

Metalogix® Replicator 7.4

Product and Concepts Guide



© 2023 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Metalogix are trademarks and registered trademarks of Quest Software Inc. and its affiliates. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend



CAUTION: A caution icon indicates potential damage to hardware or loss of data if instructions are not followed.



IMPORTANT, NOTE, TIP, MOBILE OR VIDEO: An information icon indicates supporting information.

Contents

About the Product and Concepts Guide	5
Organization of this Guide	5
Metalogix Replicator Documentation	5
Metalogix Replicator Overview	6
Replication Functionality	7
Key Features	8
Concepts	11
Groups, Connections and Map Families	11
Groups	11
Bidirectional Replication	12
One-way Replication	12
Mixed Replication	12
Connections	12
Map Families	13
Queuing and Capturing Events	14
Queuing	14
Capturing	14
Ping Backs and Secondary Events	14
Replication Modes	14
Direct	15
Adaptive	15
Picking the Right Replication Mode	15
Remapping Site Structure	15
Connectivity	15
Performance	16
Differencing Compression	16
Replication Pipeline	16
Replicating Between Editions	16
Look and Feel	16
Workflow Associations and Instances	17
Alerts	17
Web Front-ends	17
Single Web Front-end	17
Multiple Web Front-ends	17
Scenario 1	17
Scenario 2	18
Scenario 3	19
Multiple Replicator Engines	19
Parallel Processing	21
Threads	21
Factors Affecting Threading	22

Enterprise Edition	22
Metalogix Replicator Specifications	25
Editions and Features	25
Enterprise Edition	25
Cross Domain Edition	25
Standard Edition	25
Express Edition	26
Replication Feature Matrix	27
Replication Support Between SharePoint Versions	32
Look and Feel Support Between SharePoint Versions	34
SharePoint 2013 Look and Feel Support	37
Supported	37
Not Supported	37
Workflow Support Between SharePoint Versions	40
Supported Web Parts	43
Metalogix Replicator Reference Architecture	46
Metalogix Replicator Logical System Architecture	47
Metalogix Replicator Physical System Architecture	48
About	49
Contact Quest	49
Technical Support Resources	49

About the Product and Concepts Guide

The Metalogix Replicator Product and Concepts Guide provides you with product-related information such as details about each edition of the product, a reference architecture, and various concepts.

Organization of this Guide

The Product and Concepts Guide includes the following information:

1. About the Product and Concepts Guide
2. Key Features
3. Concepts
4. Metalogix Replicator Editions and Features
5. Metalogix Replicator Reference Architecture

Metalogix Replicator Documentation

Metalogix Replicator documentation ships with the Metalogix Replicator distribution. It is also available at the Metalogix Product Resource Center, <http://documents.quest.com/Metalogix-Replicator>

The Metalogix Replicator documentation set consists of the following guides:

Product and Concepts Guide

The Product and Concepts Guide provides a high-level overview of Metalogix Replicator and describes new features introduced in this release.

Quick Start Guide

The Quick Start Guide guides you through a basic installation and configuration of Replicator. After completing this guide you will have a working Replicator Network where changes made on one web application are automatically replicated to the other. You will be able to move onto the Evaluation Guide to perform various Metalogix Replicator scenario tests on your environment.

Installation Guide

The Installation Guide guides you through installation, uninstall, and upgrading of Metalogix Replicator in various SharePoint farm configurations.

Evaluation Guide

The Evaluation Guide guides you through various Metalogix Replicator scenarios, provided for the purpose of reviewing Replicator within your own SharePoint environment. Prospective Clients are encouraged to use this guide to test Replicator.

Reference Guide

The Reference Guide describes the operations and user interface available to replication administrators through the SharePoint Central Administration and Site Administration pages.

Command-line Reference

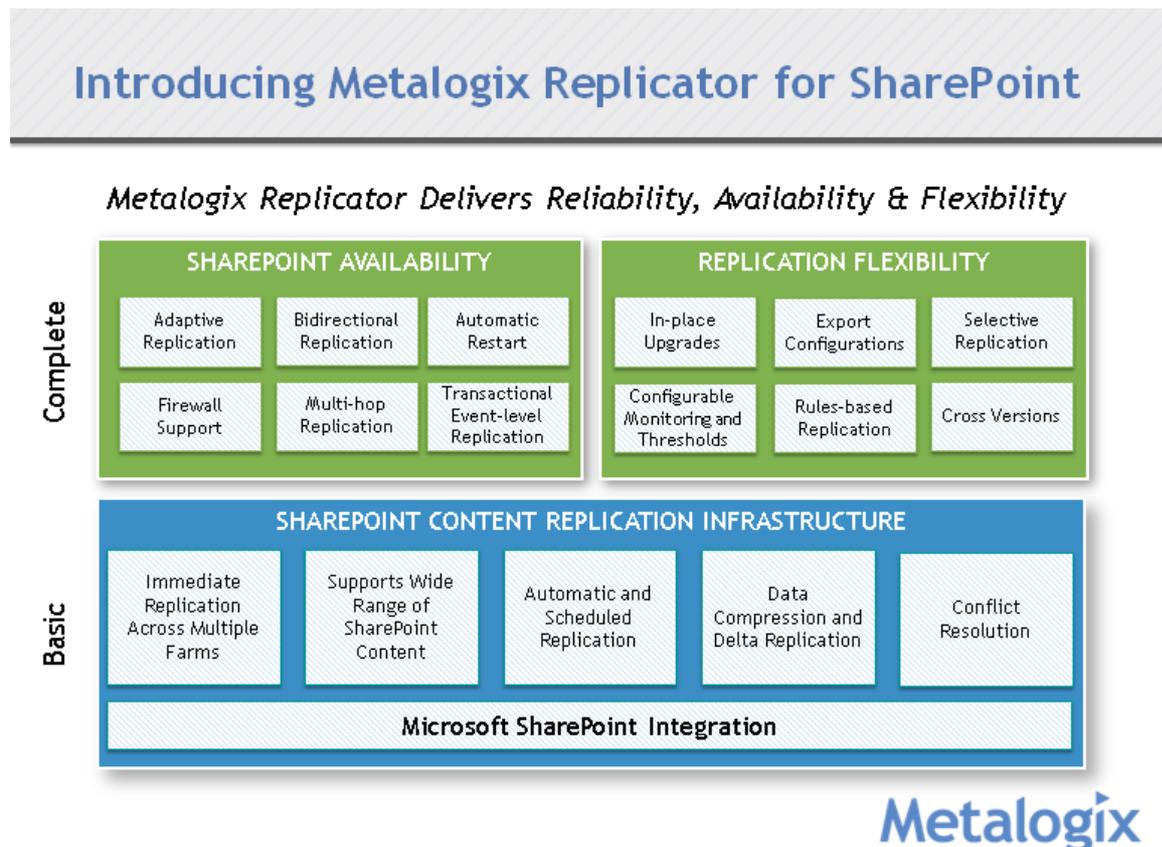
The Command-line Reference describes the PowerShell commands you can use to configure and maintain your Replication Network.

