and random bytes. It will also rename the file name multiple times before it is deleted.

Deletion Method

Normal ▾

Audit

Normal

Shred

015Z

CONTENT EXTERNALIZATION AND RECALL TIMER JOBS

Many organizations already store large amounts of documents within SharePoint. Getting that content out of the content databases and into an external storage location is a simple matter of running a Content Externalization job from SharePoint central admin. Each of the externalization jobs are attached to StoragePoint Storage Profiles, which are based on criteria you specify. These criteria can specify storing different types of content in different places. Scope of these profiles can be at a top level the content database, and at a low level a content type. Running a job that is set to externalize specified content to a cloud storage provider will move all of the content that currently exists within the scope of the profile off-premise. Conversely, a recall job will bring the content on-premise, providing organizations with key fall back and recovery methods should they grow uncomfortable with using cloud storage.

USAGE SCENARIOS

Not all content is created equally. There are some scenarios where it makes more sense to keep data on-premise and others where off-premise storage is the logical choice. Additionally, there may be scenarios where content should stay on premise for a given time period and then be migrated to cloud storage. A few of the primary use cases for externalizing SharePoint content to cloud storage are highlighted below.

CORPORATE RECORDS MANAGEMENT

Enterprise Content Management (ECM) controls the lifecycle of content from creation to disposition. There are many types of records that must be retained for a certain period of time before they can be destroyed; keeping this readily online but in cheap storage makes for an appealing choice. These records are rarely, if ever accessed by any user or system. This makes them an ideal candidate for off- premise storage. Some organizations mistakenly leverage backup tape and other similar methods for retention. This can have negative consequences, as it may require keeping older technology around and ready for use should older content have a requirement to be restored.

STATIC OR INFREQUENTLY CHANGING CONTENT

Content that does not change frequently is an excellent candidate for cloud storage. Documents such as multimedia files, scanned images, and PDFs are essentially read only documents and not subject to collaboration or frequent edits. The bandwidth requirements will be less than half of collaborative documents types such as Word or Excel files. When a user opens a document of that type from the cloud for editing they are essentially pulling the entire document to the local machine and then uploading it back to the cloud.

PRIMARY DATA STORAGE LOCATION

Using a cloud storage provider as your primary data location can be an effective way for small to medium-sized organizations to benefit. This may save thousands of dollars in hardware acquisition costs as well as ongoing maintenance. The need for backup media can be reduced significantly as the bulk of the storage is off-premise and backed up by the cloud provider. Please see the Cost-Benefit Analysis section below for a more detailed explanation.

StoragePoint features or actions that can be taken to mitigate the risks associated with storing data off- premise.

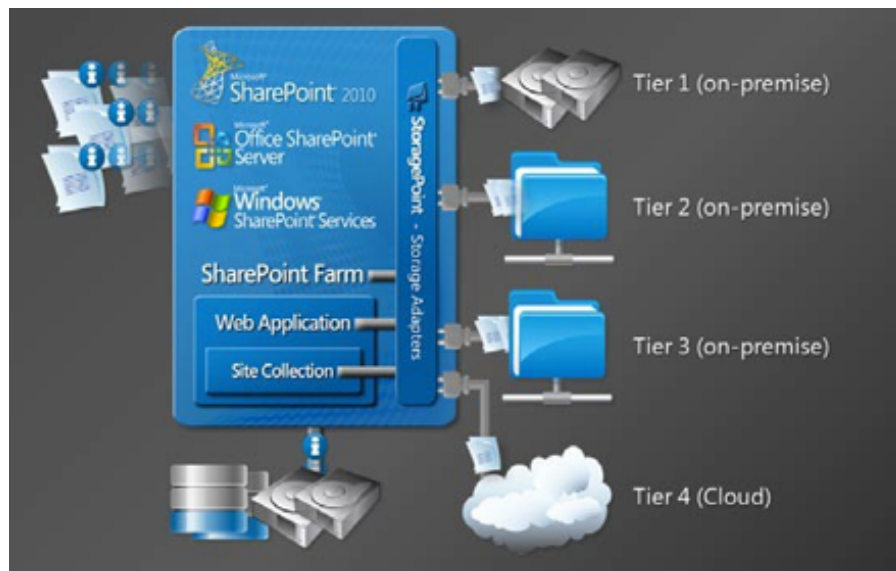
Risk Factor	Mitigating Action(s)	Notes
Secure Transmission Of Documents	Enable Encryption in the Storage Profile	<p><i>When encryption is enabled the content is encrypted prior to transmission to the cloud storage provider.</i></p> <p><i>*Many cloud storage providers allow SSL connections as well.</i></p>
Securing Content on the Cloud	Enable Encryption in the Storage Profile	<p>If the cloud provider’s security is compromised, the attacker would need to decrypt the content if it was going to be useful to them. StoragePoint provides support for AES 128 and AES 256 bit encryption.</p>
Legality and Ownership of Content	Enable Encryption in the Storage Profile	<p>Assuming a cloud provider is ordered to turn over records and data from their systems, it will still be encrypted. This can help to prevent agencies from circumventing a corporation’s own legal department when attempting to gain access to corporate records.</p>
Availability of the cloud provider	<p>Fail safe measures in StoragePoint</p> <p>Be selective about content stored on the cloud</p>	<p>If any endpoint becomes unavailable the content is simply written to the content database. Once the endpoint becomes available again it can then be externalized via externalization timer jobs.</p> <p>Externalizing content that is not used on a daily basis will also help to keep users working during provider outages.</p>
Latency	<p>Enable QoS throttling</p> <p>Asynchronous operations and Cache</p>	<p>Latency will always be a concern as you are transmitting and receiving information over the internet. Data transmission speeds are higher than ever and continue to increase. The risk here is minimal.</p> <p>Configure StoragePoint to perform its cloud operations asynchronously. This allows for quick local writes, with the content moved to the cloud after control is returned to the user.</p>
Viability of the cloud provider	Choose a well-established company	<p>Economic pressures can affect businesses large and small. When selecting your cloud storage provider, be sure to choose a company who can withstand these pressures. Carefully review the terms of service and service level agreements before committing your data to a third party.</p>

