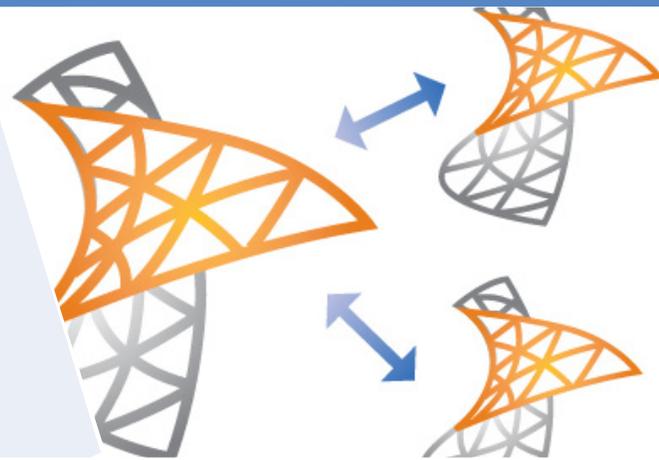


What To Look For In A SharePoint Replication Solution

Metalogix Software



INTRODUCTION

The mission of Metalogix is to enhance the use and performance of your SharePoint platform. As part of that mission, Metalogix believes every business should have a SharePoint continuity plan in place that includes SharePoint replication. To help guide you, here is our guide on what to look for when considering adding replication capabilities to your SharePoint environment.

First, a SharePoint replication solution needs to be comprehensive, enterprise-ready, and scalable. It needs to enable the immediate bi-directional replication across multiple SharePoint farms and support a wide range of content, including permissions, lists, configurations, web parts, version histories, and even metadata. Second, it should be able to trigger replication when a transaction or event occurs, to keep everything tightly synchronized. Third, as an enterprise platform, it should minimize replication bandwidth and the related overhead through compression, scheduling, and should have the ability to scale based on traffic fluctuations. Fourth, it should be able to provide granular control over what is replicated, whether it is site collections, web sites, lists, or libraries. It's an added bonus if it can replication conditionally based on metadata or on content based rules. Fifth, it should also handle "multi-hop" distribution so replication occurs beyond just one pair of SharePoint locations. Finally it should be able to replicate through corporate firewalls without requiring any specific network configuration.

IMMEDIATE REAL-TIME REPLICATION ACROSS MULTIPLE FARMS

For continuous data redundancy or expedited disaster recovery, a SharePoint replication solution should immediately improve business continuity and minimize downtime based on recovery objectives. The solution should provide instant, real-time replication with both high availability and disaster recovery services to provide a complete data protection service to ensure business operations are not disrupted.

TRANSACTIONAL EVENT-LEVEL REPLICATION

Transactional or event-level replication is typically used in server-to-server scenarios that require high throughput to improve scalability and availability, frequent reporting, and data synchronization between multiple sites, keeping your data tightly synchronized between all configured servers. Every time a specific configured event occurs, replication is triggered. It is essential to ensure near real-time replication, since replication occurs every time an action takes place within SharePoint, such as the check-in or check-out of a document.