Scenario Guides

Scenario Guides describe specific product features and how to implement and configure them for your SharePoint environment.

Metalogix Replicator Overview

Replicator allows SharePoint lists and events to be replicated to one or more SharePoint servers. This is a convenient and efficient means of making sure data is always available.



Examples of items that can be replicated are document, picture and slide libraries, announcements, links, contacts tasks, issues, custom lists, surveys, Wiki page libraries, project tasks, discussion boards and permissions. Replicator also replicates many SharePoint events, including the addition or modification of web sites and site collections, templates, users and groups.

Live network bandwidth is often at a premium, and networks compensate for this by creating local copies of information on remote web applications. Managing content in this manner allows users local access to regularly updated copies of documents without taxing available bandwidth or network resources. Many replication applications allow users to view read-only copies of documents, without permitting updates from remote sites. When they do permit this, the process requires complex caching servers to be established and maintained. With Replicator, when a file is created on one server and a user is modifying it on another, the servers do not have to communicate with each other every time the documents are locally accessed.

Replicator provides efficient, easy, multi-directional replication. SharePoint lists can be replicated to many systems, and any modifications to documents are automatically updated throughout the whole system. In effect, instead of having isolated Target Web Applications, the entire SharePoint network is interconnected and constantly updated, so that it functions as a single unit, but does not require extra bandwidth or network resources.

Replicator is very easy to access, because it installs itself as a set of solutions within SharePoint itself and is fully integrated. Because it is resident within SharePoint, the entire user interface is directly incorporated into SharePoint's Central Administration site.

Replicator Packages are transferred using the Microsoft Background Intelligent Transfer Service (BITS). Replicator communicates solely using the HTTP or HTTPS protocols. Replicator can also be set to run in offline mode, which replicates Packages directly to a system drive. In this case, only local disk access is used and it is an administrator's responsibility to transfer these packages to the target web application.

Replication Functionality

Replicator replicates objects that have an event, like adding a document, to target web applications. Changes to the system that should not be replicated, such as time zone settings, and changes that are administrative, like web part installation and site templates, are not currently replicated. Replicator also lets you select which events are replicated.

Some users may want the option to replicate document templates, in which case, each remote site uses the same letterhead on their documents. Other users may not want to replicate document templates, in which case, each remote site uses their own unique letterhead on their documents.

Replicator is a feature within SharePoint. There are also features that are enabled that we have not tested in combination with Replicator, or all of the additional features provided by Microsoft.

Version settings are replicated; however we do not synchronize version history, nor guarantee that the version numbers will match on each side, as each web application tracks its own version history. If version numbers are important, Metalogix recommends that you create a new metadata field to track that information.

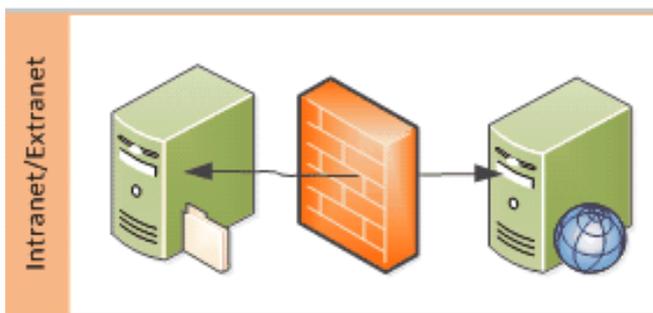
Key Features

Replicator provides these practical and useful features:

- SharePoint document libraries and documents can be easily and automatically replicated to one or more other SharePoint servers.

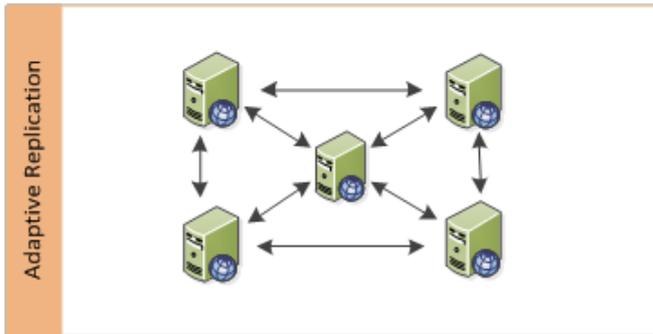


- Data is always available, conveniently and efficiently.
- Documents on source and target web applications are available locally and remotely, without the need for excessive network bandwidth or complex server caching.
- Multiple configurations are possible. This allows for a large variety of internal web application environments, with multi-directional replication.
- SharePoint Farms are interconnected. The process is very efficient, as replication can be configured to occur only when changes are introduced to documents or document libraries.
- With Replicator, SharePoint functions as a unified network without needing extra bandwidth.
- Replicator installs as a solution within SharePoint itself. Because it is resident within SharePoint, the entire user interface is directly incorporated into SharePoint's Central Administration page. Replicator can be accessed through the Application Management tab on SharePoint Central Administration.
- Replicator supports replication when one of the farms is behind a firewall and not accessible to the other. It is the ideal solution for an intranet-extranet scenario, where you want to replicate information from the extranet to your corporate intranet without opening ports or making changes to your corporate firewall.



- Replicator allows for the use of a patent pending transportation model called Adaptive mode. This mode allows your farms to be set up in a dynamic topology which is always changing. For instance, in a ship to shore scenario, the ships talking to one port can later switch to another point of communication, pulling content from the closest server available. With Adaptive applied in this

scenario the destination dynamically and automatically decides where the best source for obtaining data is located. Adaptive mode is also the lightest weight in terms of network traffic - particularly a result of its patent pending micro messaging system which is used to determine the best possible route for information retrieval.



- Replicator uses a patent pending enhancement to Remote Differential Compression. Metalogix Dynamic Compression (MDC), reduces the network traffic when determining the changed content in a file and allows for multiple hops without any global owner.

Concepts

The following section will provide you with an overview of the various concepts associated with Replicator functionality.

Groups, Connections and Map Families

This section explains the relationship between groups, connections, and map families as you configure Replicator. Groups specify which connections are used to replicate changes in a network. Connections specify how web applications interact with each other. Map Families define what SharePoint changes are replicated and how often.

Consider a company with three web applications as follows:

- Corporate Office Portal, the primary intranet site for all employees located in North America.
- London Office Portal mirrors the Corporate Office and serves all employees located in Europe.
- Extranet Portal publishes a limited amount of information from the Corporate Portal that is intended for external use.

We want to design Replicator to replicate SharePoint changes as follows:



Groups

Groups identify which connections in a Replication Network will be used to replicate changes as specified by a map family. When defining groups, consider the following relationships:

- A map family only has one group.
- A group allows you to specify which connections fit within a map family.
- A single connection can belong to multiple groups.

