Metalogix® Replicator 7.4

September 2023

Disaster Recovery Best Practices

This document describes how Replicator responds to various Continuity of Operations Plan (COOP) and Disaster Recovery (DR) scenarios. Although Replicator can be configured similarly for both scenarios, we recommend a separate configuration between your primary farm and disaster recovery farm. The disaster recovery-specific configuration is described later in this document.

Continuity of Operations (COOP) can be described through a variety of scenarios, but essentially it ensures that users have continuous access to up-to-date content. Replicator helps ensure the continuity of operations for organizations by providing up-to-date content to all end-users, regardless of geographic location. This is where geo-replication and COOP cross paths in terms of their purpose.

Disaster Recovery can be described as a tool, or safeguard, that ensures that the continuity of operations is maintained regardless of possible scheduled or unscheduled outages.

A scheduled outage can be anything that requires a planned data-center outage for an upgrade, a planned take down of a web front-end or an entire SharePoint farm.

An unscheduled outage includes any kind of fail-over scenario in which administration is not given the option to prepare or manually stop replication this includes power outages, network issues, and natural disasters.

In the event of either a planned and unplanned outage, a Disaster Recovery environment is key to have in place, as it ensures that COOP is maintained. It is important to note that our recommendation is to set up a Disaster Recovery environment using one-way replication. A DR environment is strictly for the purpose of backing up and recovering all content in the event of a planned or unplanned outage.

Replication is initially set up using a one-way connection that replicates all content regularly from your main server to your DR environment. In the event of a planned or unplanned outage on this main server, pause the one-way connection from your main server to the DR server, and direct all of your users to your DR environment. Here end-users can continue to use SharePoint as per usual, where all changes will be replicated and placed in a queue. Once your main server is back up and running, set up a connection from your DR to your main server and then pushing the event queue with all changes back over.

Redirecting Users to your DR environment

There are various ways in which this can be done; the ultimate method of redirection is up to you.

The simplest solution in this scenario is to have your IT team send out a mass email, notifying all end-users that your main server is down, and that the backup or Disaster Recovery server should be accessed for all SharePoint use until further notice.

More advanced solutions include having a load-balancer set up between your main server farm, and your Disaster Recovery farm. The load balancer can detect that the primary site is unavailable and then automatically redirect all inbound users to the DR site.

The following sections describe what happens with replication when a front-end server becomes unavailable. Any server where Replicator is installed is called a replication engine.

Single Web Front-end is Not Available

Farms with a single web front-end server are not able to serve client requests from end users or Replicator. Other Replicator instances that are still running will continue to try to notify this farm that replication packages are available for download or will try to download replication packages from this farm until it becomes available.

Front-end Server is Not Available

If a WFE not running Replicator services fails, then there is no specific Replicator configuration required since a WFE running Replication will continue to process all inbound and outbound packages.

Replication Engine is Not Available

In this scenario, users can continue to access SharePoint, but Replicator is unable to process changes made on this or other farms. Changes made to this farm are still collected in the Replicator databases and when the Replication Engine is restored, replication continues.

Replication Engine is Not Available with Multiple Replicator Servers

If you have multiple WFEs running as replication engines, then if one of these goes down, the other replication engine continues to process inbound and outbound packages.

There is more than one way to configure Metalogix Replicator for disaster recovery scenarios. This document describes the basic procedures to prepare one of these configurations. Specifically, we will create a separate and isolated web application for disaster recover (DR), configure one-way replication to the DR web application, and then synchronize it with our primary SharePoint web application. In this example http://corporate is the main web application, and http://corporate-dr is the disaster recovery web application to which the user will be redirected in the event of a disaster.

Initial Configuration

1. Install Metalogix Replicator on both the main and DR servers.
2. Configure Replicator to perform one way replication from http://corporate to http://corporate-dr.
    - With one way replication, only changes that occur on the main site will be replicated over to the DR site. Any changes made to the DR site will not be replicated over to the main site.
3. Enable replication of SharePoint user alerts using Replicator enable alerts PowerShell administration command.
    - Replicator does not replicate user alerts by default; this feature must be enabled for DR scenarios.
4. Turn off outgoing emails on the DR SharePoint farm by clearing the Outbound SMTP server field in the DR farm's SharePoint Outgoing E-Mail settings.
    - Clearing the SMTP server field for the Outgoing E-Mail settings will disable the DR SharePoint farm from sending outgoing emails triggered by alerts thus avoiding duplicate notifications.
5. Perform initial site collection synchronization using the Replicator site collection backup PowerShell command.
    - The Replicator site collection backup PowerShell command uses the SharePoint site collection backup feature to create a backup which is then packaged up and replicated to the target where the backup is then restored so that the content on the main site and DR sites are now the same. This does not, however, allow for domain remapping, user remapping,
6. Now Replicator will catch the changes made on the main SharePoint site and apply them to the DR site as they occur.

On-going Maintenance

- Run Replication Reports regularly. This will insure that you have validated and synchronized the content that is available on the DR site (http://corporate-dr).
- Perform site collection backups regularly using the Replicator site collection backup PowerShell command.

Comprehensive synchronization on the DR Site

You can do comprehensive synchronization periodically, weekly or monthly based on your needs. This ensures that the disaster recovery site has all content, structure, and permissions from your main site, failover to the Disaster recovery site.

1. When an outage occurs, whether planned or unplanned, redirect users to the DR site (http://corporate-dr).
    - Users can continue operating seamlessly, as they will have access to all of the same content that was available on the main site.
2. Turn on outgoing emails from your DR farm.
    - Do so by filling out the SMTP server field for your Outgoing email settings on your DR SharePoint farm. This will allow you to get emails triggered by alerts on your DR farm.

3. Reconfigure Replicator to perform one way replication from DR site to main site with outbound processing paused.
    - Replicator on the DR site will now capture changes to the now live DR site, but will pause replication, queuing up the changes for when the main site is ready to come back online.

Recovery

1. When the main server is ready to come back on line, reconfigure replicator on the main server to accept one way replication from the DR site. And un-pause Replicator on the DR server.
    - This will allow the changes that were captured on the DR site, during the time period that users were redirected to the DR site, to be queued over to the main site.
2. Alternatively, perform a site collection backup from the DR site to the main site using the Replicator site collection backup PowerShell command.
    - This will re-synchronize the main site with any changes made during the redirection time period.
3. Once all content from the DR site is updated on the main site, reconfigure Replicator back to initial one way replication from main to DR.
    - This will bring your environment back to its original state, where users can continue working on the main site.
4. Turn off outgoing emails from the DR farm.
    - Do so by clearing the SMTP server field in the Outgoing mail settings for the DR site. As stated previously, this will prevent notification emails from being sent out about alerts, creating duplicate emails.
5. Redirect users back to the main site (http://corporate).
    - Users can now continue using the main site, and all the changes have been synchronized, allowing for flawless continuity of operations.