SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	HPE Shadowbase for Other Servers
Release:	Gravic Version 6.810 T1123-AAK (SB REPL/OSS) T1124-AAG (SB REPL/HP UX) T1125-AAI (SB REPL/SUSE LINUX) T1126-AAH (SB REPL/RHEL LINUX) T1127-AAG (SB REPL/SOLARIS X86) T1128-AAG (SB REPL/SOLARIS) T1129-AAJ (SB REPL/MS WINDOWS) T1130-AAG (SB REPL/AIX)
Date:	December 12, 2024
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2024 (www.gravic.com)
File Name:	IPM6810_other_servers.pdf

NOTE: This softdoc covers new features and corrected problems for HPE Shadowbase for Other Servers, Version 6.810. It is available as a PDF file (.PDF).

NOTE: **Supported Release Version Updates (RVUS)**: This softdoc applies to all currently supported database and operating system versions as described by the Shadowbase Supported Versions Cross-Reference. Alternatively, visit the http://www.ShadowbaseSoftware.com web site for the list of supported platforms and databases. Contact your reseller or Gravic if you have any questions.

Table of Contents

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.

Mature Features and Deprecated Features

This section provides a list of mature features ("***mature features***") and deprecated features ("***deprecated features***") for HPE Shadowbase for Other Servers. Gravic will continue to support ***mature features*** and will continue to include support for them in future releases; however, Gravic will not enhance those features further or qualify them on newer versions of operating systems and databases as they become available. If you are using a ***mature feature***, Gravic highly recommends moving off the ***mature feature*** to the suggested replacement approach as soon as possible. New installations should avoid using ***mature features***; please discuss this issue with Gravic if you feel you must use a ***mature feature***.

Gravic will also continue to support ***deprecated features*** in the current release. However, Gravic plans to eliminate the deprecated features in a future release and *strongly* recommends that customers begin the planning and implementation to move off the feature *immediately*. Contact Gravic if you need assistance in planning your migration.

Table 1 - Mature Features in HPE Shadowbase for Other Servers, below, lists the features that are currently considered 'mature' by Gravic, along with the release/date where they were declared mature. Again, please note that the *mature* designation means the feature will not be enhanced and will eventually be deprecated (retired).

Table 2 - Deprecated Features in HPE Shadowbase for Other Servers, below, lists the deprecated features. The deprecated features are no longer being enhanced, and are on a schedule to be retired. The table also contains a column indicating the date of the last release that will contain the deprecated functionality (the last release before this end date will be the last release to contain this functionality). Please note that release plans are subject to change without notice and delivery dates are not guaranteed. However, Gravic currently expects that all general availability releases up to and including the date of the last release will contain the deprecated feature(s).

Please see Shadowbase Software Product Release and Support Policies on www.shadowbasesoftware.com for additional information regarding the Shadowbase Software Support Policy.

Table 1 - Mature Features in HPE Shadowbase for Other Servers

Feature	Description	Mature Release
Trigger Source for Oracle	Trigger source collectors for Oracle (OPCOL) are being replaced by the Shadowbase log source collector for Oracle (SBORLOG).	OSB 6.300 May, 2017
SQL Server Native Client	Objects that are reliant on SQL Server native client which Microsoft has labelled as obsolete 1. sbmswrtc 2. sbmssql	OSB 6.800 June, 2023

	These objects should be replaced with the generic ODBC executable types sbodbctrs and sbodbcdw.	
Shadowbase objects using Oracle Call Interface (OCI)	OCI executables (sboracle and sborwtc) are being replaced by ODBC executables to allow for new and expanded features (such as ARRAY processing)	OSB 6.800 June, 2023
Multi-ported DOC Writer	Multi-ported DOC Writer support will not be updated with future performance and usability enhancements. They can be replaced by single-ported DOC writers.	OSB 6.810 Dec, 2024

Table 2 - Deprecated Features in HPE Shadowbase for Other Servers

Feature	Description	Deprecated Release	Last Release Date
Execute Immediate Objects	The execute immediate (EI) mode of functions is deprecated and the execute immediate objects will be removed in a future release. Customers should take the announcement period to migrate their configurations to the Shadowbase cached objects that replace and supersede the EI functionality. The list of Shadowbase executable programs that use EI functionality are: 1. sbocons 2. sbmspre 3. sborapre 4. sbsypre 5. sbigdir 6. sbmswrt 7. sborcol 8. sbsywrt 9. sbigwrt 10. sborcol If you are using any of these executable programs (configured by SBMON), they will be retired on the date listed above.	OSB 6.300 May, 2017	May, 2019
SBDOCRD	The functionality in SBDOCRD has been replaced by SBFILE. SBDOCRD only supports execute immediate DOCs, which are also deprecated. Gravic supplied the source code for SBDOCRD to our customers for custom applications. If you are currently using SBDOCRD in a custom application, please contact Shadowbase Support for assistance migrating the application.	OSB 6.300 May, 2017	May, 2019
SQL92	The SQL database format SQL92 is deprecated and will be removed in a future update. This option is superseded by SQLSBV1 for DOC Writers and Direct Writers. SQL92 remains a current option for Open Collectors and Transaction Replay Servers in this release, but is expected to be deprecated for those objects as well in a future release.	OSB 6.600 April, 2020	April, 2022

UserExit Eclipse projects	The UserExit Eclipse projects have been deprecated in favor of the Makefiles shipped with the product for every platform but HPE OSS. If you are an existing customer that uses the Eclipse project please contact Support for assistance in migrating to a Makefile based workflow.	OSB 6.300 May, 2017	May, 2019

Supported Database and Operating System Versions

The database and operating systems versions supported by HPE Shadowbase can be found on the Gravic Shadowbase website. Please see Supported Databases and Platforms on www.shadowbasesoftware.com for additional information regarding Shadowbase support for databases and operating systems.

Important Note about Licenses to run Shadowbase

WARNING:

When moving from a version of Shadowbase for Other Servers software earlier than 6.700 to version 6.810, a new license file is required. Note that current license terms are still valid and remain in force. Contact the HPE License Manager (license.manager@hpe.com) or your local sales representative to request a new license file. DO NOT INSTALL Shadowbase for Other Servers software version 6.810 until a new license file has been obtained if you are upgrading from a version of Shadowbase for Other Servers software earlier than 6.700.

Upgrade Considerations for Version 6.810 from 6.800

This section is intended to give users an overview of changes that might impact their existing Shadowbase environments when they upgrade to this version. For this list of topics, refer to the details elsewhere in this document to understand whether this might have an impact on your environment.

1. Multi-ported Shadowbase objects will require the SHAD_MULTIPORTED_DOC_LEVEL to continue running properly.

2. Previous SHAD_UPDATE_NOT_FOUND values 2 and 3 have been reassigned to values 1 and 2 respectively.

3. If SHAD_TRANS_EXPECTED_ENDS is set greater than 1 and SHAD_DOC_BLOCKING is set to 2, the object will fail in this version because this is an unsupported combination of settings. One of the two parameters will need to be changed in order to start properly.

Special Notes for Version 6.810

1. Version 6.810 is a "general availability" (non-restricted TCF) release for Microsoft Windows Server systems. The previous non-restricted TCF release is version 6.800, which corresponds to HPE SPR version AAH for Microsoft Windows Server systems.

2. This document describes the changes since Version 6.800. If you are upgrading from a version of Shadowbase for Other Servers prior to that, please first see the version 6.800 software release document, IPM6800_other_servers, for a list of the changes in that version, as well as the softdocs for any additional intermediate versions since the version being upgraded.

3. **Due to licensing changes introduced in version 6.700, existing installations of Shadowbase prior to version 6.700 will require a new password file in order to run after the upgrade. For HPE Shadowbase customers, consult the HPE License Manager (license.manager@hpe.com) or your local sales representative for a new license to run this version.**

4. Immediately after install, sbmon should be opened and the LICENSE command should be run. This will ensure the license file is valid for this version of HPE Shadowbase.

5. **For Shadowbase on OSS, a special file called "shadtype" may be provided along with each license file. This file should be placed in the $SHAD_BASE/data directory just like the shadpass.ini file. This file's name should remain "shadtype" (all lowercase, no file extension.) This file is used to validate NonStop system types that are newer than the Shadowbase release to ensure that the license is correct in order to allow replication to run properly.**

6. As of version 6.600, SAP Sybase ASE databases are supported as a target using ODBC for cached statements only. Execute Immediate (EI) statements for Sybase are no longer supported. If you are using EI for Sybase, you cannot upgrade to version 6.600 or later without converting to using cached statements. Contact Support for consultation if this pertains to you.