The following illustrations show how group membership affects what is replicated between connected web applications. The red lines indicate that the connection is a member of the group used by the map family.

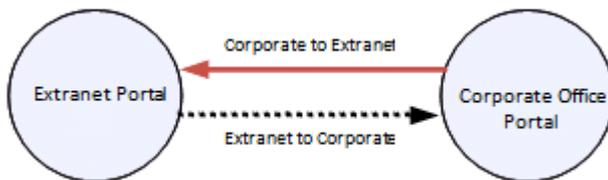
Bidirectional Replication

For bidirectional replication, both connections for a pair of web applications must belong to the same group. When you create the map family for Branch Offices Group on Corporate Office Portal, the map family is replicated to London Office Portal. Since the Branch Offices Group exists on London Office Portal, the map family also captures changes made to London Office Portal and replicates them along the same group to Corporate Office Portal.



One-way Replication

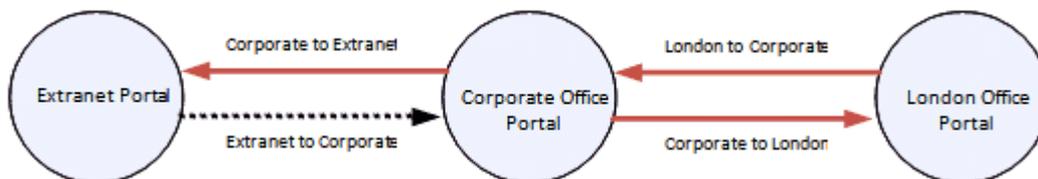
For one-way replication, the map family's group must only exist on the outbound web application. When the map family for the Extranet Group is replicated to Extranet Portal that group does not exist on extranet and so no changes are captured and replicated from Extranet Portal.



Mixed Replication

We can combine one-way and bidirectional replication in the same network. If you created the map family on London Office Portal, then it would be replicated to Corporate Office Portal. Corporate Office Portal has two outbound connections that belong to the same group; ensuring changes made on Corporate Office Portal are replicated to both Extranet Portal and back to London Office Portal.

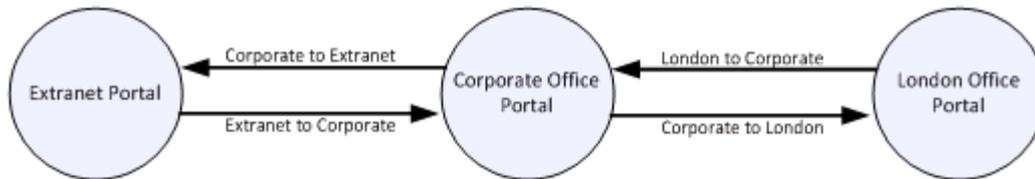
Since the Extranet to Corporate connection does not belong to the group, changes made on Extranet are not replicated back to Corporate Office Portal.



Connections

Each pair of web applications must have two Replication Connections. Replication Connections specify how one web application will talk to the other and also ensure each web application is authorized to communicate with the other.

As SharePoint content is being replicated between web applications, the Replicator instances for each web application are actively communicating with each other and delivering status updates. This is why each web application in a pair must have an outbound and an inbound connection to the other web application.



Map Families

Map families define what SharePoint changes will be automatically replicated and how often that replication will occur. In most cases, there is only one map family defined for a Replication Network, but there are also valid reasons for having multiple map families.

When you create a map family, you specify its scope. The scope determines which events are available to capture. You should select the scope of your map family to represent the events you want to capture. For example if you want to capture and replicate new site collections, then you must create your map family with a Web Application scope.

Event Group	Selected	
Web Application	<ul style="list-style-type: none"> Site Collection Add Site Collection Delete 	Web Application Scope Events
Site Collection	<ul style="list-style-type: none"> User Add User Update User Delete Group Add Group Update 	Site Collection Scope Events
Web Site	<ul style="list-style-type: none"> Web Permission Level Add Web Permission Level Update Web Permission Level Delete Web Permission Add Web Permission Update 	Web Site Scope Events
List	<ul style="list-style-type: none"> List Add List Update List Delete Recycle Bin Restore 	
List Item	<ul style="list-style-type: none"> List Item Add List Item Update List Item Delete List Item Checkout List Item Checkin 	
Workflow	<ul style="list-style-type: none"> Workflow Association Add Workflow Association Update Workflow Association Delete Workflow Instance Cancel Workflow Modify 	

Queuing and Capturing Events

Queuing and capturing are two different ways in which SharePoint events are replicated.

Queuing

When you use Replicator, you are **manually** selecting what you want to replicate. Replicator goes through your SharePoint environment, at the levels you specify, looking for items that match the events you queued, for example list item updates. Queuing is useful for initial replication scenarios or when you want to manually promote content from one site to another, such as intranet to extranet or development to testing. With queuing you can replicate the entire contents of your web application to another web application so that they are consistent from the start, even before live replication starts.

Capturing

Capturing is the **automatic** replication of changes made to a web application. As changes are made, they are captured and marked for replication.

Ping Backs and Secondary Events

It is important to note that ping backs and secondary events may occur during regular replication queuing and capturing. A ping back is when the same event type is sent back to the source. For instance, when an inbound list add event results in an outbound list add/update event for the same list.

It is expected behavior to see occasional ping backs in normal replication operation. These ping backs are not destructive as they are simply returning the same content back to the source.

The reason you may see occasional ping backs is because replicator is unable to determine if the event was caused by replicator applying an event, or caused by an end user. In this case we err on the side of caution and allow the event to ping back so we don't inadvertently stop a user generated event.

Secondary events are different event types being triggered by an inbound event. For instance, an inbound list add event could trigger outbound view update events, site navigation update events and webfile update events. We try to associate the secondary events to the inbound event and stop their replication, but if we can't associate the events to an inbound event we let it replicate as it may have been triggered by a valid user action.

Replication Modes

The replication mode determines how servers in this group send and receive replication packages from each other.

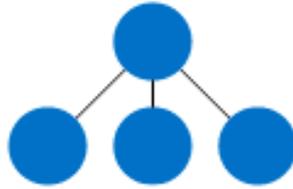
Direct

Direct replication mode uses only one path between any two web applications in the Replication Network.

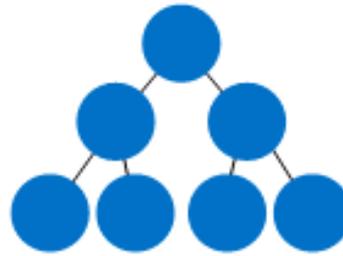
Examples of Direct Topologies



Point to Point



Hub and Spoke

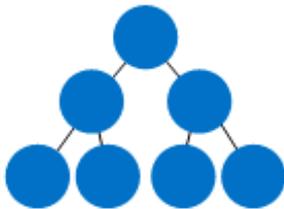


Cascading Hubs

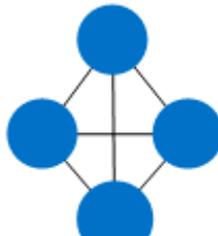
Adaptive

Adaptive replication mode selects the best source for inbound replication packages. When starting a transfer, Replicator picks the best connection to use from the available connections with the packages.

