Why use a preload (seed) drive?

If you have a large amount of data in an AppAssure CORE repository that you will be replicating to eFolder's cloud, you may want to request a preload (seed) drive from us.

To help determine if this procedure may be necessary for you, please see this knowledge base article on estimating how long the initial backup job will take when sending data over the Internet.

We recommend doing a preload any time you are backing up more than 100GB of data on a standard internet connection.

eFolder offers a round-trip preloading (seeding) service, which includes everything required to properly preload (seed) your account. Depending on the shipping option you choose, it could take a few days to receive your drive, so we recommend starting the process as soon as possible.

- Please review How to Request a Preload Drive to get started.

- If you are ready to request your preload drive, please click Request a Preload (Seed) Drive.

- More details about the preload process are available in our FAQ.

Not using AppAssure? To create a preload (seed) drive for another software package, please refer to:

- How To Create an eFolder BDR for Acronis Preload (Seed) Drive
- How To Create an eFolder Backup for Files Preload (Seed) Drive
- How To Create an eFolder BDR for Replibit Preload (Seed) Drive
- How To Create an eFolder BDR for ShadowProtect Preload (Seed) Drive
- How To Create an eFolder BDR for Veeam Preload (Seed) Drive

Additional Assistance

eFolder is committed to responsive, competent technical support. Our team strives to exceed your expectations. Contact eFolder by emailing support@efolder.net or call us at 800-352-0248 [Option 2]. Additional material is available online at http://www.efolder.net/support/ and in the eFolder Partner Portal.

The People Behind Your Cloud

Overview

This article provides instructions for preloading data to a hard drive for both AppAssure V5 and Replay V4.

AppAssure V5 Instructions to preload data to a hard drive

Important: Please note that you can only replicate or seed AppAssure recovery points to eFolder that have been created with agent encryption. **We will reject a seed that contains unencrypted AppAssure V5 recovery points.** For instructions on how to encrypt your AppAssure V5 agents, please refer to our AppAssure Cloud Replication Guide.

1. Use the AppAssure V5 Admin Console to configure replication using your assigned target CORE hostname and credentials. Make sure that you check **Use a seed drive to perform initial transfer.**
2. On the *Copy To Seed Drive* dialog that appears when setting up replication, specify the path to the disk that will store the seed data and start the operation.
3. The export operation will progress in the background. You can use the **Events** tab to monitor the details of the export operation to monitor it for completion.

Important: Compression and deduplication only occur within AppAssure when recovery points are committed to the repository. When AppAssure exports recovery points, it puts them a raw state, meaning the seed data is uncompressed and may be larger than what is stored in your local AppAssure repository. **Pay close attention** if AppAssure prompts for an additional drive before completion of the export.

VERY IMPORTANT: There must be as much free space in the Target Core repository as the total amount of data exported to the seed disk. Otherwise, the seed will *not* have enough space to import.

Notes:

- Replication of the incremental recovery points that are *not* part of the seed drive data will automatically occur over the network. This ensures that replication will not fall behind while you are waiting for the seed drive to be processed. This is one of the unique advantages of the AppAssure V5 technology that makes it well suited for large data sets.

- In AppAssure V5, you do **not** need to disable and re-enable rollup during the seed operation. This is only required for Replay V4 seeds.

Replay V4 Instructions to Preload Data to a Hard Drive

1. Use the Replay Admin Console to configure replication for each protected server on the CORE, using your assigned target CORE hostname and credentials. Make sure that you CHECK the **Pause replication on the protected server** checkbox when setting up replication.

2. Connect to the target CORE using the Replay Admin Console 4.7.1.
   Do this by using the **File** menu in the Admin Console, connecting to CORE, entering the remote hostname and credentials, and adjusting both the rollup (data retention) and virtual standby settings.

   > **Note:** You must **disable rollup** throughout the seed process (more details are given below).
   > For virtual standby, update the virtual standby image **weekly** or **daily** (not hourly).

3. EXTREMELY IMPORTANT: Use the Replay Admin Console to disable rollup for **all** agents having their data copied. You must do this on both the source CORE and target CORE using the Replay Admin Console.
a. Login to the target CORE by using the **File** menu and connecting to a Replay CORE command in the Admin Console. To disable rollup, right-click each server, choose **Properties**, go to the **Retention** tab, and uncheck **Rollup**. Rollup must remain disabled until the data has been fully consumed in the remote data center and Replay indicates that both sides are fully synchronized.
   Not properly following this step can cause you to have to redo your seed again. **Please be sure to disable rollup**. If you have any questions about the process, please submit a ticket.

4. Use the *Replay Admin Console* to copy the initial seed data to the disk. In the CORE status pane, click on the **Replication** tab to perform a **Copy** operation and select one or more servers. Click **Copy** to begin.

   VERY IMPORTANT: Put a README.txt file on the root of the drive that includes your assigned target CORE hostname (such as aa4-myhostname.aa.sc...). Without this file, we will *not* be able to determine which target CORE server should use the data. Once this has been done and the disk shippeded to us, you will be notified when data has been copied to your target CORE and that you now need to perform the following steps to finalize the procedure.

5. Use the Replay Admin Console to connect to the source CORE and resume replication.
   Do *not* enable rollup yet.

6. At this point, Replay may take a significant amount of time to commit the recovery points from the replication staging area to the repository. The length of time depends on the amount of data, as it is verifying the checksums of the data to ensure no data was corrupted while on drive. This can take several days or longer for large amounts of data (hundreds of GBs or TBs). Once this has occurred, Replay will then upload recovery points from the source to the target to fully catch up. Once replication is fully caught up, the Replay Admin Console will show a synchronized status and both sides will be in sync.

7. EXTREMELY IMPORTANT: Do not re-enable rollup on the source CORE and target CORE until replication for each agent has entered the synchronized status. Once the Admin Console shows that replication is in a synchronized status, you must log in to *both* the source CORE and the target CORE and enable rollup for each server that you seeded. To enable rollup settings, right-click the server, then select **Properties** and click on the **Retention** tab.)