7. HPE Shadowbase for Other Servers now obfuscates configuration data (in particular the passwords) for objects when it is stored in the COLLCONFIG data file. This version can read the configuration records created by prior releases and will automatically store the information in obfuscated format when the record is saved. Once the information has been obfuscated, the record is no longer usable in prior versions of Shadowbase. If you are upgrading from a version prior to 6.300 and you want to maintain the ability to fall back to a prior release, you must keep a copy of the collconfig.dat and collconfig.idx files for the prior release.

   You can, for example, install the release in a new directory and copy the data directory from the old directory to upgrade. This will maintain both the binaries and

the configuration files for the old release.

If you do need to recreate an obfuscated configuration with a prior release, you will need to drop and re-add the objects using the prior version of SBMON.

8. For Oracle databases, the SHAD_SQL_ERROR_* error processing parameters do not apply to database errors that occur while trying to read the table's schema to retrieve column related information. Instead, the SHAD_OCI_INCOMPLETE_SCHEMA specifies the action the Shadowbase process will take if it cannot read the table's schema. If, for example, you specify an error to skip using the SHAD_SQL_ERROR_EXCLUDE parameter, and that error occurs while reading schema, the SHAD_OCI_INCOMPLETE_SCHEMA will take precedence. If SHAD_OCI_INCOMPLETE_SCHEMA is set to SHUTDOWN, the process will stop even though the error is specified to be skipped.

9. The UserExit Eclipse projects have been deprecated in favor of the Makefiles shipped with the product for every platform but HPE OSS. If you are an existing customer that uses the Eclipse project please contact Support for assistance in migrating to a Makefile based workflow.

10. When building User Exits on Linux for x86_64, only binutils version 2.27 and later and GCC version 5 or later are supported due to an x86_64 change that effected relocation information.

11. The consumptive TFS has been removed in its entirety, the current replacement is the consumptive TRS, if you have any issues with the migration please contact Support for assistance in migrating.

12. Sybase references were removed from both the sbinstall and sbfile scripts. Sbinstall would previously print what Sybase version was being used but now that is not relevant and could be confusing. Sbfile took a Sybase type as an argument with COLLCONFIG and RESETSBCONFIG commands.

13. The sbinstall script for windows has been updated so that it will now fail and inform the user if SHAD_BASE is not set correctly as a system variable. This may be due to it being set as a user level environment variable or not at all but this notification should make it far easier to find a fairly common issue in new environments.

14. All of the 3rd party vendor libraries shipped with Shadowbase are now located in the $SHAD_BASE/lib/optional directory. If any errors with missing libraries are encountered when attempting to start any Shadowbase objects, look for the missing library in this location and move it up a level into the $SHAD_BASE/lib directory and start the object again.

15. The array-based TRS (sbarrtrs) and DW (sbarrdw) do not use the value entered at the "Convert Updates to Inserts when not found (Y/[N])?" prompt when adding the

object through SBMON. They support multiple options and use the value specified by the *SHAD_CONVERT_UPDATES* and *SHAD_UPDATE_NOT_FOUND* shadparm.ini parameters.

Changes in Version 6.810 (since Version 6.800)

New Features

This section provides a summary of the features added to HPE Shadowbase for Other Servers products since the previous general availability release (V6.800).

1. In addition to the array-based ODBC TRS (sbarrtrs), an array-based Direct Writer (sbarrdw) has been added as well.

2. The array-based ODBC TRS and Direct Writer have been enhanced to support "fuzzy" replication by converting updates to missing rows on the target to insert operations. See the *Array-based ODBC Statement Conversion* section for more information.

3. The array-based ODBC TRS and Direct Writer's statistics report have been enhanced to better track the performance of the processes. See the *Array-based ODBC Statistics Reporting* section for more information.

4. A new sbfile command DUMPLATENCY has been added for the purposes of reading latency data for a given Shadowbase object. This is provided as an alternative to the parameters SHAD_LATENCY_MONITORING_THRESHOLD and SHAD_LATENCY_MONITORING_FREQUENCY for obtaining latency information relating to an object.

   Usage: sbfile DUMPLATENCY <object name>
   object name - configured replication object name

5. New settings have been added to provide information to Shadowbase support to aid in diagnosing disk and network issues on Shadowbase deployments. New parameters for these settings are SHAD_ENABLE_CTREE_IO_STATS, SHAD_CTREE_IO_STATS_TIME_THRESHOLD, and SHAD_CTREE_IO_STATS_IUD_THRESHOLD. These settings are meant for coordinated use with Shadowbase support and is not recommended for regular use. For more details about these settings, refer to the *New & Modified shadparm.ini Parameters* section. These settings currently support all Shadowbase objects except trigger-based SQL Server Source Collectors, Direct Writers, and DOC Cleaners.

6. TCP/IP host names longer than 20 characters are now supported by using the SHAD_HOST_NAME parameter instead of the "ip address" field when

configuring the object in sbmon. Refer to the *New & Modified shadparm.ini Parameters* section for more information.

7. Passwords entered at sbmon prompts are no longer echoed back to the terminal when configuring or editing objects.

8. Replication objects can be configured to log warnings to their errlogs when the latency exceeds a specified threshold. Refer to the SHAD_LATENCY_MONITORING_THRESHOLD and SHAD_LATENCY_MONITORING_FREQUENCY parameters in the *New & Modified shadparm.ini Parameters* section.

9. When Shadowbase attempts to apply a delete event for a row that does not exist, it's now possible to log a warning in that situation. The default behavior remains in effect to not log any warnings. Refer to the SHAD_MISSING_DELETE_LEVEL parameter in the *New & Modified shadparm.ini Parameters* section.

10. Additional error codes will be logged in errlog messages when they are set to a significant value. This may be helpful in diagnosing errors when they occur. This feature can be disabled with the SHAD_LOG_ERROR_CODES option as described in the *New & Modified shadparm.ini Parameters* section.

11. This release of Shadowbase includes a utility for all platforms except OSS called "sblicinfo". This tool extracts information about the system where Shadowbase is to be installed for the purpose of generating a correct license for it. Refer to the *System license information extraction tool (sblicinfo)* section for details on how to use it.

Problems Corrected

This section provides a summary of the problems corrected in HPE Shadowbase for Other Servers products since the previous general availability release (V6.800).

1. The array-based ODBC TRS and Direct Writer's error reporting has been improved with respect to identifying the specific row or rows causing the error. When an error occurred applying a batch of data in the previous version, the processes frequently could not determine the specific row causing the error. The processes will now try to re-apply the data row by row where possible to determine specific dataset that was in error.

2. SHAD_UNIX_FILE_PERMISSIONS now works as intended for OSS environments.

3. The Linux COLUMNS environment variable will no longer give unexpected results to sbmon STATUS commands.

4. For a TRS/DW object that does not have Insert Not Found enabled, when an update is attempted for a row that doesn't exist and the update is skipped, in previous versions of Shadowbase this would not result in a reject record being created for the Log Server DOC. This has been fixed so that now a reject record is created as expected.

5. Explicitly setting SHAD_STOP_AT_END_OF_DOC to its default value of 0 now works as expected.

6. SHAD_SQL_ERROR_INVALID_NUM now takes effect properly for Hana BOOLEAN values.

7. The "shad" environment script for OSS now properly sets the _RLD_FIRST_LIB_PATH variable instead of LD_LIBRARY_PATH.

8. SQL Server source trigger create script now handles the case where the same table exists in multiple schemas in a database instance.

9. Some SQL errors resulting from INF statements could result in errlog messages with a misleading SQL statement. This issue has been fixed so the SQL statements in the errlog are accurate.

10. SHAD_SQL_ERROR_INVALID_NUM_SKIP was not working as expected for Windows environments. This has been fixed.

11. A problem in the errlog for TCP/IP objects was fixed so that it no longer logs an increasingly long message regarding socket sizes.

New Features Description

Array-based ODBC Statement Conversion

The array-based ODBC TRS and Direct Writer now support converting update operations on missing target rows to inserts. The options for converting update operations are controlled by the SHAD_CONVERT_UPDATES parameter. You can specify one of the following options:

- No conversion. If no conversion is selected, updates to missing rows are not converted. The existing shadparm.ini parameter *SHAD_UPDATE_NOT_FOUND* specifies the behavior if the update fails due to a missing row.
- Reapply the entire update statement batch one row at a time, converting only the specific missing rows to inserts.

For more details refer to the SHAD_CONVERT_UPDATES parameter in the *New & Modified shadparm.ini Parameters* section.

Array-based ODBC Statistics Reporting

The array-based TRS and Direct Writer statistic reporting has been significantly enhanced to provide improved insights into the statement processing performance of the object.

Statements received from the source may be merged or transformed into different statements prior to application to the target, so the array-based ODBC statistics are tracked for two statements types: as received from the source database and as applied into the target database. *SHAD_ARRAYSTATS_PRINT_CATEGORY* shadparm.ini controls which statistics are printed:

- None.
- Received statement statistics only.
- Executed statement statistics only.
- Both received and executed statement statistics.