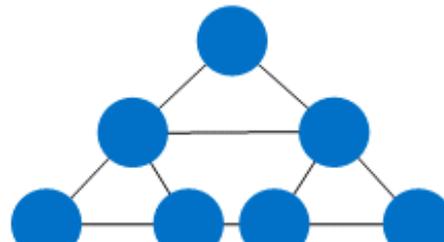
Examples of Adaptive Topologies



Cascading Hubs



Full Mesh



Cascading Mesh

Picking the Right Replication Mode

This section will cover some of the considerations you should take into account when deciding when to use Direct or Adaptive mode for replication.

Remapping Site Structure

The custom binding feature is used to remap site structure when replicating. You can use custom binding to replicate from <http://corporateoffice/public> to <http://extranet/> or from <http://londonoffice/sites/finance> to http://corporateoffice/sites/finance_london.

Custom binding can only be used in Direct mode.

Connectivity

In the event that there is intermittent connectivity between web applications, then adaptive mode can help by finding the best available path to transfer packages in a replication topology that includes more than two farms. With direct mode, replication will not continue until the connection is re-established.

Performance

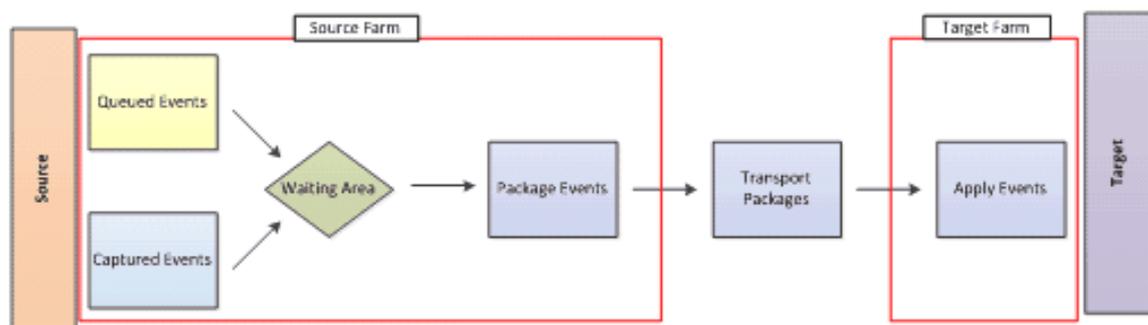
Adaptive mode can be faster than direct mode by better distributing activity on the network. Furthermore, when implemented across multiple farms, adaptive mode can use less system resources for creating replication packages, since packages are reused instead of recreated at each step of a multi-hop replication. However, Direct mode can send fewer notifications to the target farms than adaptive mode because there are fewer paths in the replication network.

Differencing Compression

When using Differencing Compression for a connection, Adaptive mode caches packages on the source and target to anticipate the changes that need to be replicated and therefore further reduces the bandwidth used to replicate documents.

Replication Pipeline

Both queued and captured events are set aside in the same space, awaiting replication. Depending on the replication schedule defined in the map family, Replicator looks for events in the waiting area and starts to process them into packages.



In the diagram above the source farm has had several events manually queued for replication by an administrator, and has also automatically captured several events for replication based on changes made on the web site level by an end user. Both types of events have been placed in the same waiting area in the replication pipeline. Both type of events will be packaged, transported and then received on the target where they are applied.

Replicating Between Editions

You can run different Replicator editions in your replication network. For example, in a hub and spoke topology, the hub can run Replicator Enterprise or Standard edition for immediate and multi-threaded replication. The spokes, with less traffic and changes, can run Express edition.

Look and Feel

Look and Feel Replication is only possible when both the source and target can process look and feel events. For example, you cannot replicate look and feel changes from a Standard edition source to a Express edition target.

Workflow Associations and Instances

Workflow associations and instances can be applied to targets running Express edition, however the Express edition farm cannot replicate these actions back to an Enterprise edition farm. If the workflow is processed through replicated SharePoint objects, for example, tasks in an Approval workflow, then you can successfully process a workflow on the Express edition farm.

Alerts

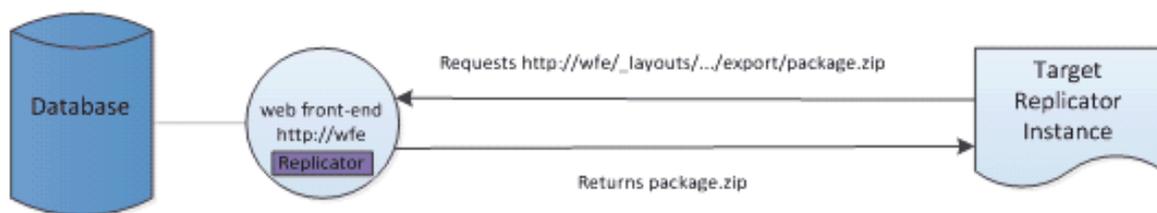
Alerts can be replicated from Standard edition farms to Express edition farms. Changes to alerts are not replicated if they occur on a farm with Replicator Express edition.

Web Front-ends

Web front ends process incoming requests on a web server. This section describes and depicts the various web front-end scenarios that are often used by companies and the way in which Replication can be set up to function with each.

Single Web Front-end

If your farm only has one web front-end server, then when the target Replicator instance requests a replication package, it makes the request from that web front-end server. IIS maps the path to the package file, then to the Replicator data folder on the disk, and then returns that file.

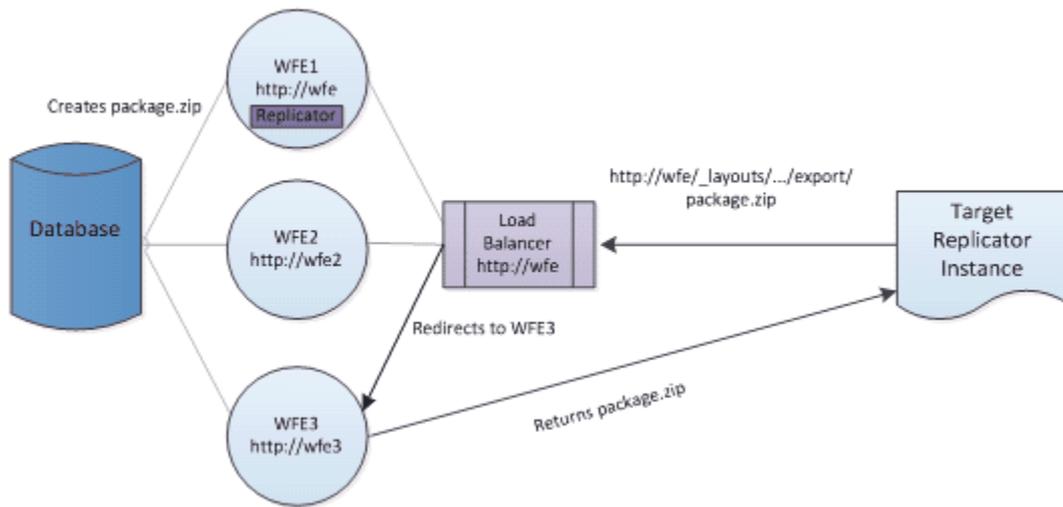


Multiple Web Front-ends

When dealing with multiple web front-ends, Replicator only needs to be installed on one web front-end. In this case Replicator captures changes, extracts from the content database, and creates the package zipping it to the Replicator data folder on a shared network drive.

