

# SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	HPE Shadowbase for Other Servers
Release:	Gravic Version 6.700 T1123-AAI (SB REPL/OSS) T1124-AAE (SB REPL/HP UX) T1125-AAG (SB REPL/SUSE LINUX) T1126-AAF (SB REPL/RHEL LINUX) T1127-AAE (SB REPL/SOLARIS X86) T1128-AAE (SB REPL/SOLARIS) T1129-AAF (SB REPL/MS WINDOWS) T1130-AAE (SB REPL/AIX)
Date:	January 5, 2022
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2022 ( <a href="http://www.gravic.com">www.gravic.com</a> )
File Name:	IPM6700_other_servers.pdf

NOTE: This softdoc covers new features and corrected problems for HPE Shadowbase for Other Servers, Version 6.700. 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

<b>Table of Contents</b>	<b>1</b>
<b>Disclaimer</b>	<b>2</b>
<b>Note for TCDs</b>	<b>3</b>
<b>Mature Features and Deprecated Features</b>	<b>4</b>
<b>Supported Database and Operating System Versions</b>	<b>6</b>
<b>Important Note about Licenses to run Shadowbase</b>	<b>7</b>
<b>Special Notes for Version 6.700</b>	<b>8</b>
<b>New Features, Problems Corrected, and Other Changes in Version 6.700 (since Version 6.600)</b>	<b>10</b>
<b>New Features in Version 6.700</b>	<b>10</b>
<b>Problems Corrected in Version 6.700</b>	<b>12</b>
<b>New &amp; Modified SBMON Commands</b>	<b>14</b>
<b>New &amp; Modified shadparm.ini Parameters</b>	<b>15</b>
SHAD_ABORT_ON_FATAL_ERROR	15
SHAD_OPCOLLECT_NCHAR_TARGET_TYPE	15
SHAD_DOC_ALLOW_USRX	16
SHAD_ENABLE_SEQUENCE_OBJECT_REPLICATION	16
SHAD_OPCOLLECT_COLUMN_WRAP	17
SHAD_TFS_TX_ID_SWAP	17
SHAD_OPCOLLECT_DELAY_LEVEL	18
SHAD_OPCOLLECT_SLEEP_TIME	19
SHAD_OPCOLLECT_EXEC_COUNT	19
SHAD_MAX_CURSORS	20
SHAD_DOC_BLOCKING	20
SHAD_INVALID_RESTART_SHUTDOWN	21
SHAD_IGNORE_MISMATCH	22
SHAD_ODBC_SQL_C_TYPE	22
SHAD_SQL_FORCE_INF	23
SHAD_SYSKEY_INDEX_FIX_LEVEL	23
SHAD_OVERRIDE_SYSKEY	24
SHAD_SKIP_INF_ERROR	25
SHAD_SQL_ERROR_INVALID_NUM_SKIP_LOGGING	26
SHAD_SQL_ERROR_INVALID_NUM_SKIP	26
SHAD_ODBC_CONCURRENCY	27
SHAD_OPCOLLECT_TIMESTAMP_CONVERSION_LEVEL	27
<b>New &amp; Modified User Messages</b>	<b>29</b>
<b>Known Problems Remaining in Version 6.700</b>	<b>30</b>
<b>Installation Instructions</b>	<b>40</b>

## 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](http://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
<b>Oracle Trigger Source</b>	Oracle trigger source collectors (OPCOL) are being replaced by the Oracle log source collector (SBORLOG).	OSB 6.300 May, 2017

**Table 2 - Deprecated Features in HPE Shadowbase for Other Servers**

Feature	Description	Deprecated Release	Last Release Date
---------	-------------	--------------------	-------------------