Commit statistics are tracked and are printed if either of the statement statistics are selected. The column definitions for each section are described in the following sub-sections:

- Received Statement Statistics
- Executed Statement Statistics
- Commit Statistics

Statistics are collected and tracked at the SQL statement level. You can specify the level of detail to be printed using the *SHAD_TRACK_PERFORMANCE* shadparm.ini parameter. Statistics can be printed for:

- Each SQL statement.
- Total by table.
- Total over all statements.

At the SQL statement display level, the statements are grouped by table name (in alphabetical) order, with a line for each statement (in alphabetical order) plus a total line. At the end of the section, the overall totals for delete, insert, and update statements are printed, along with a grand total line.

The total by table level displays one line per table (alphabetical by table name), providing the total statistics for the table, plus the overall total lines. The overall totals level just displays the delete, insert, update and grand total lines for all statements.

You can specify either a time interval (*SHAD_ARRAYSTATS_TIME_INTERVAL*) or a statement interval (*SHAD_STATS_FREQUENCY*) for the statistics. If the *SHAD_STATS_FREQUENCY* is selected, the statistics will be printed after the specified number of statements are received from the source.

Note: To ensure consistency between the as-received and as-executed statistics, the statistics are printed on a commit boundary. They will be printed when the first commit is processed after the interval has been exceeded. As a result, slightly more statements or more time may be included in the statistics than are specified by the parameters.

Received Statement Statistic Report Columns

The received statement statistics section has the following columns:

- Table Name: The table name. If the table name is too large for the column, only the rightmost portion of the name will be displayed.
- IUD: The type of statement (insert, update, or delete) or TOTAL for total lines.
- Total Prepares: The number of prepares since the TRS or DW was started.
- Total Frees: The number of times the statement was freed. SQL statements are freed when the number of statements in the SQL statement cache exceeds the maximum cache size and the statement is the least recently used.
- Delta Prepares: The number of prepares since the last time the statistics were printed.
- Delta Frees: The number of times the statement was freed since the last time the statistics were printed.
- Total Prep Time: The total time spent preparing the statement.
- Avg Prep Time: The average prepare time for the statement.
- Min Prep Time: The minimum prepare time for a single prepare.

- Max Prep Time: The maximum prepare time for a single prepare.
- Stmts Received: The number of statements (individual row updates) that were received since the last time the statistics were printed.
- Stmts Merged: The number of received statements that were merged into another statement and not executed.
- Stmts Converted: The number of received statements that were converted to other statements for execution.
- Stmts Created: The number of statements created from the converted statements.
- Statement Text:
    - For the individual statement statistics line, the statement text (if enabled)/
    - For the table statistics line, the table name in the format TABLE-<table name> Total/
    - For overall statistic lines, the statement type (insert, update, or delete) in the format All tables <statement type> statement total

Notes:

1. All times are in micro-seconds.
2. The Total Prepares, Total Frees, Total Prep Time, Avg Prep Time, Min Prep Time, and Max Prep Time values do not reset with each interval and are based upon all activity since the TRS or Direct Writer was started. All other statistics reset at each interval and represent the values since the last time statistics were printed.
3. The statement text column has either "@" (statement level statistics), "@@" (table level statistics) or "@@@" prepended to allow quick extraction of the statistics from the log.
4. All received statements are prepared the first time they are received from the source system even if they will never be used. Thus, even UPDATE statements that are always converted to DELETE/INSERT combinations will be listed and will have prepare statistics associated with them.

Executed Statement Statistic Report Columns

The executed statement statistics section has the following columns:

- Table Name: The table name. If the table name is too large for the column, only the rightmost portion of the name will be displayed.
- IUD: The type of statement (insert, update, or delete) or TOTAL for total lines.
- Total Prepares: The number of prepares since the TRS or DW was started.
- Total Frees: The number of times the statement was freed. SQL statements are freed when the number of statements in the SQL statement cache exceeds the maximum cache size and the statement is the least recently used.
- Delta Prepares: The number of prepares since the last time the statistics were printed.
- Delta Frees: The number of times the statement was freed since the last time the statistics were printed.

- Total Prep Time: The total time spent preparing the statement.
- Avg Prep Time: The average prepare time for the statement.
- Min Prep Time: The minimum prepare time for a single prepare.
- Max Prep Time: The maximum prepare time for a single prepare.
- # Batch Executes: The number of times the SQL statement was executed in batch mode with one or more row updates per execute.
- Batch Exec Time: The total time spent executing the batches.
- Avg BE Time: The average time for one batch execute.
- Min BE Time: The minimum time for a single batch statement execution.
- Max BE Time: The maximum time for a single batch statement execution
- # Rows Submitted: The total number of rows updated in the SQL statement batch executions.
- Avg Event Exec Time: The average execute time per row (Batch-Exec-Time / #-Rows-Submitted)
- Avg Batch Size: The average size (number of rows updated) for the batch executions.
- Min Batch Size: The smallest batch size.
- Max Batch Size: The largest batch size.
- # Single Executes: The number of rows that were submitted for execution a single row at a time. If the batch execute of multiple rows fails, the TRS or Direct Writer will fall back to executing the batch a row at a time to identify the failure. This value represents the number of rows that were submitted one at time when this occurs.
- Single Exec Time: The elapsed time for the single row executes.
- Avg SE Time: The average time for the single row executes.
- Min SE Time: The minimum time for the single row executes.
- Max SE Time: The maximum time for the single row executes.
- # Rows Skipped: The number of row updates that were not applied but were skipped.
- # Rows Converted: The number of applied rows that were converted to one or more different SQL statements and re-applied.
- # Rows Created: The number of rows applied that were created from another type of SQL statement.
- Statement Text:
    - For the individual statement statistics line, the statement text (if enabled)/
    - For the table statistics line, the table name in the format TABLE-<table name> Total/
    - For overall statistic lines, the statement type (insert, update, or delete) in the format All tables <statement type> statement total

Notes:

1. All times are in micro-seconds.
2. The Total Prepares, Total Frees, Total Prep Time, Avg Prep Time, Min Prep Time, and Max Prep Time values do not reset with each interval and are based upon all

activity since the TRS or Direct Writer was started. All other statistics reset at each interval and represent the values since the last time statistics were printed.
3. The statement text column has either "@" (statement level statistics), "@@" (table level statistics) or "@@@" prepended to allow quick extraction of the statistics from the log.

Commit Statistic Report Columns

The commit statistics section has the following columns:

- # Commits: The number of commits since the last time the statistics were printed.
- Tot Commit Time: The total elapsed time executing the commits.
- Avg Commit Time: The average time for a single commit.
- Min Commit Time: The minimum time to execute a single commit.
- Max Commit Time: The maximum time to execute a single commit.

Notes:
1. All times are in micro-seconds.
2. All statistics reset at each interval and represent the values collected since the last time statistics were printed.

System license information extraction tool (sblicinfo)

Each Shadowbase delivery except OSS includes a tool called sblicinfo at the highest level of the delivery package. The purpose of this tool is to run it on a platform where Shadowbase is intended to be installed and run. This tool will collect information about the system into a .cache file. That file can then be provided along with the request for a new license. This will help ensure that the license generated for the system has all the correct attributes necessary for the system and it will allow Shadowbase to run as expected.

In order to use this tool, first copy it onto the server where Shadowbase will be installed. It's a simple executable file, so it will need permission to run on that server. When it's in place, run it from a command line with no arguments. Here's an example in Windows:

```
> sblicinfo
Copyright Gravic, Inc. 1995-2022. USAGE SUBJECT TO THE TERMS OF A WRITTEN AGREEMENT.
Portions
copyrighted and licensed from third parties as specified in the included
release README file. PATENTS: SEE WWW.GRAVIC.COM/GRAVICLABS/PATENTS/PRODUCTS

File written to SYSTEM4.cache
```

This generates a file called <system-name>.cache in the same directory as the executable. This file should be gathered and sent along with the license request. This .cache file is read-only and it is not alterable once it's generated by the system. It's also unusable until

the .cache file is processed by the person processing the license request. When that happens, the .cache file will provide information like so:

```
Copyright Gravic, Inc. 1995-2022. USAGE SUBJECT TO THE TERMS OF A WRITTEN AGREEMENT.
Portions
copyrighted and licensed from third parties as specified in the included
release README file. PATENTS: SEE WWW.GRAVIC.COM/GRAVICLABS/PATENTS/PRODUCTS

System Information:
  File Version: 1
     Raw Flags: 3
     Node Name: SYSTEM4
       OS Name: Windows
System Version: Windows 10 Enterprise
System Release: 19045
           SMT: 2
        Endian: little
     CPU Count: 1
 Logical Cores: 12
   Total Cores: 6
      Checksum: e60bdb60
```

These values are useful to ensure that the license generated is accurate and no errors occur during the license request process.

