

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	HPE Shadowbase Compare for SQL
Release:	Gravic Version: 6.820 HPE NonStop Shadowbase: T1122L70 (SB Repl/Guardian)
Release Date:	November 22, 2024
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2024 (www.gravic.com)
File Name:	IPM6820-SQL COMPARE.pdf

VERY IMPORTANT: Due to licensing changes in Version 6.700, existing Shadowbase installations running versions prior to (earlier than) 6.700 will require a new license file in order to install and run Version 6.820. This is true for any Shadowbase upgrade when the prior release is before Version 6.700 and the new release you plan to install is version 6.700 or after (more recent).

Contact the HPE License Manager to request a new license file <license.manager@hpe.com>. DO NOT INSTALL Shadowbase Version 6.700 (or later) software when upgrading from a version prior to Version 6.700 until a new license file has been received.

NOTE: This release contains updated software for HPE Shadowbase Compare for SQL for the HPE Integrity Nonstop X, Virtualized NonStop, and HPE Integrity NonStop I Servers.

NOTE: If this is a TCD delivery, please see [NOTE FOR TCDs](#) for TCD delivery information.

NOTE: This softdoc applies to the HPE Shadowbase Compare for SQL component, which is used to compare SQL/MP and SQL/MX databases. Other softdocs document the releases of the other components in T1122L70, including:

- HPE Shadowbase Audit Reader,
- HPE NonStop Shadowbase Guardian replication, and

- HPE NonStop Shadowbase OSS replication (see T1123).

You are advised to reference those other softdocs for the changes related to those specific components of T1122L70.

NOTE: This softdoc covers new features and corrected problems for Shadowbase Compare for HPE Integrity NonStop I servers (H06 and J06 Guardian NonStop) and HPE Integrity NonStop X/Virtualized NonStop servers (L06 Guardian NonStop).

This softdoc is available in an Adobe PDF file (.PDF). Softdoc files for SQL Compare are named IPMnnnn-SQL COMPARE.pdf (where nnnn is the Shadowbase version number).

NOTE: HPE Shadowbase Compare for SQL internally uses the nomenclature “CS-SQL-COMPARE” to reflect its original development roots. Hence, you may notice that some examples show commentary reflecting that name.

Table of Contents

SOFTWARE RELEASE DOCUMENT (SOFTDOC).....	1
Table of Contents	3
Disclaimer	4
Note for TCDs.....	5
Special Notes for Version 6.820	5
Changes in Release 6.820	5
New Features	6
Problems Fixed	9
Upgrade Considerations for Version 6.820 from 6.802.....	9
Validating Downloaded Files	9
Installation Instructions (Shadowbase Compare for SQL)	10
HPE Release File Structure.....	10
Known Problems Remaining	11

Disclaimer

We are distributing this communication in an effort to bring important information to the attention of users of the affected products. We recommend that all users determine the applicability of this information to their individual situations and take appropriate action. We do not represent or warrant that this information is necessarily accurate or complete for all user situations and, consequently, we will not be responsible for any damages resulting from the user's use or disregard of the information provided. To the extent permitted by law, we disclaim all representations and warranties, whether express, implied, statutory, or otherwise, including the warranties of the merchantability, fitness for a particular purpose, title, and non-infringement.

We expect customers of the Shadowbase product suite to “stay current” on Shadowbase releases. This means that you, the customer, should periodically upgrade your Shadowbase software to a newer release that is under support before support ends on your current release. For most customers, this means that you will want to upgrade while your release is in ‘ACTIVE’ support. Otherwise, you run the risk of not being able to get full (or even any if the release has gone ‘OBSOLETE’ end-of-service-life) support for the version you are running.

The Shadowbase Software Policy for Software Versions is described here:
<https://shadowbasesoftware.com/support/shadowbase-software-product-release-and-support-policies/>.

We encourage all customers to periodically review this material and plan for periodic upgrades to their Shadowbase software. Contact Support if you need additional information.

Note for TCDs

TCD (Temporary Code Delivery) – A software update delivered via an SPR downloadable from an FTP dropbox. A TCD is an early version, intended for customer testing only (not production usage). A TCD by definition is restricted to certain customers. Note that a “Gravic TCD” is delivered directly from Gravic, not via HPE, but otherwise has the same attributes.

