
Index Optimization Service

<https://campus.barracuda.com/doc/43223043/>

The ArchiveOne Index Optimization Service defragments the index to improve archive search. This feature is installed on the ArchiveOne server. It can run at the same time as other activities without interrupting those activities, for example, archiving to a repository that is being optimized.

Process Repositories

Once you have an installation of the ArchiveOne Index Optimization Service, you can use it to process repositories hosted on multiple servers managed by ArchiveOne. Optimization is not by default enabled for all repositories.

1. From the **Start** menu, launch **ArchiveOne Index Optimizer Configuration**.
2. In the configuration window, click **Add** to add a repository.
3. Enter the configuration server and select the product that owns the repository, and then click **Find** to view the repositories for the selected server.
4. Select one or more, and then click **OK**.

The repositories listed on the main application window are those that are to be optimized. The first repository is examined, and a small optimization may be made, then the next is examined, and so on until all repositories in the list are complete. Upon completion, the first repository is examined and another small optimization takes place, and so on. In this way, the optimization level for each repository gradually increases until all reach the configured level of optimization.

View the Report

Click **Problem report** to view details of any problems that have occurred with any partition files in the repositories during indexing. If any problem occurs while a partition file is processed, it is excluded from optimization, then the rest of the index is optimized. This includes the case of any one message not being retrievable from the partition. Typically the optimization process returns a high degree of index optimization even if some partition files are excluded in this way.

Configure Parameters

There are three configurable parameters. click **Configure** to make adjustments:

- **Target index file size** - The index for a repository is made up of a large number of small index files, usually one per archiving run, and the act of optimizing involves merging these into a small number of large index files, for which searching is much faster. This configurable parameter specifies the index file target size. For example, if you specify 10Mb, the optimization process quickly merges all small index files into files of at least 10Mb, and there is a quick improvement in search times. If you then change this to 100Mb, it takes longer and has a smaller benefit, but ultimately would produce a better optimized index.
- **Restart delay** - The amount of time the service is idle before checking if optimization is needed. The process above describes how each repository is processed, one by one. This process is continuous, and continues until every repository has every index file larger than the **Target index file size**, at which point, there is no benefit in running optimization. The service then goes idle for the time indicated by the **Restart delay** field, and the repositories are checked again for archiving which can be optimized.
- **Directory for temporary data** - Indicates the location on the server that has enough space for temporary data. This may require up to ten times the size of the **Target index file size**.

The lower half of the application window shows the current status of the optimization process, with options to start it if it is idle and stop it if it is running and needs to be stopped.

Making a Repository Optimizable

In order for a repository to be optimizable, the following conditions apply:

- The repository must not use a [storage manager](#).
- The archive and index directories must be accessible by the ArchiveOne Index Optimization Service. This can happen in one of three ways, otherwise the repository is rejected:
 - The ArchiveOne Index Optimization Service is running on the same server as the service which owns the repository, for example, they share a **C:** drive.
 - The repository's paths are defined using UNC's, as in **\\server\share\index**
 - The server hosting the repository has administrative shares available, that is, the C: drive is accessible using **\\server\c\$**.

The optimization process makes changes to a repository's index files but not its archive (**.zip**) files. If it replaces any index files, those files are moved into a directory under the index directory including the date and time when optimization was complete. For example, **PreOptimize_120621_173037** where the first number indicates the date (2012, June, 21st) and the second indicates the time (17:30). All index files are retained in this way so that, if necessary, you can revert back to a pre-optimized index. Note that these directories may take a lot of space so you may want to back up the files to another drive and then remove the files.

If you are installing ArchiveOne Optimization Service into an environment where you have a number of repositories, consider adding them to the list of repositories to be optimized. If some are closed repositories which no longer receive archive data (such as repositories for previous years), then once

they have been fully optimized you can remove them from the list.

© Barracuda Networks Inc., 2020 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.