New & Modified SBMON Commands

This section provides a description of the SBMON operational commands that are either new or modified since the previous general availability release (V6.800).

New & Modified shadparm.ini Parameters

This section provides a description of the parameters for the shadparm.ini configuration file that are either new or modified since the previous general availability release (V6.800).

SHAD_ARRAYSTATS_PRINT_CATEGORY

Initial Version:	V6.810	**Syntax:**	SHAD_ARRAYSTATS_PRINT_CATEGORY =<num>
Last Change:	--		

Default Value	3
Other Servers	TRS (sbarrtrs only), DW (sbarrdw only)
Valid Settings	0 – OFF 1 – print statistics for RECEIVED STATEMENTS ONLY 2 – print statistics for EXECUTED STATEMENTS ONLY 3 – print statistics for BOTH executed and received statements
Basic/Advanced	Advanced

Description

This parameter controls the categories for which array-based statistics will be logged. Received statements refer to statements received from the source side, and executed statements refer to statements that have been successfully applied to the target database.

SHAD_ARRAYSTATS_TIME_INTERVAL

Initial Version:	V6.810	**Syntax:**	SHAD_ARRAYSTATS_TIME_INTERVAL =<seconds>
Last Change:	--		

Default Value	0
Other Servers	TRS (sbarrtrs only), DW (sbarrdw only)
Valid Settings	Number of seconds
Basic/Advanced	Advanced

Description

This parameter will control how often array-based replication statistics are output based on a time threshold defined in seconds.

Notes

When SHAD_TRACK_PERFORMANCE is greater than one and SHAD_STATS_FREQUENCY is not defined, SHAD_ARRAYSTATS_TIME_INTERVAL must be defined to a non-zero value in order to print statistics. SHAD_ARRAYSTATS_TIME_INTERVAL and SHAD_STATS_FREQUENCY cannot be enabled at the same time.

SHAD_CONVERT_UPDATES

Initial Version:	V6.810 **Syntax:** SHAD_CONVERT_UPDATES=<num>
Last Change:	--
Default Value	0
Other Servers	TRS (sbarrtrs only), DW (sbarrdw only)
Valid Settings	0 – Disabled 3 – Convert to INSERT operation on Not Found error.
Basic/Advanced	Advanced

Description

The SHAD_CONVERT_UPDATES parameter specifies what, if any, statement conversions the array-based TRS or DW should make if it detects one or more updates not applied during a batch due to missing rows in the target database.

Notes

1. Option 3 will reapply the updates in the original batch one at a time to identify the specific row or rows that failed. Only those rows will be converted to inserts and applied.
2. If option 3 is selected, it will override the stop/skip behavior specified by the SHAD_UPDATE_NOT_FOUND parameter. However, the logging of the successfully converted events is still controlled by setting the SHAD_UPDATE_NOT_FOUND parameter. When SHAD_UPDATE_NOT_FOUND is set to 0 or 1, events will be logged when they are converted. When SHAD_UPDATE_NOT_FOUND is set to 2, events will not be logged when they are converted.
3. This parameter overrides the value entered when the object is added through SBMON at the "Convert Updates to Inserts when not found (Y/[N])?" prompt.

SHAD_CTREE_IO_STATS_IUD_THRESHOLD

Initial Version:	V6.810 **Syntax:** SHAD CTREE IO STATS IUD THRESHOLD=<num>
Last Change:	--
Default Value	0 (Disabled)
Other Servers	All (Except Direct Writer, DOC Cleaner, and SQL Server OPCOL)
Valid Settings	0: Disable insert/update/delete based CTREE measurement threshold 1 – 2147483647: Number of inserts, updates, and deletes between CTREE measurement stat dumps
Basic/Advanced	Advanced

Description

This parameter sets a threshold based on the number of inserts, updates, and deletes processed by a given object for how often stats produced by the CTREE disk IO performance measurement system should be dumped to the relevant error log.

Notes

SHAD_ENABLE_CTREE_IO_STATS must be set for this parameter to have any effect. If this parameter and SHAD_CTREE_IO_STATS_TIME_THRESHOLD is also set, Shadowbase will fail with a configuration error.

Related

SHAD_ENABLE_CTREE_IO_STATS
SHAD_CTREE_IO_STATS_TIME_THRESHOLD

SHAD_CTREE_IO_STATS_TIME_THRESHOLD

Initial Version:	V6.810	**Syntax:**	SHAD_CTREE_IO_STATS_TIME_THRESHOLD=<num>
Last Change:	--		
Default Value	0 (Disabled)		
Other Servers	All (Except Direct Writer, DOC Cleaner, and SQL Server OPCOL)		
Valid Settings	0: Disable time based CTREE measurement threshold 1 – 2147483647: Time, in seconds, between CTREE measurement stat dumps		
Basic/Advanced	Advanced		

Description

This parameter set a time threshold for how often stats produced by the CTREE disk IO performance measurement system should be dumped to the relevant error log.

Notes

SHAD_ENABLE_CTREE_IO_STATS must be set for this parameter to have any effect. If this parameter and SHAD_CTREE_IO_STATS_IUD_THRESHOLD is also set, Shadowbase will fail with a configuration error.

Related

SHAD_ENABLE_CTREE_IO_STATS
SHAD_CTREE_IO_STATS_IUD_THRESHOLD

SHAD_ENABLE_CTREE_IO_STATS

Initial Version:	V6.810	**Syntax:**	SHAD_ENABLE_CTREE_IO_STATS=<num>
Last Change:	--		
Default Value	0 (Disable)		
Other Servers	All (Except Direct Writer, DOC Cleaner, and SQL Server OPCOL)		
Valid Settings	0: Disable CTREE IO performance measurement 1: Enable CTREE IO performance measurement		
Basic/Advanced	Advanced		

Description

This parameter controls whether or not the CTREE disk IO performance measurement system is active. If activated for an object, This parameter will cause CTREE operations to be measured/profiled in preparation to dump the info in a table at a selected interval.

Notes

This parameter is only meant to be used in coordination with Shadowbase support, in order to diagnose problems with a Shadowbase environment. Exactly one of either SHAD_CTREE_IO_STATS_TIME_THRESHOLD or SHAD_CTREE_IO_STATS_IUD_THRESHOLD must be set in conjunction with this parameter or a Shadowbase object will exit with configuration errors. When stats are dumped to the error log as a result of one of the threshold parameters, CTREE stats will be reset internally, to make sure each statistic period is

independent from others. On an actual statistic dump, in addition to CTREE specific statistics, a Shadowbase object will also log the threshold being logged on, latency between a NonStop source and the current object, the number of operations per second (including begins and commits), the number of operations per second of ctree execution, a count of all operations, a profiling of operating system flush calls, and how long the process of logging the entire statistic block took.

Related

SHAD_CTREE_IO_STATS_IUD_THRESHOLD
SHAD_CTREE_IO_STATS_TIME_THRESHOLD

SHAD_ENABLE_SEQUENCE_OBJECT_REPLICATION

Initial Version:	v6.700	**Syntax:**	SHAD_ENABLE_SEQUENCE_OBJECT_REPLICATION=<num>
Last Change:	--		
Default Value	0 (OFF)		
Other Servers	TRS | DW (OSS)		
Valid Settings	0 = Sequence object replication disabled for TRS and DW 1 = Sequence object replication is enabled for TRS and DW		
Basic/Advanced	Advanced – See Basic vs. Advanced Parameters section		

Description

This parameter controls whether or not the replication of SQL/MX internal and external sequence objects is enabled for Transaction Replay and Direct Writer Shadowbase objects. This parameter is required if the corresponding NSB source environment is configured to replicate SQL/MX Sequence Object columns.

Notes

Setting this parameter to anything other than 0 or 1 will cause the parameter to be disabled. See the *HPE NonStop Shadowbase SQLMX Manual* for more information on enabling collection of Sequence Object events on the source side. Replicating internal and external sequence objects will cause Shadowbase TRS and DW objects to apply non-standard SQL statements, therefore care should be taken to ensure that both source and target sides of replication are configured consistently. Disabling this parameter while the source side is still configured to replicate sequence object operations will cause the target TRS/DW to fail upon attempting to replay these non-standard SQL statements. Additionally any DOC that has had sequence object statements inserted into it will need this parameter to be enabled in order to replay the data.