A TCD is provided only to the specified customer for the purposes agreed between the customer and Gravic as to how it will be used. A TCD is provided subject to the following terms and conditions in addition to the existing written license governing the use of Shadowbase:

- A TCD is provided for evaluation and test purposes only for no more than ninety (90) days use, and is not to be used in production systems
- A TCD may not have been fully tested by Gravic, no warranties are implied as to its behavior
- A TCD is delivered directly from Gravic to the customer, it is not available from HPE/SCOUT
- As testing proceeds, iterative TCD deliveries may be necessary as issues are identified/resolved
- A TCD is temporary, after evaluation it is to be withdrawn from use by the customer
- After testing completes, a TCD may or may not subsequently be released as a Shadowbase TCF or otherwise be included in the Shadowbase product line

Please see <https://www.shadowbasesoftware.com/support/shadowbase-software-product-release-and-support-policies/shadowbase-software-release-glossary/> for additional information.

Special Notes for Version 6.820

- 1) **As of Version 6.700, customers are required to update their license / SHADPASS files when upgrading from a version before 6.700 to version 6.700 or later. Shadowbase will not start if an old SHADPASS license file is used. Contact your HPE Shadowbase account representative for a renewed license.**
- 2) The SQL Compare for MX component must be installed in a different subvolume from the SQL Compare for MX Remote Agent to avoid SQL compilation conflicts.

Changes in Release 6.820

This section summarizes the new features and problems fixed since the last General Availability release, version 6.802.

New Features

- 1) The PROGRESSINFO and final summary outputs now include a count of "transient" differences that SQL Compare has encountered. A transient difference can occur in a data replication environment due to replication latency: a mismatch is initially detected, but then by the time SQL Compare attempts to fetch the data to find out what specific column values do not match, the row has already been updated by the data replication engine so they now match. This row is no longer a mismatch and SQL Compare will move on.

Depending on the type of activity on the source and target tables, there can be scenarios where there are large amounts of these transient mismatches during a large Compare job. Since it's more work for SQL Compare to fetch the row and attempt to find the specific column values that do not match, these large batches of transient mismatches can significantly slow down a Compare job. By including the number of transient mismatches in the PROGRESSINFO and final summary, it may help the user understand the compare processing sequence and why their job was slower than expected.

- 2) The DELAY ITERATE <n> option has been added. This will delay for <n> hundredths of a second between each iteration when ITERATE is enabled for the Compare job. Delaying between iterations allows time for replication latency mismatch issues to be resolved during normal replication engine processing to avoid false positive mismatches.
- 3) A new STOP option has been added called STOP <n> PERMANENT MESSAGES. It can only be enabled when ITERATE is included in the job. It will work as follows:
 1. As soon as the first mismatch (column value difference or source/target record not found) is seen, another instance of the compare engine is started in parallel to do the iteration. The new instance will only check differences after the iteration delay has passed so that differences caused by replication latency are not falsely flagged as 'real' mismatches. If ITERATE is set higher than 1, and the mismatch persists, more iteration engines will also be started in parallel to the main compare process. Each iteration will only check remaining difference records again after the delay has passed relative to the previous iteration.
 2. Each difference still present after the given number of iterations is considered "permanent" as it is deemed to not be caused by replication latency.
 3. When the last iteration process has seen <n> permanent differences, the complete compare engine (i.e. all iteration processes and the main compare process as well) will be stopped and the final report will be printed.
 4. In this mode of operation, *permanent* differences (i.e. differences still present in the last iteration) will be reported in the final and intermediate reports in addition to the current differences.
 5. The CONTINUE command is not available in this mode of operation.

The feature provides a relatively fast way to find systematic problems in a replication setup/configuration. If the replication engine does not replicate all changes, there usually will be a big number of permanent differences. These are differences which do not go away over time when the replication latency time has passed. When you use the `STOP <n> PERMANENT MESSAGES` command, the compare engine will stop whenever the first n of these permanent differences has been detected.

This feature was added to assist customers that wanted to use a “fast check” approach to their compares, and stop after only a few permanent mismatches were encountered so that they could then investigate the cause and remediate it accordingly.

To turn the special checking for permanent differences off, use `STOP 0 PERMANENT MESSAGES`.