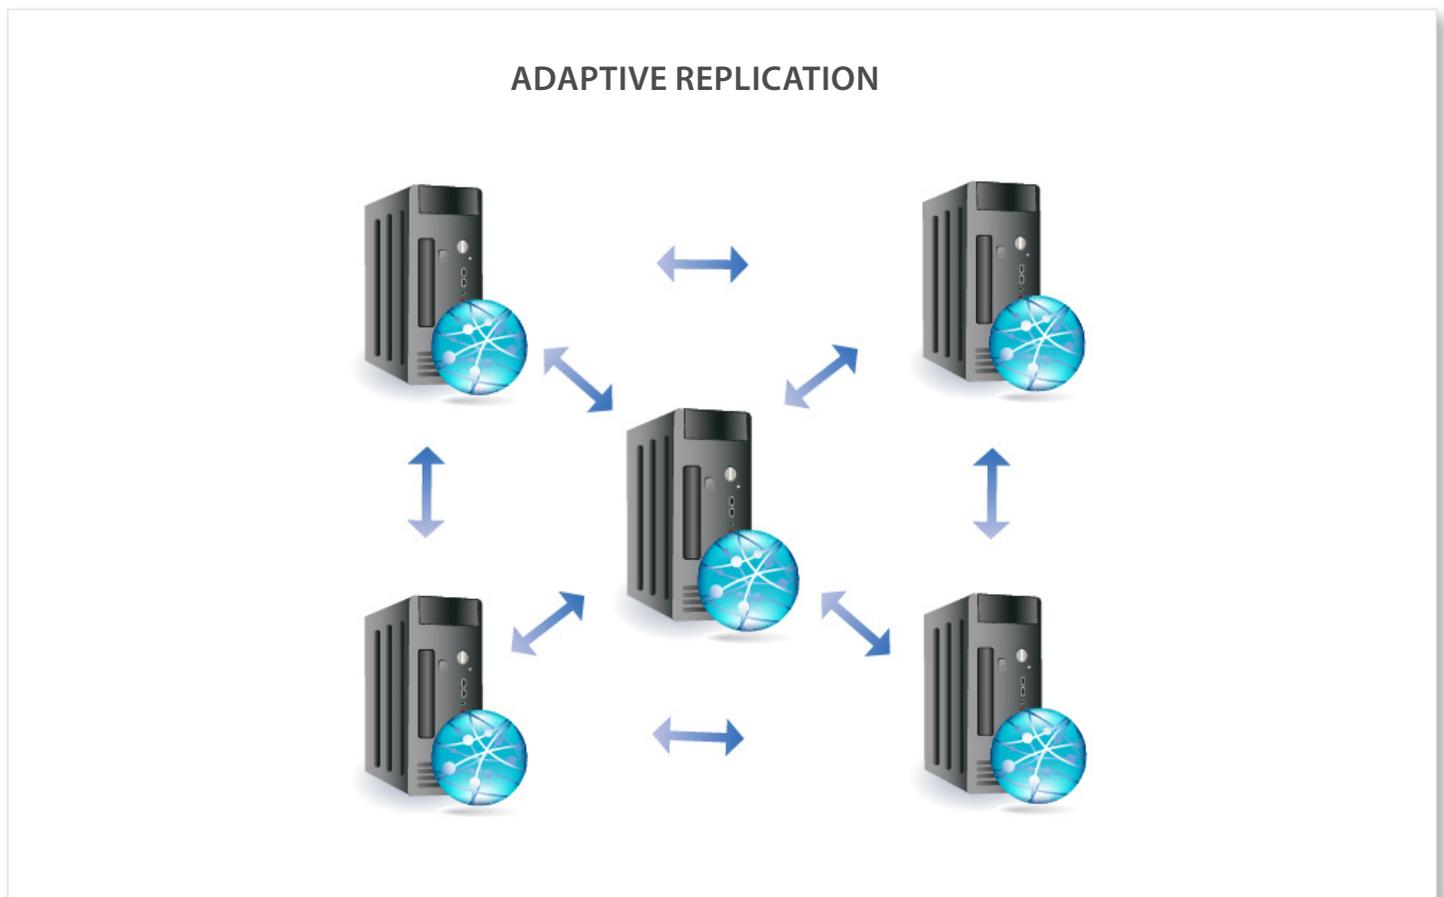
SUPPORT WIDE RANGE OF CONTENT

A successful replication solution should have the ability to bi-directionally replicate more than just documents from one server location to another, between locations and through firewalls. It should also be able to selectively replicate:

- ▶ Content – documents, list items, master pages, site content types, site columns, site templates, and document templates
- ▶ Structure – sites, site collections, libraries (document, picture, slide, ...), lists (discussions, issues, contacts, ...), and folders
- ▶ Permissions – users, groups, and permission levels (including permission settings for sites, lists, and items)
- ▶ Workflows – business processes using SharePoint out of the box, SharePoint Designer, Visual Studio, and Nintex workflows
- ▶ Web Parts – SharePoint out of the box, Bamboo, CorasWorks, OSI, and more
- ▶ Look and Feel – web parts, page layouts, site navigation, views, themes
- ▶ Social Data – notes, tags, and ratings
- ▶ Miscellaneous – alerts, feature activation, recycle bin restores

ADAPTIVE REPLICATION FOR CONTINUITY

Advanced Replication solutions have the ability to “adapt” to changes in network conditions to make sure that data gets synchronized, even if there’s a network outage or a farm outage. If one part of a network goes down, SharePoint replication software can identify a new route and then begin using that to ensure the synchronization continues.



PRIORITIZED REPLICATION

The ability to prioritize SharePoint replication is becoming more and more important as SharePoint farms extend farther and farther out and contain more and more content. Make sure the replication solution you choose lets you manage the queues so that the most important content gets replicated first. Ideally you will want the ability to prioritize based on your content structure, the content itself, or on meta-data.

MINIMIZE SHAREPOINT REPLICATION BANDWIDTH

A smart SharePoint replication solution will allow administrators reduce network bandwidth requirements by dynamically compressing data before replicating or scheduling it when bandwidth is most available. There are three key ways to manage data to optimize a SharePoint replication and a SharePoint replication solution should address all three.