SHAD_FORCE_SAVEPOINT_WITH_RI

Initial Version:	V6.810	**Syntax:**	SHAD_FORCE_SAVEPOINT_WITH_RI=<num>
Last Change:	--		
Default Value	2 (OSS), 0 (Windows, Linux, Unix)		
Other Servers	DW (SQL/MX), TRS (SQL/MX)		
Valid Settings	0 (None) 1 (Off) 2 (On)		
Basic/Advanced	Advanced		

Description

This parameter is for use when source data coming from a NonStop system is suspected to be arriving at an OSB SQL/MX target environment out of order, in a way that violates the referential integrity of the SQL statements. Setting this parameter to 2 (On) will condition the target MX database to use more advanced features to ensure that referential integrity between SQL statements is not violated. Setting this parameter to 1 (Off) will explicitly condition the SQL/MX database to not use these advanced features, while setting it to 0 (None) will result in no commands relating to this parameter being issued to the database.

SHAD_HOST_NAME

Initial Version:	V6.810	**Syntax:**	SHAD_HOST_NAME=<host name / IP>
Last Change:	--		
Default Value	NULL		
Other Servers	All		
Valid Settings	Valid host name / IP address		
Basic/Advanced	Advanced		

Description

This parameter is used to override the IP address/host name that is set for Shadowbase objects in the collconfig. This can be used to surpass the size limitation for collconfig host names.

Notes

This does NOT alter the collconfig value for the host name, it simply overrides it.

SHAD_LATENCY_MONITORING_THRESHOLD

Initial Version:	V6.800	**Syntax:**	SHAD_LATENCY_MONITORING_THRESHOLD=<num>
Last Change:	--		
Default Value	-1 (Disabled)		
Other Servers	All		
Valid Settings	-1: Disable latency monitoring 0 to 2147483647: Number of seconds of latency a transaction can have before being considered to have violated the latency monitoring threshold		
Basic/Advanced	Basic		

Description

This parameter sets the number of seconds of latency a transaction can have before being flagged by the Shadowbase latency monitoring system. If the latency of a transaction exceeds the threshold set by this parameter, and a latency message has not been output in the previous number of seconds specified by SHAD_LATENCY_MONITORING_FREQUENCY, then a message will be logged informing the user of the latency measurements that violated the threshold. Setting this parameter to -1 or any negative value will disable latency monitoring altogether for the configured object.

SHAD_LATENCY_MONITORING_FREQUENCY

Initial Version:	V6.800	**Syntax:**	SHAD_LATENCY_MONITORING_FREQUENCY=<num>
Last Change:	--		
Default Value	60 seconds		
Other Servers	All		
Valid Settings	0 to 2147483647: Number of seconds that should occur between latency monitoring messages		
Basic/Advanced	Basic		

Description

This parameter controls the number of seconds that should occur between latency monitoring messages. It can stop the error log from being inundated with messages in cases where latency spikes for a period of time, while still ensuring that the appropriate number of latency messaging appears.

Notes

Setting this parameter to 0 will tell Shadowbase to log all latency values that fall below the monitoring threshold, regardless of how recent the previous message was logged.

SHAD_LOG_ERROR_CODES

Initial Version:	V6.810 **Syntax:** SHAD_LOG_ERROR_CODES=<num>
Last Change:	--
Default Value	1 (Enable)
Other Servers	All
Valid Settings	0: Disable CTREE error code logging 1: Enable CTREE sysiocod and errno logging when an error occurs
Basic/Advanced	Advanced

Description

This parameter controls whether errno and sysiocod will be logged. Previous behavior was to not log either but this will log them if they are set to anything other than 0 in the event of an error.

SHAD_MISSING_DELETE_LEVEL

Initial Version:	V6.810 **Syntax:** SHAD_MISSING_DELETE_LEVEL=<num>
Last Change:	--
Default Value	0
Other Servers	All
Valid Settings	0 – OFF 1 – Log and continue 2 – Log and shutdown
Basic/Advanced	Basic

Description

This parameter will log delete statements that fail when attempted to be applied to the target side. Depending on the settings, replication will either continue as usual or stop and shutdown. This can be useful for tracking delete-not-found errors where certain statements do not already exist in the target table.

SHAD_MULTIPORTED_DOC_LEVEL

Initial Version:	V6.810 **Syntax:** SHAD_MULTIPORTED_DOC_LEVEL=<num>
Last Change:	--
Default Value	0
Other Servers	All
Valid Settings	0 – OFF 1 – Enable multiple port use
Basic/Advanced	Advanced

Description

This parameter will control whether or not a user can define multiple ports within a Shadowbase object.

Notes

During configuration, objects will not allow the entry of more than one port when SHAD_MULTIPORTED_DOC_LEVEL is not manually set to 1.

SHAD_TRACK_PERFORMANCE

Initial Version: V3.933 **Syntax:** SHAD_TRACK_PERFORMANCE =<num>	
Last Change: V6.810	
Default Value	0
Other Servers	All
Valid Settings	0 – OFF 1 – Enable performance statistic collection 2 – Enable performance statistic collection with additional prepared statement information. For array-based Shadowbase objects, statistics will display statement, table, and grand totals. 3 – **Array-based Shadowbase objects only:** Enable performance statistic collection, statistics will display table and grand totals. 4 – **Array-based Shadowbase objects only:** Enable performance statistic collection, statistics will display grand totals.
Basic/Advanced	Advanced

Description

This parameter controls whether performance statistics are written to the log file. The statistics will be logged to the object's errlog as determined by the SHAD_STATS_FREQUENCY parameter value. For array-based Shadowbase objects, there is the option to use SHAD_ARRAYSTATS_TIME_INTERVAL instead. **One of these two frequency parameters must be defined for SHAD_TRACK_PERFORMANCE values greater than 0.**

Notes

For array-based objects, depending on the setting, the log will display the summed statistics for each level. This includes the grand totals and the sum for each individual table. The highest level of logging (2) will display the grand totals and table totals, and it will also display the statistics for every statement in each table.

Related

SHAD_STATS_FREQUENCY
SHAD_ARRAYSTATS_TIME_INTERVAL
SHAD_ARRAYSTATS_PRINT_CATEGORY

SHAD_UPDATE_NOT_FOUND

The values for SHAD_UPDATE_NOT_FOUND have changed from v6.800. Option 1 was original reserved for converting updates to delete / insert pairs, which was introduced with this release and is controlled by the SHAD_CONVERT_UPDATES parameter. Old options 2 and 3 have become options 1 and 2, respectively.

Initial Version:	v6.800	Syntax:	SHAD_UPDATE_NOT_FOUND=<num>
Last Change:	v6.810		
Default Value	0		
Other Servers	TRS (sbarrtrs only), DW (sbarrdw only)		
Valid Settings	0: Stop 1: Skip and log 2: Skip without logging		
Basic/Advanced	Advanced – See Basic vs. Advanced Parameters section		

Description

This parameter specifies how the array-based object behaves when it detects that not all rows in a batch of updates were applied. This indicates that one or more of these rows do not exist in the target database.

Depending on the batch size and ODBC implementation, the object may not have enough information to determine which rows were skipped. In this case, the TRS will log a message with as much information as is available.

Notes:

If SHAD_CONVERT_UPDATES option 3 is selected, it will override this parameter's stop/skip behavior. Missing rows will be converted to INSERT operations and retried, even if SHAD_UPDATE_NOT_FOUND specifies stopping. In this case, if SHAD_UPDATE_NOT_FOUND is set to 0 or 1, the specific missing rows will be logged. If set to 2, the missing rows will not be logged.

New & Modified User Messages

This section provides a description of the important user messages that are either new or modified since the previous general availability release (V6.800).

Due to modifications to internal Shadowbase timestamp handling, transaction timestamp values generated by DOC Writer objects in response to the SHAD_DOC_GENERATE_AUDIT_LEVEL parameter will now display with microsecond rather than second precision in the output of SBFile calls.

Known Problems Remaining in Version 6.810

1. There is a one-to-many relationship between the SSQLD000.dat file and the series of SSTMNCxxx.dat files within a given cached SQL statement DOC database. If the SSQLD000.dat or SSQLD000.idx files are deleted or otherwise modified, replication may fail. If the SSQLD000.dat file becomes unusable or is accidentally removed, contact support for assistance and resolution to this issue.

   Note: The above condition causes replication to the target database to fail. However, the target database is not adversely affected; target database corruption does not occur.

2. Audit Log: The Audit Log image column SHAD_EVENT_TIMESTAMP reflects the wall clock time in which the Shadowbase NonStop Consumer process replicated the event to the HPE Shadowbase for Other Servers DOC database. This column is meant to reflect the NonStop audit trail event timestamp. That is, this timestamp does not represent the events source database activity time, but rather the time the event was replicated to the Open Server DOC database. This issue will be changed in an upcoming Shadowbase NonStop release, such that the SHAD_EVENT_TIMESTAMP column will contain the time the event was recorded in the HPE NonStop system audit trail.