<b>Execute Immediate Objects</b>	<p>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:</p> <ol style="list-style-type: none"> <li>1. sbocons</li> <li>2. sbmspre</li> <li>3. sborapre</li> <li>4. sbsypre</li> <li>5. sbigdir</li> <li>6. sbmswrt</li> <li>7. sborcol</li> <li>8. sbsywrt</li> <li>9. sbigwrt</li> <li>10. sborcol</li> </ol> <p>If you are using any of these executable programs (configured by SBMON), they will be retired on the date listed above.</p>	OSB 6.300 May, 2017	May, 2019
<b>SBDOCRD</b>	<p>The functionality in SBDOCRD has been replaced by SBFILE. SBDOCRD only supports execute immediate DOCs, which are also deprecated.</p> <p>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.</p>	OSB 6.300 May, 2017	May, 2019
<b>SQL92</b>	<p>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.</p>	OSB 6.600 April, 2020	April, 2022
<b>UserExit Eclipse projects</b>	<p>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.</p>	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](http://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 an earlier version of Shadowbase for Other Servers software to version 6.700, 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](mailto:license.manager@hpe.com)) or your local sales representative to request a new license file. **DO NOT INSTALL** Shadowbase for Other Servers software version 6.700 until a new license file has been obtained.



## Special Notes for Version 6.700

1. Version 6.700 is a “general availability” (non-restricted TCF) release. The previous non-restricted TCF release is version 6.600, which corresponds to HPE version AAC for HP-UX and IBM AIX, Oracle Solaris, and Oracle Solaris x86 systems, version AAD for Microsoft Windows Server and Red Hat Enterprise Linux (RHEL) systems, version AAE for SUSE Linux Enterprise Server (SLES) systems, and version AAF for HPE OSS systems.
2. This document describes the changes since Version 6.600. If you are upgrading from a version of Shadowbase for Other Servers prior to that, please first see the version 6.600 software release document, IPM6600\_other\_servers, for a list of the changes in that version.
3. **Due to licensing changes introduced in this version, existing installations of Shadowbase prior to this version 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](mailto: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 perform some simple verification that the license file is valid for this version of HPE Shadowbase.
5. **For Shadowbase on OSS, a special file called “shadtype” will 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 new NonStop system types that are newer than the Shadowbase release to ensure that the license is validated 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. Version 6.700 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 psABI 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. Sbininstall 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 3<sup>rd</sup> party vendored 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.

## **New Features, Problems Corrected, and Other Changes in Version 6.700 (since Version 6.600)**

### ***New Features in Version 6.700***

This section provides a summary of the features added to HPE Shadowbase for Other Servers products since the previous general availability release (V6.600).

1. Support for collection of the Oracle types `TIMESTAMP WITH TIMEZONE` and `TIMESTAMP WITH LOCAL TIMEZONE` has been added to HPE Shadowbase Log Source for Oracle. Note that for `TIMESTAMP WITH TIMEZONE`, there are limitations for target replication. Due to the format this type is collected in, `TIMESTAMP WITH TIMEZONE` can only be replicated into Oracle target columns of the same type without enabling the `SHAD_OPCOLLECT_TIMESTAMP_CONVERSION_LEVEL` parameter or changing the NLS settings in the Oracle source database. `TIMESTAMP WITH LOCAL TIMEZONE` is collected in a more standard SQL `DATETIME` format and can therefore be replicated into most `DATETIME/TIMESTAMP` target formats without any extra parsing. The most reliable way to collect this data in a format compatible with other database targets is to alter the `NLS_TIMESTAMP_TZ_FORMAT` Oracle NLS parameter to match your desired format.
2. Support for the use of User Exits that can modify data with HPE Shadowbase Log Source for Oracle has been added. This can allow for column filtering on replication. To use user exit functionality, `SHAD_DOC_ALLOW_USRX` parameter must be set to 1 and DOC blocking must be disabled for the OPCOL object.
3. `SQLSBV1` has been added as a statement type for TRS and OPCOL in order to facilitate the deprecation of `SQL92`.
4. Support for replicating Oracle column names that include spaces has been added with the use of the `SHAD_OPCOLLECT_COLUMN_WRAP` parameter to set the delimiters required by the target.
5. Support for PostgreSQL target replication via ODBC has been added for Windows and Red Hat Linux clients.
6. Support for MySQL target replication via ODBC has been added for Windows clients.
7. Added target side support for the replication of SQL/MX Sequence Objects from one SQL/MX table to another SQL/MX table. Note that to use this feature, a version of T1122(SB REPL/Guardian) supporting source side SQL/MX Sequence Object replication is also needed.

