

Mailbox and Journal Repositories Node

<https://campus.barracuda.com/doc/43222679/>

A repository is an archive of one or more mailboxes, which consists of a set of data files stored either on a [storage manager](#) or on hard drive, and a set of index files which must be stored on a hard drive. Multiple repositories are independent and can have different configurations. For example, they can have different storage and different retention periods. A repository can be removed without affecting other repositories.

A repository typically holds multiple mailbox archives. Typically there are only a few repositories that hold all the mailbox data for all mailboxes.

If you have run [ArchiveOne Admin](#), you have already set up a repository. This may be sufficient for your needs; create additional repositories if you want to configure the mailbox archive in multiple types of storage or configure them differently.

There are two types of repositories, Mailbox and Journal, listed on the corresponding nodes. Mailbox repositories only store data archived from standard mailboxes and Journal repositories only from journal mailboxes.

If you have multiple ArchiveOne installations on multiple servers, you can add a reference in the current **Mailbox Repositories** node to a repository controlled elsewhere. That is, the repository can be archived to by this product but its configuration is held and managed by another ArchiveOne Admin. This is referred to as a *linked repository*; the original is referred to as the *link parent* and the linked one is referred to as the *link child*.

Once you have linked to a repository, you can use the link child in exactly the same way as any repository for archiving. Note that there are some differences in behavior:

- Attempting to delete a link child only unlinks it;
- Linked repositories cannot be renamed;
- Restoring an index for a repository can only be done on the link parent.

You cannot link to a repository that is using a [storage manager](#) unless the link parent and link child are using ArchiveOne installed on the same server.

To create a new repository:

1. In ArchiveOne Admin, right-click **Mailbox Repositories**, and then click **New Mailbox Repository**.
2. The Wizard launches.
3. Follow the onscreen prompts to enter the repository characteristics.
4. Click **Finish** to create the repository.

5. Click **Mailbox Manager**, right-click the mailbox that is to use the new repository as its archive, and click **Properties**.
6. In the **Mailbox Properties** dialog box, select the new repository from the **Repository** drop-down menu, and then click **OK**.
For additional information, refer to [Group Mailboxes into Repositories in Step 1](#).
7. To link to an existing repository, right-click **Mailbox Repositories**, and then click **Add Linked Repositories**.
8. The Wizard launches.
9. Follow the onscreen prompts to select the **Manager** and the **Configuration server** the manager is processing. Click **Next**.
10. Accessible repositories display in the list. Select the repositories to which you want to link, and then click **Done**.

Note that your logged in user and the user running the ArchiveOne Service (typically called *ArchiveOneAdmin*) must have full read/write access to the add-ins share on the link parent server, the manager-specific subdirectory beneath it, and all files and directories beneath that.

To manage Journal repositories, run the [Journal Archiving Wizard](#).

Repository configuration includes the following; note that some features may not display based on your installation and configuration:

- [Storage manager](#) - Specify the storage manager the repository is to use.
- [Index Location](#) and [Archive Location](#) - Specify the Index Base Directory and Archive Base Directory.
- [Retention period](#) - If one is specified, archived mail created more than the specified time period is marked as expired, will no longer appear in standard search results, and the message links for the mail are deleted. Note that these messages remain in the archive and an administrative search for expired mail can find and retrieve these messages.
- [Rate control](#) - Allow you to control the rate of data flow to avoid overloading your Exchange server.
- [Index subdirectory name](#) and [Archive subdirectory name](#) - Initially set to **1**, however, change these values if you want to extend the index or archive onto a new hard drive mounted as a different subdirectory of the Index Base Directory or Archive Base Directory of the repository.

Each mailbox can only be archived to one repository. If you have archived a mailbox to one repository and now want to change it to another, you cannot move the mail already archived. However, you can change the archive location for new mail archived from the mailbox.

If a repository is using a storage manager and is not a link child, right-click on the repository and click **Restore index** option; note that this option is disabled if you are not using a storage manager. For more information, refer to [Managing Long-Term Configuration](#).

Multiple repositories do not affect [search permissions](#). If an NT user has rights to search into the archive of a mailbox, they can do so regardless of which repository the mail has been archived to, even if the mail is spread over multiple repositories. When defining an administrative search, you can specify which repository to search which improves search time.

Items queued for archiving into a repository are processed on one of the following events:

- An archive policy which has written data to the repository has completed;
- At the start of each hour;
- When you right-click a repository, and then click **Start Archiving Now**.

There are several features which are specific to repositories using [Amazon S3 Disaster Recovery Storage Manager](#):

- The right-click context menu includes **Convert to local disk-based** to convert the repository from Disaster Recovery to local.
- It right-click context menu includes **Verify** which compare the S3 stored files to the local files and reports whether they match.
- The **Verify** process runs on a regular basis. Right-click on **Repositories**, and click **Set Repository Validation Schedule** to configure when this process is run.
- Any non-Storage Manager repository includes the **Convert to DR** option to convert it to an Amazon S3 Disaster Recovery repository.

Click [Role definition](#) in the toolbar to assign [How to Define Role Behavior](#).

You can view numeric and graphical representations of the history of repository size changes. To view the history of all updates to all repositories, right-click **Repositories**, and then click **History**. To view the history of all updates to a single repository, expand **Repositories**, right-click on the repository, and then click **History**.

If you are running multiple instances of the ArchiveOne Admin application, changes made on this node in one instance are not visible on the other instance. If you suspect there have been changes, restart your instance to ensure it has up to date data.

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.