3. DOC Writer and Source Collector restarts the TRS/TFS even if the TRS/TFS was stopped by SBMON. When enabled, the DOC Writer and the Source Collector will monitor and restart TRS/TFS if it stops running. If a TRS/TFS was manually stopped by an SBMON STOP command (normal shutdown), the DOC Writer and or Source Collector will continue to restart the TRS/TFS instead of leaving it in a stopped state.

4. ***The SBMON ROLL command must not be used on actively replicating objects or DOC corruption may result***. DOC rolls generated internally by the DOCW or collector object are handled correctly. However, there is a risk that a DOC roll triggered by a user issuing the ROLL command may do so while the replication object is in a critical state. If a manual SBMON ROLL command is required, shut down the relevant DOC writing replication object(s) (e.g., OPCOL, DOC Writer) and all database user sessions for source collection objects prior to issuing the ROLL command.

5. Use of Reserved Words as target SQL Table Column Names. In particular, the following reserved words are not supported for HPE Shadowbase for Other Servers target replication:

   AND
   WHERE
   VALUES

6. The HPE Shadowbase for Other Servers, Oracle Log Based Source Collector (*sborlog*) has the following limitations:

    a. HPE Shadowbase Log Source for Oracle does not support replication of tables with the following column types:

        BLOB
        CLOB
        NCLOB
        BFILE

    b. Selection criteria for the source collection is limited to including/excluding table names and users. These selections can be combined to select all but a specified set of tables for a user (specify the user and the set of tables to exclude), all tables for all users except for a specified set of users (specify the tables to be included and the user or users to exclude), and a subset users and tables (specifying both the tables and users). However, more complex selections may require multiple SBORLOG processes to be configured. If, for example, Users 1 and 2 both have tables named A and B, and you want to collect data from User 1's table A and User 2's table B, you will need to configure two SBORLOG processes.

    c. Shadowbase Log Source for Oracle currently only collects DML events (Inserts, Updates, and Deletes).

    d. HPE Shadowbase Log Source for Oracle does not support collecting tables with table names that are reserved words.

    e. Detailed collection stats are not supported (shadparm.ini parameter SHAD_OPCOLLECT_LOG_STATS=3). If you specify SHAD_OPCOLLECT_LOG_STATS=3, no statistics will be collected.

7. When configuring triggered-based replication for Microsoft SQL Server on a Windows 2008 R2 or later server, the generated trigger scripts may not be written to the Shadowbase data directory due to permissions problems. SQL Server writes the trigger scripts to the directory. If the SQL Server user does not have the correct permissions, the configuration will fail. The SQL Server user also needs execute access to the bin directory within the installation to collect data.

    SQL Server needs Full Access (F), Object Inherit (OI) and Container Inherit (CI) permissions to the Shadowbase data and bin directories. If you set the permissions on the installation directory *(%shad_base%)*, both directories will inherit the permissions. You can check and set these permissions logged on as an administrator by using *icacls* in the command prompt. To check permissions on the Shadowbase installation directory:

```
CD %shad_base%
icacls *
```

and look for the SQL Server owner. If needed, grant the correct permissions on the Shadowbase base directory before configuring triggered based replication:

```
CD %shad_base%
icacls /grant:R <user>:(OI)(CI)F /T
```

8. Under certain transaction profiles when replicating from Other Servers to HPE NonStop Guardian, the Consumer will stop with an EMS message error message (#2017):

   ```
   SBOS-TO-NSK COVERSION BUFFER OVERRUN, SET
   SHAD_REMOTE_MAX_EVENTS BETWEEN 100 TO 400
   IN SHADPARM.INI
   ```

   This typically occurs if there are many empty transactions (transactions with no associated database modifications) sent to the NonStop Consumer. If this occurs, set the SHAD_REMOTE_MAX_EVENTS parameter in SHADPARM.INI to between 100 to 400 events, e.g.:

   ```
   SHAD_REMOTE_MAX_EVENTS=200
   ```

9. On Windows, when configured to “roll on size”, the DOC will not always roll at the correct size and can exceed the specified roll size significantly. This is related to a file system size reporting issue. Hence, you may need to set the roll size tens of MB’s less than you otherwise would.

10. Due to limitations in configuration record sizes, DNS names cannot be longer than 20 characters. If the DNS name exceeds 20 characters, use the dotted IP address instead.

11. When replicating from the NonStop in a multi-ported DOC environment, the SUSPENDUPD/RESUMEUPD command ***cannot be used*** if the SHAD_TRANS_EXPECTED_ENDS parameter is set to a value greater than 1. The SHAD_TRANS_EXPECTED_ENDS parameter is not required in configurations where the NonStop Shadowbase is sending to a single multi-ported DOC Writer. If the configuration includes multiple DOC Writers and Direct Writers, SHAD_TRANS_EXPECTED_ENDS is a required parameter.

    Only one commit is sent for a SUSPENDUPD command. If SHAD_TRANS_EXPECTED_ENDS is greater than 1, the DOC Writer will leave the SUSPENDUPD command in an uncommitted state, preventing the DOC files from being removed by the DOC cleaner.

12. Due to shell incompatibilities, the configuration scripts for HPE Shadowbase Log Source for Oracle do not work on Solaris 10 and AIX. If you need to configure HPE Shadowbase Log Source for Oracle for Solaris 10 or AIX, please contact Shadowbase Support.

13. If the SQL Server Native Client version 10.0 is installed on Windows 2008, the TRS and Direct Writer will fail as they cannot load the SQL Server client DLLs. This issue is still under investigation. Two work-arounds exist: either install a different version of the SQL Server Native Client, such as version 11.0, or use set up an ODBC DSN (data source name) and use the generic Shadowbase ODBC process.

14. We have seen the Oracle database driver process's memory use grow over time for certain configurations of HPE Shadowbase Log Source for Oracle when using Oracle 11 and 12 databases. This issue seems to be tied to certain Oracle patch levels. If this occurs in your configuration, report the instance to Shadowbase Support; you may need to apply some Oracle patches. Note that there is a workaround using the SHAD_OPCOLLECT_RECYCLE_CONNECTION and the SHAD_OPCOLLECT_QUERY_EXECUTION_LIMIT parameters. Enabling the SHAD_OPCOLLECT_RECYCLE_CONNECTION will cause the database connection to Oracle to be closed and reopened after SHAD_OPCOLLECT_QUERY_EXECUTION_LIMIT queries. This also causes the Oracle database driver process to restart, releasing the memory.

    Setting SHAD_OPCOLLECT_RECYCLE_CONNECTION will enable the connection recycling. SHAD_OPCOLLECT_QUERY_EXECUTION_LIMIT defaults to 250 queries before the connection is recycled, you can reduce it (resulting in less memory usage but poorer throughput) or increase it (resulting in better throughput but more memory usage) as your needs dictate.

15. Shadowbase connects to Microsoft SQL Server databases using the default setting for the AutoTranslate ODBC parameter, which is on. If Shadowbase is running on a different server from the SQL Server database and the two servers are using different ANSI code pages, character data stored in char, varchar, and text fields will automatically be converted by the ODBC driver. The ODBC driver performs the conversion by converting the data to UNICODE based upon the Shadowbase server's ANSI code page and then back to character fields from UNICODE using the SQL Server database's code page.

    If you want to disable the conversion, you need to setup an ODBC database source connection (DSN) with AutoTranslate configured off and to configure Shadowbase to use the DSN instead of connecting directly.

16. On Solaris servers, you may have to specify the LD_LIBRARY_PATH in your environment to have Shadowbase start correctly. If Shadowbase fails to start and logs a message in the error log similar to:

```
2015-05-06 14:41:45 -[14406] Critical Error: Cannot load library
```

```
(libOCIEIDLL.so) - ld.so.1: shadowbase: fatal: libclntsh.so.11.1: open
failed: No such file or directory
```

you will need to set the LD_LIBRAY_PATH environmental variable to either $SHAD_BASE/lib or $ORACLE_HOME/lib.

17. HPE Shadowbase Log Source for Oracle does not detect schema changes to replicated tables and may use an incorrect schema if the table is updated while Shadowbase is running. Added columns may not be picked up, and deleted columns may result in Shadowbase stopping.

18. When DBS Mapping is enabled, the TFS is not sending fully qualified datetime data to the NonStop for UPDATE statements, which can result in conversion errors on the NonStop.

19. Internal testing uncovered a number of limitations on the size of columns, tables, and statements:

    a. There is a limit to the size of a row in the DOC database which limits the size of statements (for EI Docs), cached statements (for cached DOCS) and statement data (for cached DOCS) to approximately 56K bytes.

    b. Table names are limited to 80 characters.

    c. Column names are limited to 74 characters.