8. Added the Visual Studio project for User Exit development back into the Windows release of HPE Shadowbase.
9. Added support for updated versions of the MySQL ODBC driver on Windows and Red Hat Linux.

## ***Problems Corrected in Version 6.700***

This section provides a summary of the problems corrected in HPE Shadowbase for Other Servers products since the previous general availability release (V6.600).

1. DBS mapping under certain circumstances would include quote characters in data inappropriately. This has been fixed so that they will not be passed in to the table.
2. There was an issue when editing a DB2 collector configuration where the prompt for IBM MQ username did not register the previously entered username and was followed by two consecutive prompts for the IBM MQ Password. This has been corrected so that the previously set username is listed as the default value and the password prompt will only print once.
3. Previously when running an Oracle Log Source environment that was also being archived by DataGuard, the CONTINUOUS\_BOTH mining mode would only replicate the data in the archived redo logs. This was due to a deprecated CONTINUOUS\_MINE Oracle option which HPE Shadowbase no longer relies on so now CONTINUOUS\_BOTH mining will replicate both archived and online logs.
4. The handling of NULL values in DBS mapping has been corrected so that when replacing NULL values using DBS mapping, the data entered will no longer be inappropriately wrapped by single quotes
5. When editing an executable and changing it to an sborwrtc type, the previous DSN value would be displayed despite it not being used by sborwrtc. This has been fixed to fully remove the DSN.
6. When forwarding from one DOC to another using a TFS, the TX ID would be inadvertently flipped. Now both DOCs will show the same TX ID.
7. In OSS Direct Writers, sockets would stay open but not be used on the target if the source side consumer fails. This would become a growing concern if the failures happened frequently as the zombie sockets would remain open until the Direct Writer was restarted. This was fixed so that sockets will close when the connection is broken
8. SHAD\_SHADTRS\_DUP\_TX\_LOGGING was behaving differently than described and was also not being logged on startup. Now the value it is set to will be logged properly and it will behave the way the documentation describes.
9. SHAD\_DISPLAY\_INF logic has been fixed to properly show when an update is modified to an insert for the insert not found processing.
10. The SBPUTCOLUMNOPEN function has been fixed so that setting the length value to -1 in OSB User Exits will now actually result in a NULL value in the respective column.

11. Log Source for Oracle has been improved to where RAW data type columns can be collected and replicated to NSB and OSB target databases without a previously present error in replicating updates and deletes.
12. The SBMON LIST command has been enhanced when being run against OPCOL and DOC objects to more accurately reflect the progress of TRS and TFS objects against the relevant DOC sequences.
13. There was an issue with ODBC drivers for Microsoft SQL Server where newer versions were not being recognized, but HPE Shadowbase should now attempt to use any of the official Microsoft ODBC drivers.
14. An issue where INVALID\_NUM\_PROCESSING would not properly work for MySQL targets has been corrected.
15. Resolved an issue where DB2 and SAP HANA Boolean types were being treated as numeric values in INVALID NUM PROCESSING.

## ***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 (V6.600).

1. When entering User/Schema name for Oracle in sbmon, the name is no longer case sensitive to match how Oracle itself treats the names. If case sensitivity is required, add double quotes around the value.
2. When configuring a DOC or Collector, the default for Transaction Processing has been changed from Yes to No for all objects except SQL Server collectors and DB2 collectors.

## 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 (V6.600).

### SHAD\_ABORT\_ON\_FATAL\_ERROR

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_ABORT_ON_FATAL_ERROR=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	0 (OFF)		
<b>Other Servers</b>	All		
<b>Valid Settings</b>	0: Shadowbase objects will shutdown normally in the case of a fatal error. 1: Shadowbase objects will raise SIGABRT after fatal errors in an attempt to generate a corefile 2: Shadowbase objects will raise SIGABRT after any exit call in an attempt to generate a corefile		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

#### Description

This parameter controls how a Shadowbase process responds to certain error and exit scenarios. It is intended to cause a coredump in order to obtain a core file that can be used by Shadowbase Development to produce a solution. Setting this parameter to 1 indicates that Shadowbase objects should raise the signal SIGABRT in an attempt to generate a coredump file on fatal errors. Setting this parameter to 2 does the same for all exit calls. The default of 0 (off) will cause objects to exit normally.