MICROSOFT REMOTE DIFFERENTIAL COMPRESSION (RDC)

Microsoft Remote Differential Compression allows data to be synchronized using byte-level differencing technology to minimize the amount of data that is sent through network connections. Instead of replicating the entire document, RDC identifies, isolates, and replicates only the changed portions of a file, rather than replicating the changed file in its entirety. This dramatically cuts down on the amount of data that gets replicated over the network and allows SharePoint to be deployed across even low bit-rate connections. RDC is most beneficial for synchronizing large files, when only a small amount of data has changed, or if there is a slow, or expensive (i.e. satellite) connection between servers.

ZIP COMPRESSION

All content should be compressed prior to replication. The more the content is compressed, the more bandwidth requirements are reduced.

THROTTLE MODE AND SCHEDULED REPLICATION

Throttle mode and scheduled replication give administrators additional control over bandwidth usage by manually replicating when needed or by scheduling replication ahead of time. Many companies prefer to replicate larger data packets when bandwidth consumption is at its lowest, typically between 2:00 am and 5:00 am.

MICROSOFT BACKGROUND INTELLIGENT TRANSPORT SERVICE (BITS)

Make sure to look for background intelligent transport service (BITS). If a replication package is interrupted en route, BITS will automatically resume the transfer once the network is restored. Built using core Microsoft technologies to transport content from one server to another, BITS uses compressed packages and a restartable protocol for maximum network efficiency. The result is no packet loss and the replication process doesn't need to start over from the beginning.

SELECTIVE OR GRANULAR REPLICATION

For additional flexibility and performance, many firms require the ability to choose what to replicate instead of replicating everything at once. Selective replication lets managers decide what specific content gets replicated and what does not, based on a specific set of attributes or metadata (i.e. replication of changes that occur in a SharePoint Web Application, Site Collection, Web Site, List or Document Library.)

In the case of remote offices, selective replication is critical. If SharePoint managers want only geographic content replicated, they can define it according to specific structure or metadata. On the other hand, if SharePoint managers want to have the entire web application replicated, including the presentation layer, they can make the remote user experience virtually identical to the user experience at headquarters.

RULE-BASED REPLICATION

Control is key when determining what to replicate, and even what not to replicate. Advanced replication solutions should give SharePoint managers the ability to automatically prevent replication based on specific criteria. For example, say a user uploads a document to a document library. It will not replicate to other servers until it has been approved for broader distribution.

BI-DIRECTIONAL REPLICATION

Even the most basic SharePoint replication solutions should be able to support bi-directional replication. Bi-directional replication ensures that content is synchronized in both directions so that content at all locations is up to date. The site structure, permissions, and data should all be bi-directional.