20. There is an issue configuring HPE Shadowbase Log Source for Oracle collection using the shadconfig configuration script if you enter an incorrect username/name/Oracle SID combination when specifying the connection information using Option 1. If you do not exit the script but instead correct the connection information, subsequent entries made to the SHAD_USERS_INCLUDE, SHAD_TABLES_INCLUDE, SHAD_USERS_EXCLUDE, and SHAD_TABLES_EXCLUDE table entries will include a trailing space for the name, which will cause the data selection criteria to be incorrect.

    If you suspect that the criteria may be incorrect, you can issue the following select command against the appropriate table:

```
SELECT * FROM SHAD_TABLES_INCLUDE
WHERE TABLE_NAME LIKE '% ';

SELECT * FROM SHAD_TABLES_EXCLUDE
WHERE TABLE_NAME LIKE '% ';

SELECT * FROM SHAD_USERS_INCLUDE
WHERE USERNAME LIKE '% ';

SELECT * FROM SHAD_USERS_EXCLUDE
WHERE USERNAME LIKE '% ';
```

The select should return no rows. You can correct the issue using the following update command against the appropriate table:

```
UPDATE SHAD_TABLES_INCLUDE
SET TABLE_NAME=RTRIM(TABLE_NAME)
WHERE TABLE_NAME LIKE '% ';

UPDATE SHAD_TABLES_EXCLUDE
SET TABLE_NAME=RTRIM(TABLE_NAME)
WHERE TABLE_NAME LIKE '% ';

UPDATE SHAD_USERS_INCLUDE
SET USERNAME=RTRIM(USERNAME)
WHERE USERNAME LIKE '% ';

UPDATE SHAD_USERS_EXCLUDE
SET USERNAME=RTRIM(USERNAME)
WHERE USERNAME LIKE '% ';
```

21. HPE Shadowbase Log Source for Oracle may incorrectly stop with the following messages:

```
*** Potential Data Loss Detected
*** SQL_REDO for last event processed in previous query does not match
current query
***  SCN [<scn>]
***  RBASQN [<rbasqn>]
***  RBABLK [<rbablk>]
***  RBABYTE [<byte>]
***  SSN [<ssn>]
***  Previous SQL_REDO[<sql statement>;]
***  Current  SQL_REDO[<sql statement>;]
***  SHAD_OPCOLLECT_RESET_SCN is DISABLED, stopping
***  Enable SHAD_OPCOLLECT_RESET_SCN to continue processing prior to restart
Performing shutdown
```

    This can occur if your text columns have binary 0 (ASCII NUL) characters in them. To continue past this error, set the SHAD_OPCOLLECT_RESET_SCN=1 parameter in the SHADPARM.ini file.

22. Special considerations are required when using a TFS as a relay server. The figure below shows a sample configuration where a TFS is used as relay server. Events are replicated from the source Oracle database into a DOC using the log source collector. They are read from the DOC and send to a second DOC via the relay TFS and DOCW. Finally, they were read from the relay DOC and applied in SQL/MX using another TFS and a Direct Writer.

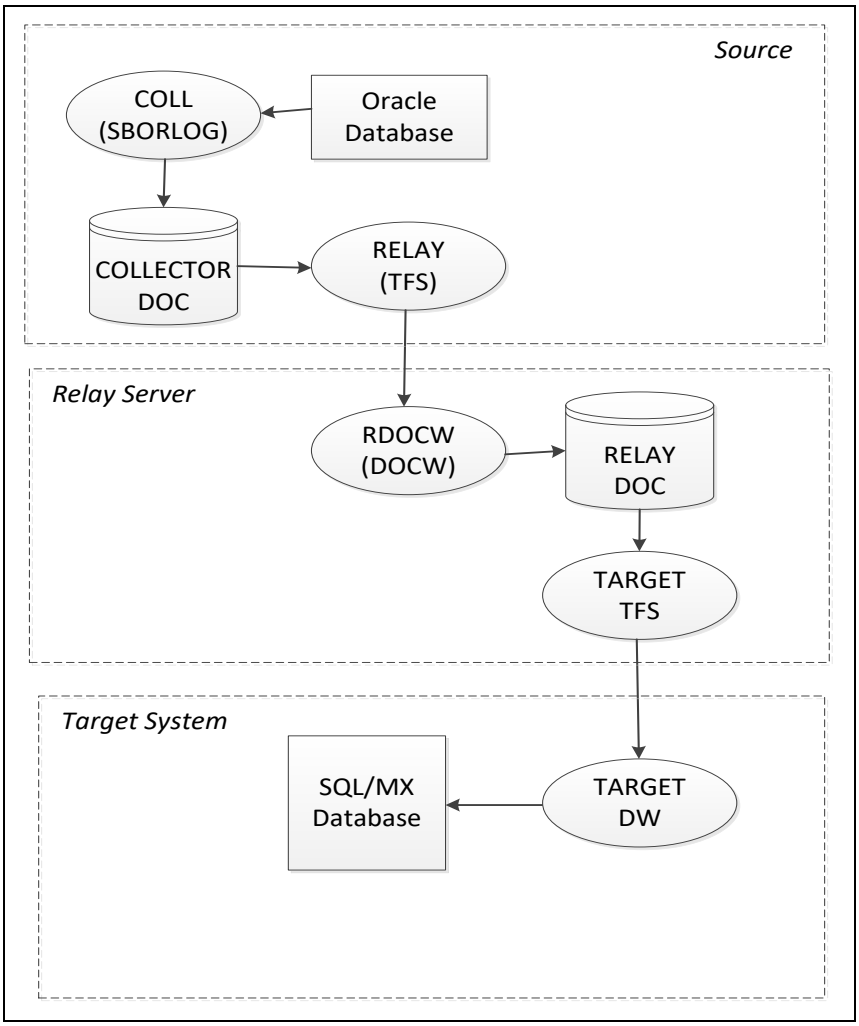

Figure 1 - Sample Relay Server Configuration

In a configuration such as this one, the relay TFS restarts must be coordinated with the collector restarts. Whenever the collector is restarted, the relay TFS should also be restarted to prevent duplicate data from being replayed.

23. The new format DOC, which eliminates the TRANS DOC, is not compatible with Pulse processing. If you are planning on using Pulse processing, continue to use the old format DOC with trans processing enabled.

24. When replicating to MySQL, the ODBC library looks in /tmp by default for mysql.sock. Set the socket attribute in the odbc.ini section to the desired location, such as in the example below:

```
[myodbc64]
Description     = ODBC for MySQL
Driver          = /usr/lib64/libmyodbc5.so
Setup           = /usr/lib64/libodbcmyS.so
FileUsage       = 1
socket          = /var/lib/mysql/mysql.sock
```

Alternatively, make sure that the mysql.sock file is either located in, or is symlinked to /tmp/mysql.sock. Due to limitations of UnixODBC, the ODBC library that HPE Shadowbase uses to interface with ODBC databases on Linux systems, the socket file must reside in the /tmp directory. You can find out the current location of the mysql.sock file by running `mysqladmin -p -u <user-name> variables` and finding the output under the 'socket' column. The socket file can then be symlinked to the /tmp directory with some version of the command `ln –s /var/lib/mysql/mysql.sock /tmp/mysql.sock`. The symlink may have to be recreated every time the system is rebooted or whenever the system cleans out the /tmp directory.

25. There is an issue in Shadowbase Doc Writer and Direct Writer support for SAP ASE target, also known as Sybase, such that users should discontinue use of SQL92 as the source SQL type and instead switch to SQLSBV1.

26. There is an issue with the Shadowbase Transaction Replay Server where a single TRS can only apply to a single database if the target database is MySQL. If you need to replicate into multiple MySQL databases you will need multiple TRS objects.

27. We have reports of a potential CTREE UpdateTrans error 101 occurring in a DOC Writer which causes a disconnection with the upstream consumer. In this situation, the Consumer and DOC Writer may continue to connect and disconnect. In this case, the workaround is to stop and restart the DOC Writer to get replication to continue as usual.

28. There is an issue on Windows 2016 where a TRS stopping can have Critical Errors that may cause problems upon the TRS restarting. This is caused by an issue between Shadowbase and Windows Services. The workaround for this is to set the Windows Service handling the TRS to restart itself after an error which will bring the TRS back up even if there were critical errors

29. There is an issue relating to the maximum statement cache depth of OSB Log Source objects where each Log Source Collector is limited to 1000 unique statements. A workaround to collect more than 1000 statements is to configure multiple Log Source collectors to keep each one under 1000 statements.

30. There are certain cases where when running a TRS or other HPE Shadowbase object from a service on Microsoft Windows Server 2019 or 2022, the errlog will be empty, however the object will be running fine. This problem is intermittent and nondeterministic. Restarting the object will correct the problem and cause the errlog to be recorded as expected.