#### Notes

Setting this parameter to 2 will cause objects to coredump whenever they are stopped in any way other than from a normal “stop” command in *sbmon*.

### SHAD\_OPCOLLECT\_NCHAR\_TARGET\_TYPE

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCOLLECT_NCHAR_TARGET_TYPE=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	0		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0: Oracle Target 1: Nonstop Target		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

#### Description

This parameter controls how a Shadowbase Log Source Collector for Oracle object processes and stores Oracle NCHAR type columns. Setting the parameter to 0 for Oracle Target will cause the Collector to wrap NCHAR column substitution variables with the Oracle function UNISTR() for storage in the statement cache. This allows the collected data to be replicated into an Oracle target without any additional data transformation.



Setting the parameter to 1 for Nonstop Target will cause the Collector to convert the Unicode codepoints collected from Oracle NCHAR data type columns into its binary representation. This allows the collected data to be replicated into SQL/MX or SQL/MP targets.

## Notes

At this time, replication to targets other than Oracle, SQL/MX, and SQL/MP is not supported

## SHAD\_DOC\_ALLOW\_USRX

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_DOC_ALLOW_USRX=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	0 (OFF)		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0: USRX disabled for Shadowbase Log Source for Oracle 1: Allow user exits for SBLs for Oracle		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

## Description

This parameter controls whether or not user exits will be used for HPE Shadowbase Log Source Collector for Oracle. The default of 0 maintains the previous functionality but setting it to 1 allows for modification of input data during replication.

## Notes

Setting this parameter to anything other than 0 or 1 will cause the object to error and exit. DOC blocking must be disabled to enable User Exit functionality for the Collector by setting SHAD\_DOC\_BLOCKING to 0.

## SHAD\_ENABLE\_SEQUENCE\_OBJECT\_REPLICATION

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_ENABLE_SEQUENCE_OBJECT_REPLICATION=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	0 (OFF)		
<b>Other Servers</b>	TRS   DW (OSS)		
<b>Valid Settings</b>	0: Sequence object replication disabled for TRS and DW 1: Sequence object replication is enabled for TRS and DW		
<b>Basic/Advanced</b>	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 6.700 SOFTDOC for HPE NonStop Shadowbase for more information on enabling collection of Sequence Object events on the source side.

## SHAD\_OPCOLLECT\_COLUMN\_WRAP

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCOLLECT_COLUMN_WRAP=<char><char>
<b>Last Change:</b>	--		
<b>Default Value</b>	No Wrap Characters		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	Single Character: The character will be added at the beginning and end of each column name Two Characters: The first character will be added before the column name and the second will be added to the end		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls how column names from Oracle should be wrapped for the target. It is only necessary with column names that need to be delimited such as multiple word names with spaces.

### Notes

If more than two characters are provided only the first two will be used. Spaces should not be used as they will read as the character to wrap column names with.

## SHAD\_TFS\_TX\_ID\_SWAP

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_TFS_TX_ID_SWAP=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	1(ON)		
<b>Other Servers</b>	TFS		
<b>Valid Settings</b>	0: Behavior from v6600 and prior where transaction ids may be reversed when using a TFS replicating to a DOC 1: TFS will fix the transaction id swap so that both DOCs match		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls how a TFS with both DOC source and target will handle transaction ids. Prior versions of HPE Shadowbase would report flipped transaction ids in the two DOCs but with this parameter set to 1 they will match.

### Notes

A TFS with a Guardian target will not be impacted by this parameter.

**SHAD\_OPCCOLLECT\_DELAY\_LEVEL**

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCCOLLECT_DELAY_LEVEL=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	1		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0: No delay is added into the OPCollect cycle 1: A sleep call will be made in the OPCollect cycle based on related parameter settings 2: A call to the Win32 api function SwitchToThread will be made in the OPCollect cycle 3: A call to the Win32 api function SwitchToThread as well as a sleep call will be made in the OPCollect cycle 4: A call to the Win32 api function SwitchToThread will be made in the OPCollect cycle. Additionally, diagnostic information will be printed to the OPCOL log.		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