Scenario 1

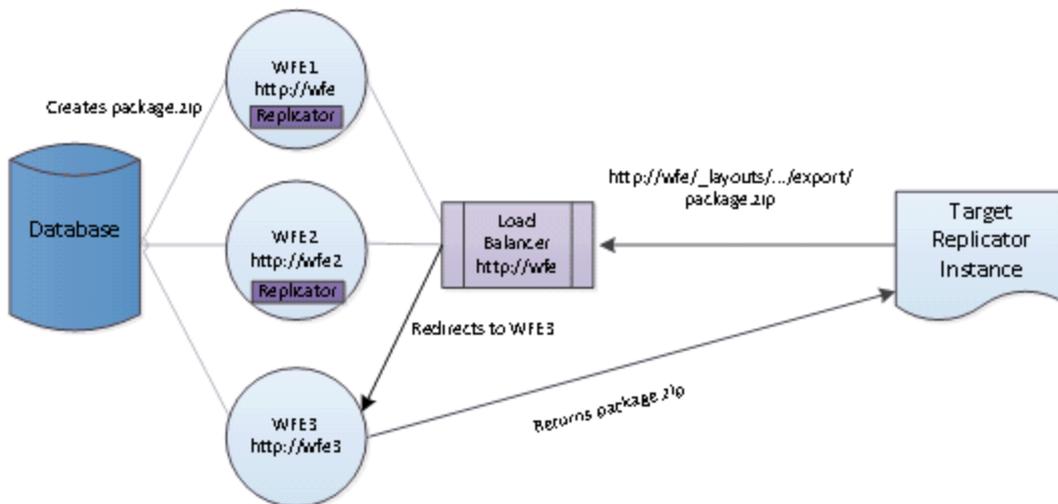
If your web server has multiple web front-ends, a load balancer is often used in order to redirect traffic to a specific web front-end. Since replicator does not know which web front-end will be sending the package to the target, it is important that all web front-ends have IIS configured to map the Replicator data folder to a shared network folder.



In the picture above, the load balancer is configured to receive incoming requests at <http://wfe> and forward them to any of the three available web front-ends. In the diagram, the request has been sent to WFE3, at <http://wfe3>. From this point the process works in the same way as the single web front-end example. IIS on WFE3 is configured to use the Replicator data folder for the export path and the package file is sent to the Target Replicator Instance that requested the package.

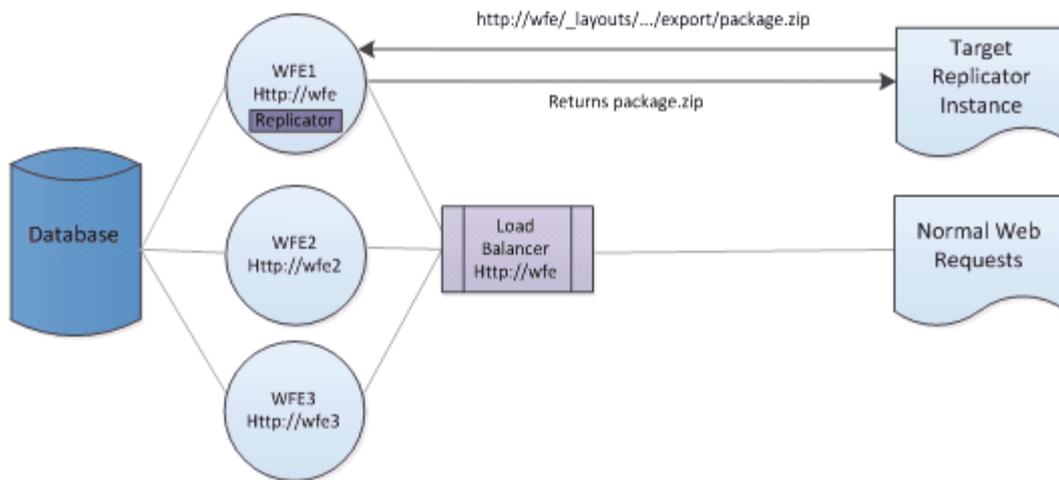
Scenario 2

If you wish to install additional Replicator engines to your other WFEs, like WFE2 at <http://wfe2>, then this can increase your thread count, and the ability of Replicator to manage multiple actions simultaneously. For more information on threads and parallel processing, see the sections below titled [Parallel Processing](#) and [Threads](#).



Scenario 3

If you are not using a load balancer, or do not want to set up your replication system to work through a load balancer, then you can set up replication to work with a single specific web front-end at all times. In this case, our target Replicator is configured to communicate directly to the server where Replicator is installed, WFE1 in our example. WFE1 is the only front-end that will service Replicator requests, meaning we do not need or use the IIS virtual directories on WFE2 or WFE3. A load balancer can still be used to redirect normal web traffic coming to <http://wfe>.



The disadvantage of such a configuration is that you limit your system to a single WFE for all Replicator requests. This can add additional strain onto that single WFE and could therefore affect the performance of replication, or clients accessing that server.

Multiple Replicator Engines

With Metalogix Replicator Enterprise Edition, you can install Replicator on multiple front-end servers in the same farm. The benefit of doing so is:

- Instead of one server managing the load of creating and applying events, the load can be distributed over multiple web front-end servers.
- In the event that the front-end server where the Replicator services are running becomes unavailable because of an outage, then Replicator will continue to create and apply packages using the Replicator services on the other web front-end servers.

i **NOTE:** Replicator is designed to continue capturing events as they occur, even if the front-end server running the services is not available. These events will merely be captured and placed in a queue until the server is back up and able to replicate them.

Parallel Processing

Parallel processing involves the use of multiple threads in order to perform various computing actions simultaneously. The number of threads available determines the number of actions that can be run at the same time.

Parallel processing involves the handing of tasks simultaneously. However, it is important to note that some tasks must occur in sequence. For example, within a site we process all events in the order that they are captured, on the outbound, and the order that they are received, on the inbound. One thread is therefore used within a site to perform this processing. When a list is added, it must be processed prior to the addition of an item to that list, so that there is a list present on the target to add the item to. This means that a thread will be used to process the addition of the list, and once this processing is completed, the same thread will be able to process the addition of an item to that list. These two cannot be performed simultaneously, and therefore do not require the use of more than one thread, the same thread.