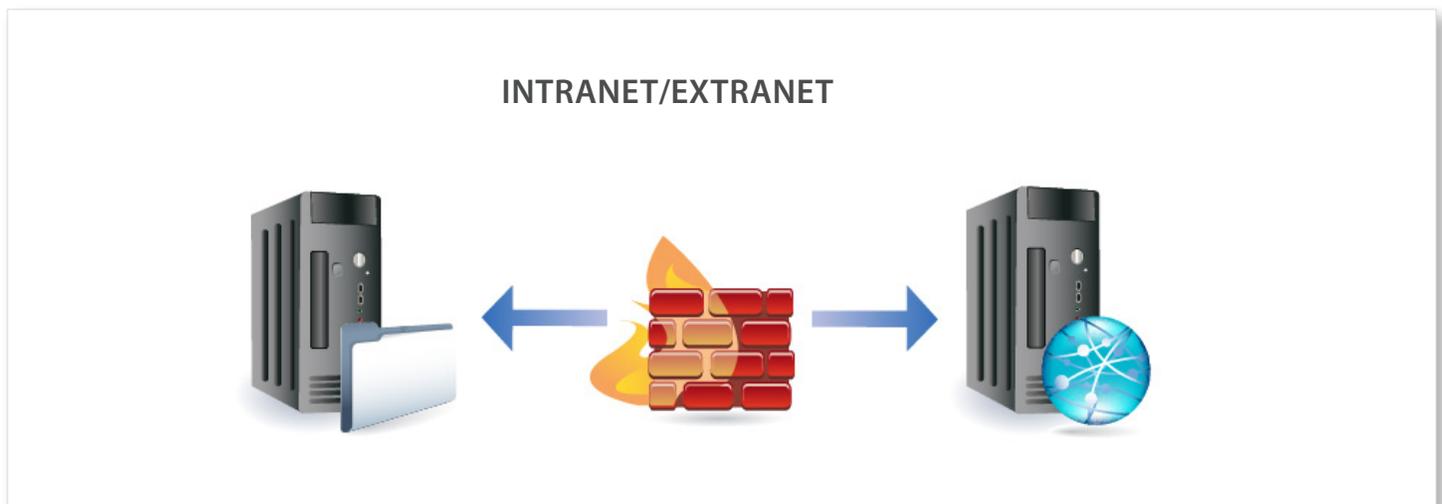


Replication environments that leverage full bi-directional methods provide additional flexibility by avoiding requirements for any single server to be seen as the master server. Any server can function as a master server, in the case of recovering from a disaster, and content, including sites, and site collections can be created anywhere in the environment, and replicated to all configured nodes. Bi-directional replication also ensures that changes made to your remote farms during an outage are synchronized back to the primary farm when it is restored.

REPLICATION THROUGH FIREWALLS

Cross-firewall SharePoint replication must maintain strong security. For that reason, SharePoint managers should be able to selectively replicate specific content and be able to configure bi-directional replication through a firewall, not just one-directional replication.

Many products offer content replication, but only a few offer bi-directional replication through a firewall. Additionally, the ideal replication solution would not require any changes to the corporate firewall (ports) to work. This capability is essential for global organizations that cannot afford to be restricted by environmental and physical boundaries during the course of business.



SHAREPOINT INTEGRATION

Replication solutions that fully integrate into the SharePoint interface provide seamless management through SharePoint, at the site and Central Administration screens. Ideally, managers need a replication tool that does not require dedicated servers and allows the application to scale for clustered environments.

CONFLICT RESOLUTION

Because many users may be working on the same document at the same time, make sure the SharePoint replication solution you choose provides a number of options to deal with conflicts. If two users, on two different servers make changes at the same time to copies of a single document, and then replicate the updated document, at least one server will report a conflict. A successful replication solution allows you to efficiently manage these conflicts and gives managers the flexibility to specify different ways to handle them when they occur. For example, when replicating bi-directionally from Washington to London to Paris, you may want to configure the following conflict handling strategies:

Changes from Washington are always kept, even if older than the change from London. If a conflict is detected in a change made in either London or Paris, the document owner is notified and must select the version to keep.

FULLY INTEGRATED SOLUTIONS

Replication solutions that are fully integrated with SharePoint typically offer the following advantages:

- ▶ Easy to install and configure
- ▶ Interface is part of SharePoint, no interface-related learning curve for users
- ▶ Behave like a regular part of SharePoint
- ▶ Integrated with Microsoft technologies (the Microsoft stack)
- ▶ Better visibility into SharePoint events
- ▶ Increased performance due to using the native SharePoint API
- ▶ Leverage the existing SharePoint admin credentials

WEB BASED SOLUTION

A successful replication solution should be able to replicate all required content using only HTTP or HTTPS.