**Description**

This parameter controls whether or not a sleep or thread yield is issued in the MSSQLCollector OPCollect function. This functionality has been added to avoid the issue of MSSQL OPCOL objects failing to roll the DOC in situations where large transactions have been held open for extended periods of time. The default value 1, will result in a sleep call being made every SHAD\_OPCCOLLECT\_EXEC\_COUNT calls to OPCollect for a length of SHAD\_OPCCOLLECT\_SLEEP\_TIME milliseconds. A value of 2 will call SwitchToThread instead of the existing Shadowbase sleep, and is uninfluenced by SHAD\_OPCCOLLECT\_SLEEP\_TIME. A value of 3 will do both actions of options 1 and 2 by calling both SwitchToThread and sleep. A value of 4 will call SwitchToThread and print out diagnostic information intended to aid with the debugging of a Shadowbase environment.

**Notes**

This parameter only applies to trigger-source OPCOL and not log-source.

**Related Parameters**

SHAD\_OPCCOLLECT\_SLEEP\_TIME  
SHAD\_OPCCOLLECT\_EXEC\_COUNT

## SHAD\_OPCOLLECT\_SLEEP\_TIME

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCOLLECT_SLEEP_TIME=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	1000 ms		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0 – 4294967295 (UINT_MAX)		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls the length of the sleep call to be made in the OPCollect function of MSSQLCollector. If SHAD\_OPCOLLECT\_DELAY\_LEVEL is set to 1 or 3, this parameter will cause a sleep call for the number of milliseconds provided as the parameter value.

### Notes

This parameter only applies to trigger-source OPCOL and not log-source.

### Related Parameters

SHAD\_OPCOLLECT\_DELAY\_LEVEL  
SHAD\_OPCOLLECT\_EXEC\_COUNT

## SHAD\_OPCOLLECT\_EXEC\_COUNT

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCOLLECT_EXEC_COUNT=<num>
<b>Last Change:</b>	--		
<b>Default Value</b>	100,000		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0 – 4294967295 (UINT_MAX)		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls how many calls are made to OPCollect before either a sleep or a yield thread call is made, depending on the settings of related parameters. OPCollect will be executed the number of times specified by the value of this parameter before sleeping or yielding the thread according to SHAD\_OPCOLLECT\_DELAY\_LEVEL.

### Notes

This parameter will be disabled if SHAD\_OPCOLLECT\_DELAY\_LEVEL is set to 0. This Parameter only applies to SQL Server trigger-source OPCOL and not log-source collectors .

### Related Parameters

SHAD\_OPCOLLECT\_DELAY\_LEVEL  
SHAD\_OPCOLLECT\_SLEEP\_TIME

## SHAD\_MAX\_CURSORS

<b>Initial Version:</b>	V2.800	<b>Syntax:</b>	SHAD_MAX_CURSORS=<num>
<b>Last Change:</b>	V6.700		
<b>Default Value</b>	1000		
<b>Other Servers</b>	Cached TRS   Cached DW		
<b>Valid Settings</b>	1 -1000: Maximum cursors per process.		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter can be used to control the maximum allowable cursors per process.

### Notes

When using Oracle, if the HPE Shadowbase Oracle user account has DBA access, or otherwise has access to “V\$PARAMETER”, Oracle’s setting is used and this parameter is ignored.

Cursor reuse can occur when this parameter is set to an unsuitably low value, and can have a negative impact on performance. Information on whether cursors are being reused can be found by configuring the SHAD\_STATS\_FREQUENCY parameter appropriately.

As of version 6.700, max cursors were increased from a default of 50 to 1000.

### Related Parameters

SHAD\_STATS\_FREQUENCY

## SHAD\_DOC\_BLOCKING

<b>Initial Version:</b>	V3.930	<b>Syntax:</b>	SHAD_DOC_BLOCKING=<num>
<b>Last Change:</b>	V6.700		
<b>Default Value</b>	2		
<b>Other Servers</b>	DOC		
<b>Valid Settings</b>	0: Disable combining of DOC rows (No DOC blocking) 1: Enable combining of DOC IUD Event Rows (DOC Blocking, Version 1) 2: Enable combining of all DOC Event Rows (DOC Blocking, Version 2)		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter is used to control whether a single DOC row can contain multiple SQL statements from the same transaction. There are three variations of DOC blocking:

- Disabled (SHAD\_DOC\_BLOCKING=0): Each event is written to a separate record in the DOC.
- Version 1 (SHAD\_DOC\_BLOCKING=1): Insert, Update, and Delete events for a transaction may be combined in a single DOC record. BEGIN, COMMIT, and ABORT events are still written in individual records.