Notes:

- If you change any parameters relevant to searching for permanent differences (like number of iterations, delay between iterations etc.) you must always save the run before starting it, otherwise the setting will not work.
- With `STOP <n> PERMANENT MESSAGES`, the iterations looking for permanent differences will run in parallel to the main compare process. This means that the last iteration will be started last, but will typically finish first. If you set `REPORT LEVEL ITERATE` you will see this effect.

The new report level options `ITERATIONS/NO-ITERATIONS` have relevance only in connection with the `STOP <n> PERMANENT MESSAGES` option. If `STOP <n> PERMANENT MESSAGES` is set, Compare internally starts iteration processes running in parallel with the main SQL Compare program. When `NO-ITERATIONS` is in effect, the Compare process reports statistics only after all iterations are finished. When the report level is set to `ITERATIONS`, each iteration process (including the primary compare run) will report statistics at the end of each iteration. This end can be reached either if all rows have been compared or the last iteration has still found the number of differences specified in `STOP <n> PERMANENT MESSAGES`. The default for this option is `NO-ITERATIONS`.

- 4) `SET OUT FILE CODE` sets the format of the OUT file to 101 (edit type file, default value) or 180 (C type binary file). The command must be given before the `OUT FILE` command is given. If the OUT file already exists, the existing format will be preserved. The command stays in effect for the duration of the session. It can be given multiple times with different values before changing to a different OUT file.
- 5) `SET DBS PRI` and `SET PCE PRI` set the priority of the DBSRV and PCE processes (respectively) started by SQL Compare. Valid values for the priority are 0 to 199, where 0 means the process inherits the priority from the ancestor (default behavior).
- 6) An uninstall feature has been added to allow users to remove SQL Compare and all of its support files from the system. This is useful, for example, for product cleanup after

performing trial and Proof of Concept test runs of the product. The installation of SQL Compare may create the following entities:

- SQL catalogs (and schemas for MX) if chosen to do so
- Several SQL tables used by the product internally
- SQL demo tables
- SQL tables for checkpoints and storing iteration data
- Edit files containing run descriptions (will be created when saving a run)
- Executable files and text files that come with the installation package

The location of these entities on the system (disks, subvolumes) is decided at install time by editing the associated INI file.

In order to uninstall SQL Compare for either SQL/MP or SQL/MX, or either of the SQL Compare remote agents, run the installer program with the argument `--UNINSTALL--`. For example, in order to uninstall SQL Compare for SQL/MP, the command would be:

```
RUN INSTALL --UNINSTALL--
```

The uninstall procedure will remove all the support entities from the system. The procedure uses the INI file(s) configured created during installation, so the same INI file is required in order for the uninstall to be successful.

When the uninstall is done, it will leave the installer program, the INI file, and the SQLCMPR PAK file. This allows the user to simply run the install program again to reinstall it using the same INI file settings as before if desired.

Notes:

- If Compare for SQL/MP and SQL/MX are installed in the same location, uninstalling either the MP or the MX version of the product will make the product completely unusable as each of the procedures will delete common files used by both versions from the system.
 - It may occur (for example for security reasons) that the uninstall procedure can not do everything it needs to do, in that case it will report an error message which will need to be checked (and possibly resolved) manually.
 - If there is a dedicated SQL catalog created for SQL Compare, the uninstall steps will try dropping this catalog. If there are both the remote agent and the base Compare product installed, this will only work when both components are uninstalled in either order. The uninstallation of the first component will give an error message, which in this case may be ignored.
- 7) The name prefix for all spawned processes can be set with the PNPREFIX setting. It can be set to either `$<XX>` where `<XX>` is one or two characters, or it can be set `OFF` to use system-generated process names. The default is `OFF`.

Problems Fixed

- 1) Certain SQL/MP tables have partition breaks with mismatching formats. These tables have caused problems for SQL Compare in prior versions. Those issues have been fixed.
- 2) Step 5 was missing from the logged description when running INSTARA. That description has been fixed.

Upgrade Considerations for Version 6.820 from 6.802

There are no special considerations for upgrading from version 6.802 to version 6.820.

Validating Downloaded Files

Gravic provides an SHA1 checksum and the size of the release package (file) in bytes. The purpose of this is so that the user can confirm that the release package they downloaded is valid and free from corruption/tampering.