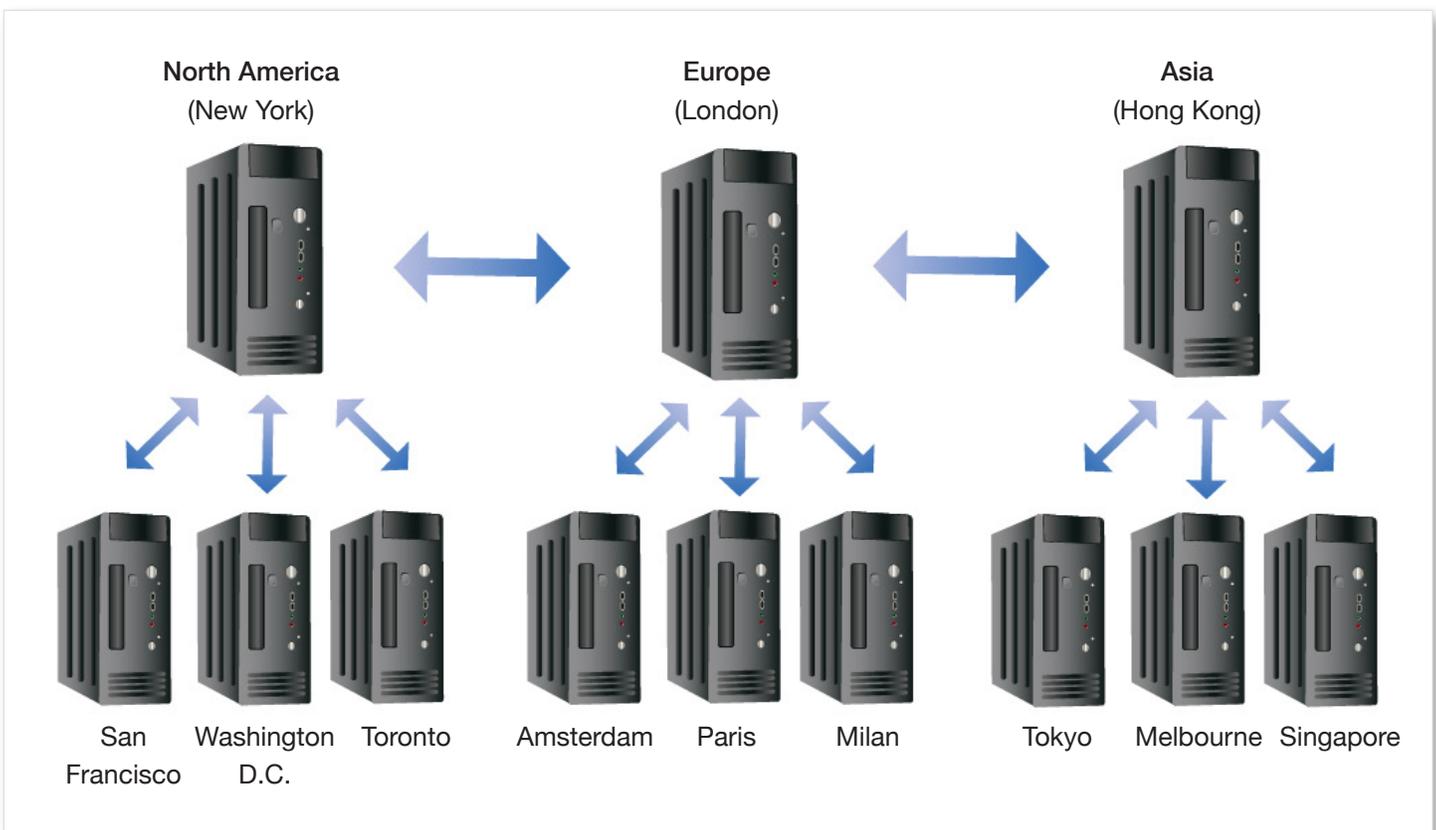
CONFIGURABLE MONITORING AND THRESHOLDS

Look for a SharePoint replication solution with flexible monitoring functionality, with configurable thresholds. Complete monitoring should be available to see the status of any given package and the state of each server at any time. By presenting replication data in a simple and straightforward manner, even the most complex networks become much easier to understand, allowing administrators to accurately regulate bandwidth usage.

Administrators can manage and manipulate replication servers directly by launching the administrative screens, and run various reports on detailed environmental information and thresholds for performance. Flexible monitoring solutions allow administrators to monitor and track replication activity and alerts, visualize the status of server farms and their connections, and display and configure a layout map of all of the farms in a replicated environment.

MULTI-HOP DISTRIBUTION

Advanced SharePoint replication solutions support multi-hop distribution for better efficiency in larger environments. When a corporation has multiple locations involved in SharePoint replication, it is better to set up a multiple stage replication instead of a central hub with all remote locations tied to that single hub. The network paths will differ from hop to hop and it is usually more efficient to make the long haul connections only once and then have a secondary hub to continue to burst information at the local level. Consider the scenario shown in the following graphic:



In this example, the corporate head office is located in New York, with additional North American servers in San Francisco, Washington and Toronto. In Europe, there are servers located in London, Amsterdam, Paris and Milan. In Asia, there are servers located in Hong Kong, Tokyo, Melbourne and Singapore.

It is neither practical nor cost-efficient to replicate a large package of information from the New York server to every other server directly. In this case, the optimal replication structure is a bridged hubs model, as shown in the graphic above. Multi-hop capabilities for replication are a key requirement that allows flexibility for changes to structure and permissions at either end of the chain.

BENEFITS OF A SHAREPOINT REPLICATION SOLUTION

- ▶ Information is continually updated, ensuring that all users have access to the most current versions at all times
- ▶ If disaster strikes, you can quickly redirect your users to the latest content on the remote farm.
- ▶ Servers can be replaced quickly, and the process remains transparent to users
- ▶ Strategic partners can securely access current content
- ▶ Enhanced availability of documents
- ▶ Improved productivity in remote locations
- ▶ Timely distribution of documents to remote locations
- ▶ Selective movement of information, based on business rules

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 | 1 202.609.9100