- Version 2 (SHAD\_DOC\_BLOCKING=2): All events for a transaction, including BEGIN, COMMIT, or ABORT may be combined in a single DOC record.

## Notes

Records with multiple, combined events are referred to as “Multi OPs” in sbmon filestats.

When using cached DOCs, statement data expansion can occur for in-bound data representing the ‘null’ state. In this case, the DOC database would store the character string "NULL" to represent the ‘null’ state, which would increase the total byte length by four for each column instance processed.

Version 2 of DOC Blocking was introduced in release 6.230. DOCs created with Version 2 of DOC Blocking are not compatible with releases prior to 6.230. Version 2 of DOC blocking cannot be used with multi-ported DOC writers, or with Oracle or DB2 source collectors.

As of version 6.700, the default DOC Blocking level is 2.

## Related Parameters

SHAD\_DOC\_BLOCKING\_FLUSHING\_CACHE  
SHAD\_DOC\_BLOCKING\_MAX\_RECORD\_LENGTH

## SHAD\_INVALID\_RESTART\_SHUTDOWN

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_INVALID_RESTART_SHUTDOWN=<num>
<b>Last Change:</b>	V6.700		
<b>Default Value</b>	1		
<b>Other Servers</b>	TRS   TFS		
<b>Valid Settings</b>	0: Remain idle on a nonexistent restart point 1: Shutdown if idle on a nonexistent restart point		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

## Description

This parameter controls how Shadowbase reacts to having a TFS or TRS set to a restart point that is not found in the associated DOC/OPCOL.

## Notes

Setting this parameter to 1 will cause TRS/TFS objects to shutdown when started with a nonexistent restart point, while setting it to 0 will cause the objects to use the old behavior of remaining idle until action can be taken.

## SHAD\_IGNORE\_MISMATCH

<b>Initial Version:</b>	pre v2.600	<b>Syntax:</b>	SHAD_IGNORE_MISMATCH=<num>
<b>Last Change:</b>	V6.700		
<b>Default Value</b>	1		
<b>Other Servers</b>	DOC   DW		
<b>Valid Settings</b>	0: Disabled, shutdown on mismatched object names. 1: Enabled, ignore mismatched object names.		
<b>Basic/Advanced</b>	Basic		

### Description

This parameter allows control over whether HPE Shadowbase requires matched object names. When this parameter is set to “0”, HPE Shadowbase will require object names to match, and will shutdown if they do not. In a NonStop to Open replication environment, this means the name of the object receiving data on the Open side (DOC or Direct Writer) must match the NonStop collector process name. In an Open to Open replication environment, the name of the object receiving data from the TFS must match the name of the Open Collector (example: OPCOL).

If this parameter is set to “1”, HPE Shadowbase will ignore mismatched server names, and will continue processing.

### Notes

As of version 6.700, the default value is 1.

## SHAD\_ODBC\_SQL\_C\_TYPE

<b>Initial Version:</b>	v6.700	<b>Syntax:</b>	SHAD_ODBC_SQL_C_TYPE=<num>
<b>Last Change:</b>			
<b>Default Value</b>	0		
<b>Other Servers</b>	TRS   DW		
<b>Valid Settings</b>	0: SQLBindParameter is called with SQL_C_BINARY as the C data type 1: SQLBindParameter is called with SQL_C_CHAR as the C data type		
<b>Basic/Advanced</b>	Advanced		

### Description

This parameter controls how Shadowbase target objects bind to character type data columns in ODBC targets. Setting the parameter to 0 causes it to bind as BINARY, which can allow binary data to be replicated as-is. Setting the parameter to 1 causes it to bind as CHAR data, which is how every character data column was treated until this point.

### Notes

Replicating binary data and multibyte characters with a CHAR binding can result in the characters being changed to fit the one byte character width that Shadowbase ODBC support currently handles.

