

NonStop AutoSYNC – Frequently Asked Questions

This FAQ is designed to address general product questions. To obtain more detailed information, please consult the NonStop AutoSYNC User's Guide, available in the [HPE NonStop Technical Library](#). For more specific technical questions not covered here, contact the HPE GNSC support organization. If you are registered as an HPE Passport user, you can also consult the [NonStop eServices Portal](#) to access the Knowledge database.

USAGE

Our organization uses RDF and Shadowbase to replicate our database, so why do I need AutoSYNC?

RDF and Shadowbase replicate audited data, but for complete disaster recovery, you also need to keep your unaudited data, application objects, configuration files, TAcl macros, OSS files and shell scripts in sync with your backup system to be ready for a takeover, should one be necessary.

You can also use AutoSYNC to replicate files from one environment to another on the same system or between systems, such as from development to QA or from QA to production.

Can I use AutoSYNC for a system migration?

Yes, you can. AutoSYNC can operate over a TCP/IP connection, which allows you to move your files to a system with the same Expand name.

PREREQUISITES

What hardware and software are required to use AutoSYNC?

AutoSYNC runs on all supported NonStop servers and versions of Guardian supported by HPE.

It uses standard utilities such as BACKUP and RESTORE. It needs no other special software product.

Is AutoSYNC privileged?

AutoSYNC is not privileged. However, in certain circumstances, the synchronization processes might need to be licensed to perform certain functions, for example, to preserve the license or progid attribute of an object file when it is moved. See paragraph "Special Cases: Licensed Synchronization Programs" in the "Installation" chapter of the NonStop AutoSYNC User's Guide for details.

Do I need to have an AutoSYNC license on source and destination systems?

Yes, you have to have AutoSYNC installed and licensed source and destination systems. Moreover, the versions of the AutoSYNC software have to be the same on all the systems.

SCOPE OF FUNCTIONALITY

Does AutoSYNC update files on the destination?

No. AutoSYNC replaces the entire file on the destination.

What types of files can be synchronized with AutoSYNC?

- Unaudited files and SQL/MP tables
- OSS files, directories and links
- Optionally, audited files and SQL tables - Note that audited files are better suited for record-by-record replication using RDF or Shadowbase
- Files that do not exist on the destination
- Files that are "older" on the destination
- Optionally, files that simply have different mod timestamps on source and destination
- Example: an object file that has been SQL compiled on the destination after it was moved will have a more recent mod timestamp than the source
- Files that are not "ZZ" files or swap files
- Files not opened for exclusive or update access in source
- Optionally, unaudited source files opened for update (non-exclusive open)

Does AutoSYNC work over TCP/IP connections?

AutoSYNC works over Expand or TCP/IP. See chapter "Synchronization over TCP/IP" in the NonStop AutoSYNC User's Guide for details on setting up AutoSYNC over TCP/IP.

Can AutoSYNC synchronize SQL/MX tables?

Synchronizing SQL/MX tables is not supported.

Does AutoSYNC synchronize partitioned files and tables?

Yes, it does. You can use a MAPNAMES option (similar to RESTORE) to tell AutoSYNC where the secondary partitions go on the destination system.

What happens if an object file I am trying to synchronize is currently running on the destination?

You can configure AutoSYNC to rename the destination file prior to replacing it.

What happens if I purge files on the source? Do they get purged on the destination?

By default, the files purged on the source do not get purged on the destination. You can specify the PURGE option on the fileset definition to tell AutoSYNC to do it.

To purge subvolumes, specify PURGESUBVOL.

Can I exclude some files from being synchronized?

Yes, you can configure a list of files that AutoSYNC will not synchronize in a file specified by the EXCLUDE option.

Is there a way to SQL compile my object files after they have been synchronized?

Yes. AutoSYNC has a trigger mechanism that allows you to execute a TACL script or command or an OSS shell command after a file has been synchronized or after an entire file set has been synchronized. See chapter “Triggers” in the NonStop AutoSYNC User’s Guide for details on how to set up triggers.

What happens if the synchronization of a file fails?

By default, AutoSYNC has a recovery feature that saves a copy of the file on the destination system before it attempts to replace it. If the synchronization fails for any reason, the original file is restored. The synchronization is retried at the next interval.

If I can do synchronizations on the fly, what is the advantage of using AutoSYNC to do this instead of FUP?

There are quite a few advantages to using AutoSYNC instead of FUP to move files:

- It can be done over TCP/IP as well as Expand
- AutoSYNC synchronizes Enscribe files and SQL tables. The user can specify a mapping file to move alternate key files, partitions and SQL indexes to different locations on the destination
- AutoSYNC synchronizes only the files that have changed, according to the options the user specifies (AUDITED, TIMEEXACT, so on)
- AutoSYNC synchronizes OSS files
- A trigger command can be executed after the file has been moved, such as SQL compilation of SQL programs
- AutoSYNC does extensive error reporting
- AutoSYNC has scheduling capabilities that allow for automatic replication at configured intervals

How do I make sure that files that do not belong to a user are not moved by AutoSYNC?

AutoSYNC follows the Guardian and OSS and Safeguard rules as configured on the source and destination systems. The synchronization processes run under the userid of the sync user and do not have additional access privileges.

I don't have any idea how much change there is in my environment. How do I know how dynamic my environment is?

Autosync has a reporting function that tells you how many files require synchronization and the amount of data that will be moved.

You can run the report as often as you like on any file set and evaluate the amount of change your files are undergoing.

See chapter “Reporting Function” in the NonStop AutoSYNC User’s Guide for details.

Who can use AutoSYNC?

For one-time synchronizations:

- Any user who is authorized to use FUP can use AutoSYNC, unless the system administrator explicitly prevents it

For scheduled synchronizations:

- A user must be specifically configured as Syncuser to be able to create synchronization file sets, Syncfilesets that define the parameters of synchronizations for a set of files.
- A Syncuser is configured by the AutoSYNC Administrator.
- The Administrator is designated by the AutoSYNC Authorizer.

- The AutoSYNC Authorizer is the user under which all synchronizations occur. AutoSYNC “logs down” to the Syncuser ID before it performs any synchronization, to ensure that the Syncuser only synchronizes files within his/her security privileges.
- If the authorized Syncuser are all in the same group, then user manager ID can be the Authorizer
- If the authorized Syncusers belong to different groups, then SUPER.SUPER must be the Authorizer.

MANAGEMENT & DIAGNOSTIC TOOLS

When the synchronization occurs automatically, how do I know it worked?

When the Administrator adds a Syncuser, a log file is configured for that user. The details of each synchronization are stored in the user's log file. The log files reside on the source system and both the source and destination sync processes log to that file.

For one-time synchronizations, the user can specify a LOGFILE where AutoSYNC logs the results of the synchronization.

Errors and catastrophic events are sent to an EMS collector. The EMS collector is configurable.

EMS template and an EMS filter are supplied with the product to assist in displaying AutoSYNC events.

The STATUS command displays current info on the status of the synchronization.

See chapter “Problem Resolution” in the NonStop AutoSYNC User’s Guide for details.

How can I test my scheduled synchronizations?

Prior to configuring the file sets to synchronize, you can test the set up as follows:

- To test connectivity, file set parameters and triggers (if present), use the SYNCHRONIZE command, specifying the exact same options and destination parameters you intend to configure.
- To see which files and the volume of data (in megabytes) will be moved, use the REPORT feature, also specifying the exact options and destination parameters.