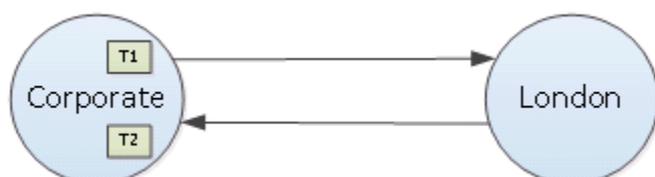
Threads

As mentioned previously, the number of threads available, determines the number of actions that can be run at the same time.

Vertical Scaling of your environment allows you to have multiple threads on a server. Having multiple threads is convenient for simultaneously performing multiple tasks.

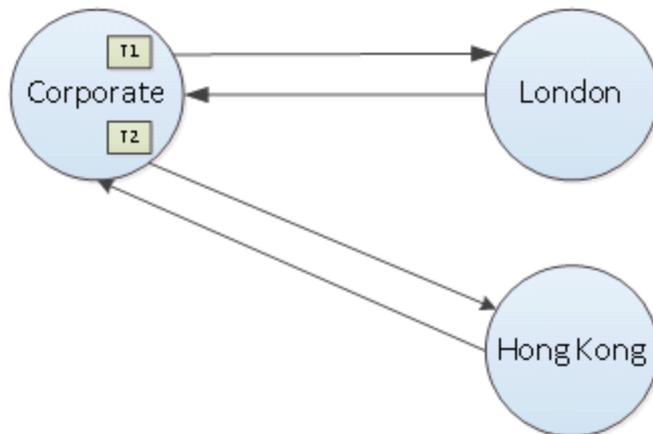
This is particularly useful when dealing with multiple site collections or sites. With multiple site collections or sites, you can have a different thread performing the processing for each site.

Replication functions optimally if it is capable of handling inbound and outbound processing on a web application simultaneously. For instance, in the following environment, an outbound event leaving Corporate has one thread available for processing, labeled T1, and an inbound event coming into Corporate from London has one thread available for processing, labeled T2. In total two threads are available for real-time bidirectional replication to take place. This environment would function with one available thread as well; however, processing would take longer as the same thread would be required to process outbound and inbound events on the web application.



If we add another farm into the mix, then replication may be slower with just two threads. In such a scenario Replicator will wait for all outbound processing to finish before inbound packages can be applied. This in turn slows down replication.

In the following scenario two threads may cause for slower processing, as an outbound event would require one thread to handle sending the package to London from Corporate, labeled T1, and another thread to send the package to Hong Kong from Corporate, labeled T2. This means that the threads are not available to processing potential incoming packages from London or Hong Kong, until outbound processing as completed. This in turn hinders the pace of replication, causing it to be slower.



In Replicator, the number of threads available for parallel processing is determined by the edition of Replicator. Express edition has 1 available thread, Standard has 2, and Enterprise has a default of 10 threads running on each replication engine, with the ability to increase to an unlimited amount of threads.

Factors Affecting Threading

Consider the following factors found within your SharePoint environment that will use threading:

- Outbound or inbound processing.
- Map families.
- Connections.
- Site collections.
- Classification queue numbers.

Each of these factors can be processed independently to distribute the load across multiple threads and will, therefore, determine the edition of Replicator that will best suit your environment. They also, therefore, determine how many threads you can take advantage of at a time.

Enterprise Edition

Enterprise Edition has a default of 10 threads working at the same time. This number can be increased to an unlimited number of threads operating at the same time. This is particularly convenient in a hub-and-spoke architecture, where Enterprise edition is set up at the hub. With Enterprise edition you will be able to handling incoming and outgoing processes for multiple spokes, allowing Replicator to function in real-time, even in more strenuous scenarios.

It is, however, important to keep in mind that we are limited by the number of threads that you can take advantage of, as replication works in sequence order.

Enterprise edition is also the only edition which allows horizontal scaling. With this edition you can use multiple engines, thereby allowing you to disperse the load placed on any one machine by dividing the load of processing between them.

Advantages for having multiple replication engines in a farm include:

1. In the event that one of your servers stop, you have another which will insure the continuity of replication operations. This ensures high availability on your farms.
2. You spread the load of processing packages between multiple replication engines, allowing your machines to work faster due to parallel processing.

For example, you can shape your farm to have two Enterprise Replication engines running in parallel. You can specify that each engine handle a maximum of five threads at a time, thereby evenly distributing the load placed on each engine. In turn, both engines will be able to work faster, optimizing both your system as well as the functionality of Replicator.

i | **NOTE:** that horizontal scaling is entirely transparent to the user, and therefore the user has no control of it.

Metalogix Replicator Specifications

This section summarizes the capabilities of Metalogix Replicator.

Editions and Features

Built from the ground up and based on Microsoft technologies, Replicator has matured to fully support mission critical applications. Our customers use our products in a wide variety of business applications including disaster recovery, mobile platform data replication, and fully distributed fringe data networks. The use of Replicator varies as widely as the variety of SharePoint deployments. To meet the needs of this diverse customer base, Replicator is available in several editions, with additional support for replicating SharePoint look and feel features.

Enterprise Edition

The Enterprise edition provides a complete feature set that is applicable to customers with mission critical, high volume replication requirements, including support for multiple processing servers, transaction logging, and programming APIs. This edition is ideal for WAN replication. It provides a heavy-duty version of Replicator which allows for scenarios like ship-to-shore Maritime replication, and Military Tactical SharePoint replication. It works ideally for global replication, allowing for a global mesh topology setup.

Cross Domain Edition

The Cross domain edition provides all of the features available in the Enterprise edition of Replication, with the addition of SFTP settings. The SFTP settings provide support for Cross Security Domain replication, which required communication with an intermediary server that does package inspection. This new transport method provide a secure method of transporting replicated content which is FIPS compliant.

Standard Edition

The Standard edition provides customers with an ideal solution for geo-replication of content, and the perfect solution for an organization that has a global user base with a need for instant global availability of SharePoint content. Standard edition will enable you to do this with a simple hub-and-spoke replication environment setup.

Standard edition also includes support for the replication of look and feel settings and changes between the same major editions of SharePoint. For example, you can replicate list views from

SharePoint 2010 to SharePoint 2010 and from SharePoint 2013 to SharePoint 2013. Standard edition also supports the live processing of workflows on multiple farms.

Express Edition

The Express edition of Replicator is the ideal solution for small-scale replication environments. This includes intranet to extranet SharePoint environments, departmental SharePoint environments, and basic farm to farm continuous migration scenarios.

Replication Feature Matrix

The following table details the features that are enabled in each edition of Metalogix Replicator.