## SHAD\_SQL\_FORCE\_INF

<b>Initial Version:</b>	v4.030	<b>Syntax:</b>	SHAD_SQL_FORCE_INF=<num>
<b>Last Change:</b>	v6.700		
<b>Default Value</b>	0		
<b>Other Servers</b>	TRS		
<b>Valid Settings</b>	0: Normal Behavior 1: Perform an INSERT event for every UPDATE that affects 0 rows 2: Always perform an INSERT event for every UPDATE		
<b>Basic/Advanced</b>	Advanced		

### Description

MySQL handles UPDATE statement “not found” cases without returning an error code, preventing the TRS configuration option “Convert Updates to Inserts when not found” from working as expected. When this parameter is set to 1, HPE will check to see if 0 rows were affected by a given UPDATE, and if so will convert it to an INSERT and reapply. When this parameter is set to 2, HPE Shadowbase will treat all UPDATE events as though the row they were updating was not found, and immediately reapply them as INSERTS. This setting can be useful for scenarios where triggers on the target table cause the wrong number of rows affected by an event to prevent INF from working properly.

For users replicating out to Windows servers, the following must be considered. If your database instance is configured with ‘Implicit Transactions’ and/or ‘No Count’ for transactions, then INF processing will fail, unless you enable this parameter.

### Notes

**When setting this parameter to 2 with a SQL Server target, it may be required to enable SHAD\_ODBC\_MARS\_LEVEL as well in order to handle the multiple results sets returned.**

## SHAD\_SYSKEY\_INDEX\_FIX\_LEVEL

<b>Initial Version:</b>	v6.700	<b>Syntax:</b>	SHAD_SYSKEY_INDEX_FIX_LEVEL=<num>
<b>Last Change:</b>			
<b>Default Value</b>	0		
<b>Other Servers</b>	TRS   DW		
<b>Valid Settings</b>	0: No call is made to the database relating to this workaround. 1: The workaround is explicitly disabled. 2: The workaround is explicitly enabled.		
<b>Basic/Advanced</b>	Advanced		

### Description

This parameter is meant to be a workaround to resolve an internal issue detected under certain rare circumstances when replicating SQL/MP and SQL/MX



SYSKEY key values (e.g., when replicating key-sequenced clustering key tables and preserving the primary key values). For Shadowbase to exactly replicate the source SYSKEY value to the target, the Shadowbase SHAD\_OVERRIDE\_SYSKEY parameter must be enabled. However, under certain rare instances, the SHAD\_OVERRIDE\_SYSKEY ‘override’ may need an ‘override.’

We are publishing this parameter in the extremely unlikely case it will be needed to be enabled. Do not alter this parameter from its default value unless explicitly instructed by Support to do so.

## Related Parameters

SHAD\_OVERRIDE\_SYSKEY

## SHAD\_OVERRIDE\_SYSKEY

<b>Initial Version:</b>	v6.205	<b>Syntax:</b>	SHAD_OVERRIDE_SYSKEY=<num>
<b>Last Change:</b>	v6.700		
<b>Default Value</b>	2		
<b>Other Servers</b>	Cached DW   Cached TRS (ODBC – SQL/MX only)		
<b>Valid Settings</b>	<p>0: OFF. Conditions the ODBC/MX connection to disable exact SYSKEY replication; system generated values will be used.</p> <p>1: ON. Conditions the ODBC/MX connection to enable exact SYSKEY replication. SYSKEY values must be specified and the supplied values will be used for all insert, update, and delete statements.</p> <p>2: NONE. No connection conditioning will be performed. The default behavior for the connection will be used.</p> <p>3: INSERTS only. Conditions the ODBC/MX connection to enable exact SYSKEY replication for inserts, while disabling the logic for updates and deletes (as it is not needed for these statement types). This setting provides a workaround for issues encountered in the field for updates and deletes and should be used when the customer wants to preserve the source SYSKEY value in the target row’s key.</p>		
<b>Basic/Advanced</b>	Advanced		

## Description

This parameter enables or disables exact SYSKEY replication for cached TRS and Direct Writer objects replicating to SQL/MP or SQL/MX tables via ODBC/MX connections. In all other configurations the parameter should be omitted or the default of 2 (NONE) should be used. When set to 3, it enables this parameter to only apply to inserts, while being disabled for updates and deletes, as this resolves an internal issue encountered when processing updates and deletes.