- For HPE Shadowbase releases, this information is published in the HPE SOFTDOC that corresponds with the specific SPR version of the software. The HPE SOFTDOC for a specific SPR version is available on the HPE SCOUT portal. This information can also be found on the Gravic Shadowbase website here (<https://www.shadowbasesoftware.com/releases>).
- For non-HPE Shadowbase releases, for example Shadowbase releases obtained directly from Gravic (regardless if an HPE TCD release or a direct Gravic licensee release), this information is included in the body of the Gravic email that provides the link to download the software. This information can also be found on the Gravic Shadowbase website here (<https://www.shadowbasesoftware.com/releases/gravic>).

All Shadowbase releases are provided as Windows format download files. The user should validate the downloaded release file size by comparing it with the published release file size.

In order to validate the SHA1 checksum of the downloaded file, users can run the following command on a Windows system where the file was downloaded in order to generate the checksum in their own environment:

```
certutil -hashfile <downloaded filename> SHA1
```

Users can compare the certutil-generated checksum value against the published checksum value to make sure they match. Note that any other SHA1 generation tool can also be used to create the checksum value on the downloaded file. The certutil method is just one available option for Windows environments, and is included as part of Certificate Services.

Installation Instructions (Shadowbase Compare for SQL)

- 1) Transfer the HPE Shadowbase Compare for SQL installation files to the host system using FTP or any other file transfer product capable of transferring ASCII and binary files. Binary transfer the SQLCMPR file from your PC (if you obtained the files from a Gravic FTP site) or DVD. Then ASCII transfer the RINSTARX, RINSTALL, RINSTLMX, and RINSTLRA files from the same location.
- 2) To begin the installation process, run the RINSTARX, RINSTALL, RINSTLMX, and RINSTLRA files to unpack the installation files for the desired HPE Shadowbase Compare for SQL component(s):

RUN RINSTARX to unpack the installation components for the SQL/MX Remote Agent.

RUN RINSTALL to unpack the installation components for the SQL/MP Compare program, SQLCMPE.

RUN RINSTLMX to unpack the installation components for the SQL/MX Compare program, SQLCMPMX.

RUN RINSTLRA to unpack the installation components for the SQL/MP Remote Agent.

- 3) Follow instructions in the Installation section of the ***HPE Shadowbase Compare for SQL Manual***. Each component has an installation INI file that must be edited with the desired system specific settings before the corresponding installation program is run.

HPE Release File Structure

If this release is obtained through HPE Scout, the downloaded file is a single self-extracting zip file that is intended to be run on a Windows PC. When the zip file has extracted itself, it will create a folder structure on the PC with all of the various pieces of Shadowbase in appropriate folders. This folder structure is as follows:

CONTENTS

Each component's installation files are in a separate directory in the self-extracting ZIP file or on the product DVD under the directory

\T1122H06-<SPR ID>. This file set contains the following files:

File or Directory	Description
\readme.txt	This file.
\SBEnterpriseManager	Directory containing the files required to install SEM on a PC.
\TNS-E	Directory containing the installation files for the HPE Integrity NonStop i versions of NS Repl, SAR, and Compare.
SBAuditReader	Directory containing the installation files for SAR.
SBCompare	Directory containing the installation files for Compare.
SBReplication	Directory containing the installation files for NonStop Guardian replication.
\TNS-X	Directory containing the installation files for the HPE Integrity NonStop X versions of NS Repl, SAR, and Compare.
SBAuditReader	Directory containing the installation files for SAR.
SBCompare	Directory containing the installation files for Compare.
SBReplication	Directory containing the installation files for NonStop Guardian replication.

Known Problems Remaining

- 1) Since the TCP/IP remote agent relies on the NonStop LISTNER program to establish connections between the client and server processes, certain settings such as TRGCPUS and CPULIMIT are ignored since LISTNER controls them. It is being investigated that a future version of SQL Compare may employ a custom listener program which will be able to assist in controlling these settings.
- 2) For long SQL/MX floating point column values, the precise values of the last digits are unpredictable when being inserted or selected from a table (generally, due to rounding by the file system/SQL engine). If SQL Compare attempts to compare these values, it may result in unexpected behavior depending on how the MX engine provides the values to SQL Compare. Hence, compares of high-precision floating point columns may yield differing results across runs. A case has been opened with HPE regarding this issue.