<p>Bandwidth costs</p>	<p>Compress Content</p> <p>Limit to long term storage</p> <p>Limit Full-Text Crawls</p> <p>Enable BLOB caching on the SharePoint WFE</p>	<p>Most cloud providers charge for incoming and outgoing bandwidth. While this is a nominal charge it can add up quickly when an organization has multiple users accessing the same files. StoragePoint can compress uploaded content prior to transmission. This can reduce bandwidth usage by up to 80 percent. If you treat the cloud storage platform as an archive then much of the content will likely never be retrieved.</p>
------------------------	--	--

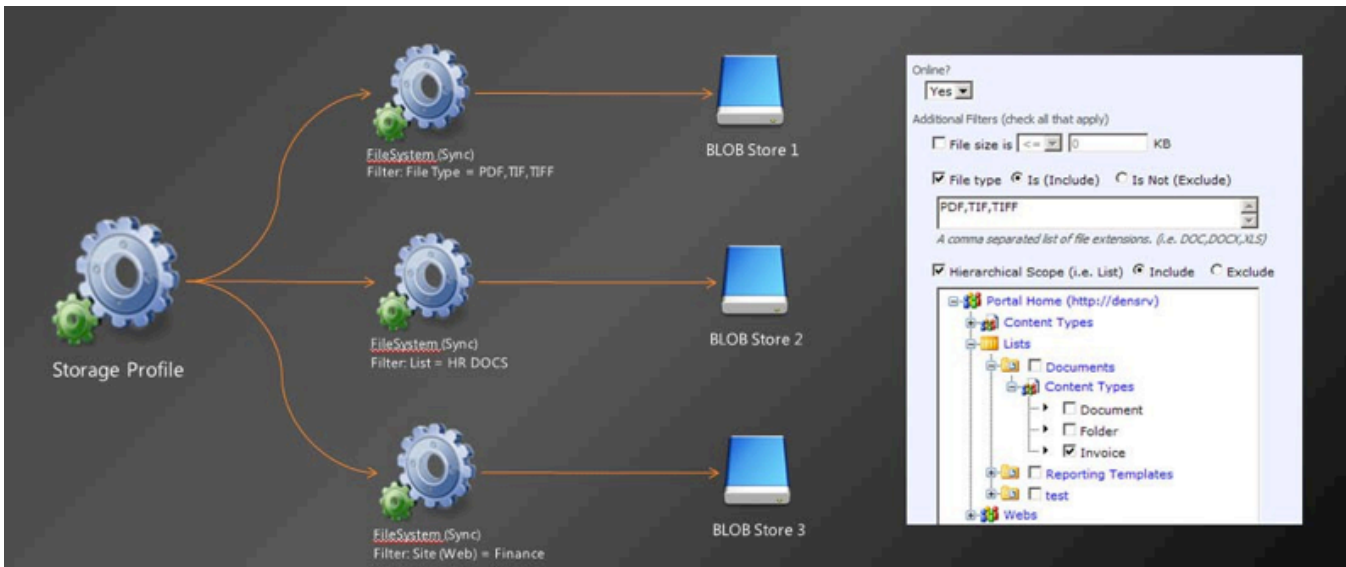
STORAGEPOINT SCENARIOS

Tiered Storage with StoragePoint

A typical enterprise implementation of SharePoint will be housed on or make use of expensive storage because of the heavy reliance on SQL. However with the broad spectrum of content that is typically resident in SharePoint not all of it warrants the use of such high end storage. Every company has different taxonomies and classifications of data. Any solution that helps to prioritize or tier the content must be flexible and tier on their terms.



With StoragePoint you can define only one or many rules that govern the placement or movement of content from one tier to another. Each of these rules is granular. With StoragePoint, companies can determine their own criteria to move content between tiers.



Criteria can be content type, location, size, metadata properties or even age. Content moving based on aging can be a very important part of an overall tiered strategy



Traditional Archiving vs. Tiered Storage

Typical SharePoint archival solutions will migrate content to another location. This could be on-premises, off-premises, or even tape or other media. When archived content needs to be utilized, it is first migrated back to primary storage and then available for consumption by the requesting user or application. This can often be problematic but has been a necessary choice with the significant cost of the online storage that has been used for SQL. With StoragePoint and the ability to appropriately tier storage based on content requirements, organizations can now get past these limitations by placing “archived” content in an appropriate low cost medium such as NAS or Cloud based storage.