Category	Feature	Express	Standard	Enterprise	Cross-Domain
Replication Scope	Full Web Application Replication	✓	✓	✓	✓
	Site Collection Replication	✓	✓	✓	✓
	Web Site Replication	✓	✓	✓	✓
	Out-Of-The-Box Web Parts	✓	✓	✓	✓
	Custom and third party Web Parts		✓	✓	✓
	Look & Feel Replication (Web Part Pages, Views, Themes, Branding, Navigation)	✓	✓	✓	✓
	Workflow Association	✓	✓	✓	✓
	In-process Workflow Replication		✓	✓	✓

Category	Feature	Express	Standard	Enterprise	Cross-Domain
	Out-Of-The-Box Workflows	✓	✓	✓	✓
	SharePoint Designer Workflows	✓	✓	✓	✓
	Visual Studio Workflows	✓	✓	✓	✓
	Nintex Workflows	✓	✓	✓	✓
	Offline Replication Export		✓	✓	✓
	Offline Replication Import	✓	✓	✓	✓
Replication Control	Cross-Domain Replication	✓	✓	✓	✓
	SFTP Transport Method				✓
	User Remapping		✓	✓	✓
	List Remapping		✓	✓	✓
	Web Site Relocation		✓	✓	✓
	Priority Queues			✓	✓
	Replicate based on	✓	✓	✓	✓

Category	Feature	Express	Standard	Enterprise	Cross-Domain
	Metadata (Rules)				
	Rules based redirection			✓	✓
	Automatic Event Capture	✓	✓	✓	✓
	Immediate Replication of all changes		✓	✓	✓
	Scheduled Replication	Every 5 min or Longer	✓	✓	✓
	Manual Replication	✓	✓	✓	✓
Replication Architecture	Maximum Outbound Connections per Web Application	2	Unlimited	Unlimited	Unlimited
	Multiple WFE Support per Farm	✓	✓	✓	✓
	Replication Engines per Farm (max. for scalability and failover)	1	1	Unlimited	Unlimited
	Replication Threads per Engine (max.)	1	4	Unlimited	Unlimited
	Web Service API & Object Model Support			✓	✓

Category	Feature	Express	Standard	Enterprise	Cross-Domain
Monitoring and Tooling	Replicator Central Administration	✓	✓	✓	✓
	PowerShell Cmdlet Support	✓	✓	✓	✓
	Event Processing History Log			✓	✓
	User Alerts		✓	✓	✓

*Express and Cross Domain Edition support only the Schedule-ReplicatorMap command for Powershell.

Replication Support Between SharePoint Versions

The following table depicts the general support for replication between different SharePoint versions.

Source	Target	Result
2007	2007	Supported.
	2010	See sections below.
	2013	See sections below.
	2016	See sections below.
2010	2007	See sections below.
	2010	Supported.
	2013	See sections below.
	2016	See sections below.
2013	2007	See sections below.
	2010	See sections below.
	2013	Supported.
	2016	See sections below.
2016	2007	See sections below.
	2010	See sections below.
	2013	See sections below.
	2016	Supported.

Look and Feel Support Between SharePoint Versions

The following table depicts the combinations of SharePoint editions that are supported by Replicator for look and feel. Look and feel events include changes to themes, site navigation, views, and pages.

Source	Target	Result
2007	2007	Supported.
	2010	Limited support. Target site has SharePoint 2010 appearance.
	2013	Supported, except themes and master page list items.
	2016	Supported, except themes and master page list items.
2010 (from 2007 upgrade, visual upgrade not performed)	2007	Not Supported.
	2010 (from 2007 upgrade, visual upgrade not performed)	Limited support. Target site has 2007 appearance.
	2010	Not Supported.
	2013	Not Supported.
	2016	Not Supported.
2010	2007	Not Supported.
	2010	Supported.
	2013	Supported except themes and master page list items.
	2016	Supported except themes and master page list items.

Source	Target	Result
2013	2007	Not Supported.
	2010	Not Supported.
	2013	See sections below.
	2016	See sections below.
2016	2007	Not Supported.
	2010	Not Supported.
	2013	Not Supported.
	2016	See sections below.

SharePoint 2013 Look and Feel Support

The following is a list of the supported and the not supported look and feel features that are replicated.

Supported

Event Group	Events
Legacy Look and Feel Features	<ul style="list-style-type: none"> • Master Page • Title, Description, Logo • Welcome Page • Tree View/Quick Launch control • Navigation
Change the Look	<ul style="list-style-type: none"> • Change to out of the box looks • Change current look
Design Manager (Publishing Feature)	<ul style="list-style-type: none"> • Create Design Package • Edit Display Templates • Edit Page Layout

Not Supported

Event Group	Events
Composed Looks Gallery	<ul style="list-style-type: none"> • Change out of the box looks • Create New Look

Event Group	Events
Publishing Look and Feel Features (Requires Publishing features to be enabled)	<ul style="list-style-type: none"> • Page Layouts and Site Templates • Device Channels
Image Renditions (Publishing Feature requires BLOB Cache to be enabled)	<ul style="list-style-type: none"> • All Supported

Workflow Support Between SharePoint Versions

The following table depicts the combinations of SharePoint editions that are supported by Replicator for workflow events. Workflow support includes out of the box, SharePoint Designer, Visual Studio, and Nintex workflows.

Source	Target	Result
2007	2007	Supported.
	2010	Not Supported.
	2013	Not Supported.
	2016	Not Supported.
2010	2007	Not Supported.
	2010	Supported.
	2013 running on Windows Workflow Foundation 3 (built-in)	Supported.
	2013 running on Windows Workflow Foundation 4 (installed with Workflow Manager)	Not supported.
	2016	Not Supported.
2013	2007	Not Supported.
	2010	Not Supported.
	2013 running on Windows Workflow Foundation 3 (built-in)	Supported.

Source	Target	Result
	2013 running on Windows Workflow Foundation 4 (installed with Workflow Manager)	Not Supported.
	2016	Not Supported.
2016	2007	Not Supported.
	2010	Not Supported.
	2013 running on Windows Workflow Foundation 3 (built-in)	Not Supported.
	2013 running on Windows Workflow Foundation 4 (installed with Workflow Manager)	Not Supported.
	2016	Supported.

Supported Web Parts

Metalogix Replicator replicates the following web parts between the same versions of SharePoint. For example, the Microsoft Picture Library Slideshow can be replicated from a SharePoint 2010 farm to another SharePoint 2010 farm.

Vendor	Partial List of Supported Web Parts
Microsoft	<ul style="list-style-type: none"> • Advanced Search Box • Colleagues Web Part • Content Editor • InfoPath Form • Site Aggregator • User Tasks
Agreeya QuickApps 6.6 (SharePoint 2013 support only)	<ul style="list-style-type: none"> • qCalendarView • qCaptionDisplay • qCascadingMenu • qChartView • qDiscussionView • qExcelViewer • qItemDisplay • qListForm • qListView • qManagement • qMediaView • qPageRedirector • qPanelMenu • qSIChartView • qSIListForm

Vendor	Partial List of Supported Web Parts
	<ul style="list-style-type: none">• qSIListView• qDynamicLayout• qHelpLink• qMultiSelector• qSelector• qSISelector

A complete list of the web parts supported by Replicator is available by running the PowerShell `PS>Write-ReplicatorSupportedWebParts` command. For more information on this command see the [Metalogix Replicator Command Line Reference Guide](#).