31. There is an issue attempting to configure a DB2 source collector on a Windows system to connect to an MQ Queue Manager on a remote server. This issue prevents the collector from establishing a connection to the Queue Manager. The workaround is to configure the DB2 source collector locally on the server where the MQ Queue Manager is configured.

32. Replication to DB2 columns with data types of BINARY, VARBINARY, and BLOB is not currently supported.

33. There is a constraint for Microsoft SQL Server tables called IGNORE_DUP_KEY that causes the database to suppress duplicate primary key errors. Avoid setting this constraint to `ON` for target replication to SQL Server as it can cause target replication objects to fail without a descriptive error message.

34. There is an issue with updates when replicating Oracle data of type LONG RAW. Inserts using that data work correctly but updating any data will result in any columns of LONG RAW type in that row to be cleared.

35. There remains an issue with the replication of MySQL columns with the type of BLOB. It is recommended that the target column type be changed to TEXT or another supported target character data type.

36. There is an issue with the replication of PostgreSQL columns with the type of BYTEA. It is recommended that an alternative target column type be chosen instead.

37. Pulse events do not currently function using DOC Blocking Level 2 due to a fault when inserting the event which causes the Pulse to be ignored. The workaround for this is to use DOC Blocking Level 1 or no DOC Blocking when Pulse events are required.

38. The Microsoft SQL Server 'SQL_VARIANT' data type has limited support and has not been fully tested.

39. There is an issue with the Visual Studio project files that are intended to be used for building custom user exits on Windows. They will load in Visual Studio properly but for building the dll files you will need to run nmake using the CustomUE-MakeFile.mak file from an Administrator Command Prompt. Using Visual Studio to build will create the files but result in an error at runtime stating that the entry point for USRXDLLINIT is missing.

40. Daily license checks will not be performed on DOCs waiting for a connection. If the connection is not established until after the check would have occurred it will be performed at the time of the connection. If a DOC is left in the waiting state this could cause the license to expire without any prior warning upon connection.

41. When replicating to Microsoft SQL Server with the data types BINARY, VARBINARY, and IMAGE, the data must be entered as hex values with no leading '0x'. A user exit will be needed to modify the values for this to work. Microsoft has also announced that the IMAGE data type will be deprecated in a future version of SQL Server and has begun encouraging using VARBINARY to store large data instead.

42. DBS mapping does not allow for $VOL.SUBVOL.TABLE syntax in the [DBS+] command. This may be a problem when replicating MP -> MP with syskey columns via ODBC/MX if the fully qualified $VOL.SUBVOL.TABLE table name is specified in the DBS PATHNAME. If DBS mapping is required, because an MP table has an MX reserved word as a column name, the TRS will fail.

    The workaround for this is to create a SQLMP ALIAS in the MX target environment to point to the target MP table and reconfigure the source NSB environment to use this alias value in the PATHNAME instead of the Guardian fully-qualified filename. The data will need to be re-collected from the source side and sent to the target OSS side again after the change is made, but that will allow the shaddbs file to have a normal table name in the [DBS+] line which will allow everything to work properly.

43. Replication to SQL/MX target table columns with data types of BINARY or VARBINARY are not currently supported.

44. Replication of fractional values in to MySQL TIME column types fails with a "Fractional truncation" error. This is due to an issue with certain versions of the ODBC MySQL driver. DATETIME and TIMESTAMP columns can handle up to 6 digits of precision appropriately.

    In addition, other MySQL ODBC driver versions cause Shadowbase to fail to replicate values into DATETIME target columns with a "Date overflow" error. For these issues and other similar MySQL issues, be sure to use the latest available ODBC driver.

45. Modifying the COLUMNS environment variable may cause the SBMON STATUS command to report objects as stopped incorrectly.

46. DB2 source events with BOOLEAN type are not processed by asnqcap and therefore cannot be replicated by Shadowbase.

47. There is an issue with the SBMON LIST command output for Direct Writers where the count of inserts, updates, deletes, and total transactions does not always update in sync with when transactions are processed. Replicating additional transactions will cause previously uncounted events and transactions to be added to this count.

48. When using SQLServer trigger source on Windows, the SHAD_BASE environment variable must be set to the current SHAD_BASE location. If it points to a previous install, the OPCOL EXTPR process will not start correctly. When set up correctly, non-trigger source environments use the proper SHAD_BASE whether it is set as an environment variable or not.

49. For array-based processing, DB tracing does not work on individual statements but rather for the full array.

50. SHAD_SKIP_DOC_OPEN_ERRORS does not work in environments with No Trans Docs. In order to get similar functionality, the TRS restart point will need to be advanced past the errors that need to be skipped.

51. When replicating batches of inserts that contain a mixture of successful and failed inserts, Array-based TRS and DW objects can generate primary key violation errors due to all inserts in the batch being replayed, including previously successful inserts.

52. TRS and DW objects configured by SHAD_LATENCY_MONITORING_THRESHOLD to log replication latency when it reaches the level of the configured threshold are programmed to also log a message when the latency falls below the configured latency threshold, but currently this feature is only present in the TRS.

53. The use of SHAD_TX_NAME is not supported for ODBC TRS and Direct Writer objects. When configuring an Oracle to Oracle bidirectional replication environment, it is recommended to use Oracle OCI TRS or DW objects (sborwrtc, sboracle), as they support the important ping-pong avoidance parameter SHAD_TX_NAME.

54. The Oracle datatype INTERVAL is currently not supported for ODBC Transaction Replay Servers and Direct Writers. It is recommended to use Oracle TRS and DW target objects over ODBC TRS and DW objects when replicating to an Oracle target database.

55. The use of Array Processing TRS and DW objects as part of a SQL Server Bidirectional replication configuration is currently unsupported. Use the standard ODBC or MSSQL TRS and DW objects instead.

56. The use of triggers that execute DML operations when replication DML operations are executed against a target table is not currently supported with Array Processing TRS and DW objects. Use ODBC or MSSQL TRS or DW objects if a configuration requires the use of data manipulating triggers on target tables.

57. There remains a bug in Array Processing objects' handling of unique constraint violations, including primary key violations, relating to the SHAD_SQL_ERROR_... parameters. In order to properly log and skip unique constraint violations that occur during batch insertion of data with Array Processing objects, one should rely on the SHAD_PKDUP… family of parameters rather than setting SHAD_SQL_ERROR_.... Setting SHAD_SQL_ERROR_... will cause Shadowbase Array Processing objects to log and then stop replication in this release.

58. The shadparm parameter SHAD_SKIP_INF_ERROR is currently not supported for Array Processing objects.

59. Array Processing replication is only supported for SQL Server targets on Windows.

60. SHADTRS inserts are processed using the traditional SQLExecute method rather than being batched together using Array Processing. Depending on the structure of the transactions being replicated, enabling SHADTRS processing may throttle the object to run as slowly as if it were a normal non-Array object. SHADTRS processing should be disabled in this scenario to achieve higher performance.

61. SHAD_SQL_ERROR_INVALID_NUM and SHAD_SQL_ERROR_MAPTO processing are not supported for Array Processing.

62. The WHERE clause used for all update and delete events must always include the same columns in every event replicated for Array Processing. Normally, this should match the primary key of the target table. If any of the target tables have no primary key, or if the WHERE clause in the updates and deletes does not match the primary key for any reason, this can be overridden by using the SHAD_KEY_COLUMNS_FILE parameter.

63. For certain SQL database errors during Array Processing, the individual event that caused the error cannot be determined, and therefore it is not logged to the errlog along with the failure message. In situations like this, it may be desirable to temporarily switch the object over from Array Processing to normal processing and start it again. The object will fail with the same error again, but this time it will provide the user with standard error details for the purpose of troubleshooting.

64. An explicit ODBC DSN for the target database must be configured before starting an Array Processing object.

65. If SHAD_DB_TRACING is enabled for an Array object, no information about the individual events applied to the target database will be logged.

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

To start the installation, use the following instructions:

1. Copy the .zip file containing the software release onto the server where Shadowbase will be installed. If using FTP to move the files onto the server, be sure to use "binary" mode.
2. Unzip the .zip file on the server.
   a. Alternatively, unzip the .zip file on a PC first, then binary FTP all the contents to the server where Shadowbase will be installed.
3. Run the `run_first` script on the Shadowbase server.
   a. On Linux/Unix/OSS, run `. run_first` from a command prompt.
   b. On Windows, run `run_first.bat` from a command prompt.
4. The script will prompt for an acknowledgement of the license format change in this version. Say `Y` to the prompt to continue. If anything other than `Y` is entered, the run_first script will need to be run again.

After the license format change has been acknowledged, the installation can proceed using the normal instructions. To complete the installation and configuration of Shadowbase please see the HPE Shadowbase Other Servers Target Installation and Configuration Manual (for replication target installations) and the HPE Shadowbase Other Servers Source Installation and Configuration Manual (for replication source installations).