With StoragePoint all SharePoint features and functions as well as third party applications that use supported API's still behave as if the content is still contained within the SharePoint environment. This provides a number of benefits over traditional archiving solutions. Search results are comprehensive

and meaningful, without the need for 3rd party solutions or customizations. Users do not need to search an archive volume as well as their normal search scopes to find the information they are looking for.

Records Retention in the cloud

Many organizations must keep certain records for long periods of time, up to 7 years plus. This presents a significant administrative overhead. Take the following situation:

A company has a records retention policy of 7 years and they regularly archive content from SharePoint or other platforms to backup tape. At a later point in time, 6 years, they need to retrieve a certain set of content back. What are the chances that the current technology they are using, either hardware or software, will be able to support a restore from technology 6 years previous? What are the probabilities that they would be able to successfully find and retrieve all the necessary content?

With StoragePoint you can keep all records that have retention requirements online and accessible while utilizing low cost off-premise cloud storage. Each piece of content can be flagged or somehow designated through any of the previously mentioned methods and moved to cloud based storage where it can be easily accessed when needed.

Increase SharePoint availability with Cloud Storage

One possible problem facing organizations with SharePoint is regularly monitoring available disk space

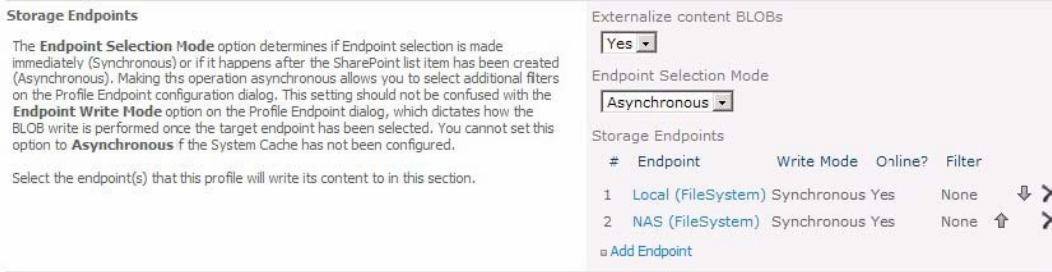
on the SQL server that stores the SharePoint content database. Too many times this has caused outages for those organizations when the disks near or reach capacity. StoragePoint and the infinite storage capacity of cloud based storage can significantly help this situation.

With StoragePoint you can leverage asynchronous endpoint selection which can take advantage of more than one endpoint in a profile. StoragePoint can use a primary endpoint that can be your regular on- premise storage and then use a secondary endpoint if the first is unavailable. This secondary endpoint can easily be cloud based storage with limitless capacity. In this scenario the cloud based storage would only be used when necessary. Additionally with StoragePoint's shallow copy ability all content that was written to the cloud can be easily migrated back to the premise based storage when it becomes available. StoragePoint will use endpoint capacity monitoring that will take an endpoint off line when it gets near capacity.

Endpoint Status

The screenshot shows a configuration window for 'Endpoint Status'. It includes a 'Status' dropdown menu set to 'Online'. Under 'Generate warning notification if:', there are two checked options: 'there is less than 10 MB % of free space' and '25 or more successive errors are encountered'. Under 'Automatically take endpoint offline if:', there are two checked options: '25 or more successive errors are encountered' and 'there is less than 5 MB % of free space'. There is also a section for 'Send Offline Notifications to:' with a checkbox for 'Use Notification Defaults' and a text area for 'Additional Contacts'.

Specify criteria for endpoint capacity management.



Add secondary endpoints to the storage profile. This secondary endpoint can be used in conjunction with the primary or as a fall back endpoint in case the primary goes offline for any reason.

STORAGEPOINT CONSIDERATIONS WITH CLOUD STORAGE

Full-Text Indexing and Cloud Storage

Content that is placed on the cloud should not be frequently crawled by the full-text indexer. Incremental crawls should be used sparingly as well. A full-text crawl will amount to pulling down every piece of content and an organization could incur large bandwidth usage charges.

Excluding content from a full-text crawl is configured in SharePoint Central Admin under the Shared SSP administration page. Rules can be set to exclude content based on metadata or site collection. It is important to design a solid taxonomy to indicate to the full-text indexer that this content should be excluded.

CONCLUSION

Cloud based platforms have become an increasingly viable option that many organizations are exploring. However each organization must determine if cloud based storage is a good option for them. SharePoint can be a good starting point to explore the cloud possibilities. Without StoragePoint, the cloud options are not quite as appealing. By leveraging StoragePoint, organizations can make a move to the cloud in gradual steps, while at the same time knowing that many of the concerns that cloud based computing present are alleviated by taking advantage of compression and encryption (both “in transit” and “at rest”). For more information visit www.metalogix.com.

ABOUT METALOGIX

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

METALOGIX

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