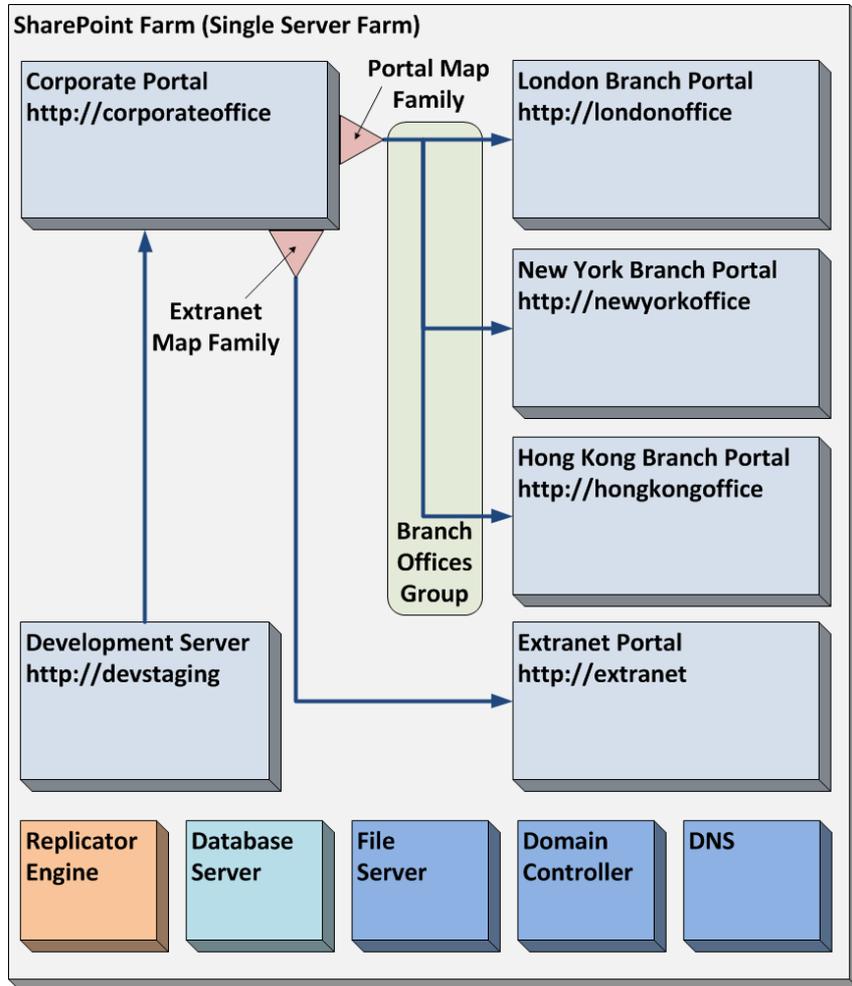
Metalogix Replicator Reference Architecture

The following diagrams illustrate the Metalogix Replicator Reference Architecture used to explain how Replicator operates as well as provide a concrete example that is referenced in the Metalogix documentation and training courses.

i **NOTE:** The diagrams below are used for reference only. They are not intended to represent best practices from a SharePoint deployment and configuration perspective. For example, the DNS, Active Directory and SQL Server, servers should be installed on their own - both physical and virtual servers. The File Server, used for Replicator's export, import and cache folders, should use high reliability storage (RAID 1 mirrored drives or higher).

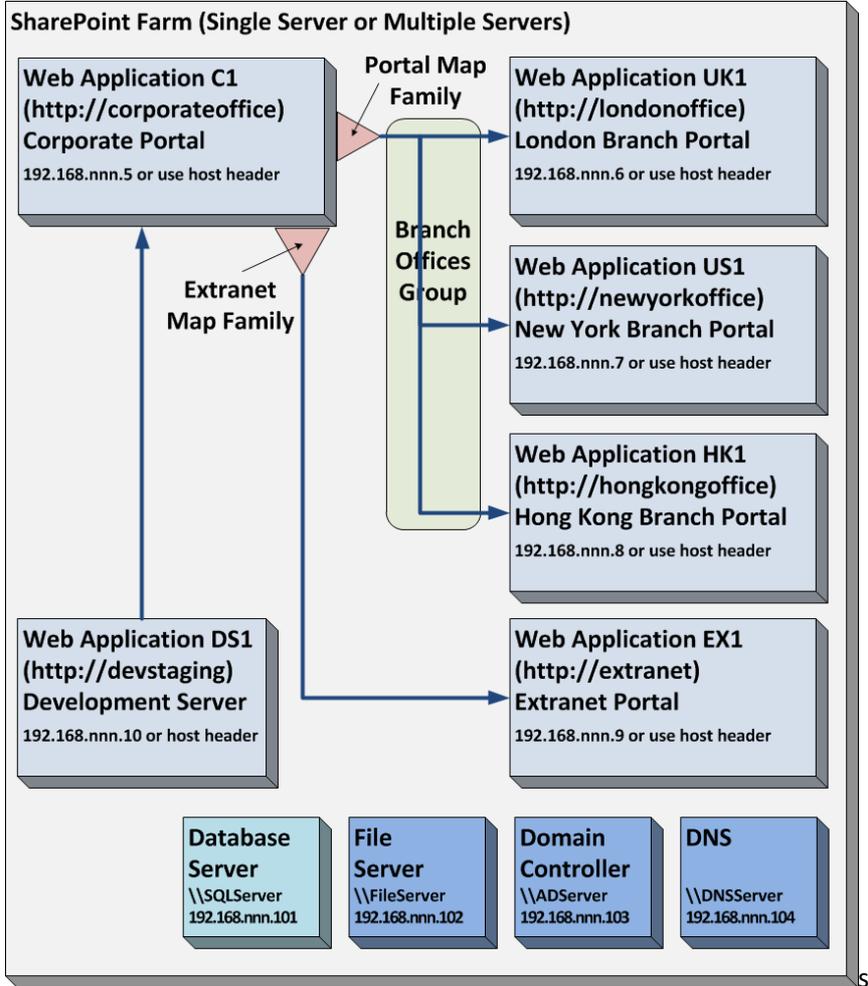
Metalogix Replicator Logical System Architecture

The following diagram depicts the Metalogix Replicator Logical System Architecture.



Metalogix Replicator Physical System Architecture

The following physical system architecture is useful for configuring evaluation lab and training course environments.



About

We are more than just a name. We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece — you — to the community, to the new Quest.

Contact Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical Support Resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions

- Chat with support engineers online
- View services to assist you with your product