## Notes

Regarding the workaround to allow inserts and updates/deletes, it is recommended that you use EITHER the 3 (inserts) setting on this parameter, OR to set this parameter to 1 and the SHAD\_SYSKEY\_INDEX\_FIX\_LEVEL to 2

(on). Either one of these options fulfils the workaround, so it is unnecessary to set both.

## Related Parameters

SHAD\_SYSKEY\_INDEX\_FIX\_LEVEL

## SHAD\_SKIP\_INF\_ERROR

<b>Initial Version:</b>	v6.600	<b>Syntax:</b>	SHAD_SKIP_INF_ERROR=<num>
<b>Last Change:</b>	v6.700		
<b>Default Value</b>	0 (off)		
<b>Other Servers</b>	TRS   DW		
<b>Valid Settings</b>	0: OFF 1: ON (log errors) 2: ON (no logging)		
<b>Basic/Advanced</b>	Advanced		

## Description

This parameter controls how a Shadowbase process responds to certain failures on INF events. When a NonStop Consumer sends an audit compressed update for a row that does not exist on the target, it gets turned into an insert due to INF processing, and the insert fails due to the statement missing 1 or more necessary column values, this parameter controls the behavior. When it is set to the default of 0 (OFF), it will cause the TRS or DW to log a message and fail in this circumstance. If it is set to 1 (ON), it will log a message, skip the failed statement, and continue running. If it is set to 2 (ON, no logging), it will continue past failed insert statements without logging any messages to the error log.

## Notes

In prior versions of HPE Shadowbase, the default behavior was to log each of the problems encountered, and skip these types of errors. The default behavior has now been changed to make this situation an error, due to the potential for data loss if the root causes of this issue are not addressed. This parameter is joined by improved logging so as to help users diagnose this issue more easily.

## SHAD\_SQL\_ERROR\_INVALID\_NUM\_SKIP\_LOGGING

<b>Initial Version:</b>	v6.700	<b>Syntax:</b>	SHAD_SQL_ERROR_INVALID_NUM_SKIP_LOGGING=<num>
<b>Last Change:</b>			
<b>Default Value</b>	0		
<b>Other Servers</b>	Cached TRS   Cached DW		
<b>Valid Settings</b>	0: Full logging is enabled for Invalid Num processing 1: Minimal logging is enabled for Invalid Num processing 2: Logging is disabled for Invalid Num processing		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls how Shadowbase handles logging of invalid num processing. With logging disabled, there will be no record of when invalid num processing modifies values. With the parameter set to 1, it will be logged when the invalid num logic is being run, and when it is set to 0 the full statement information will be logged each time the invalid num logic is utilized.

### Related Parameters

SHAD\_SQL\_ERROR\_INVALID\_NUM

## SHAD\_SQL\_ERROR\_INVALID\_NUM\_SKIP

<b>Initial Version:</b>	v6.700	<b>Syntax:</b>	SHAD_SQL_ERROR_INVALID_NUM_SKIP=<num>
<b>Last Change:</b>			
<b>Default Value</b>	0		
<b>Other Servers</b>	Cached TRS   Cached DW		
<b>Valid Settings</b>	0: Fail when attempting to set a non-nullable column to NULL 1: Skip statements with invalid num columns that cannot be set to NULL		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter controls how Shadowbase handles non-nullable numeric columns with invalid num logic. By default, if invalid num logic attempts to set a non-nullable numeric column to NULL, it will cause an error and fail. This parameter allows the entire statement with the incompatible invalid num value to be skipped and logged so as to not bring down the environment.

## SHAD\_ODBC\_CONCURRENCY

<b>Initial Version:</b>	V4.031	<b>Syntax:</b>	SHAD_ODBC_CONCURRENCY=<num>
<b>Last Change:</b>	V6.700		
<b>Default Value</b>	-1		
<b>Other Servers</b>	TRS   DW		
<b>Valid Settings</b>	-1: Use DBMS default mode. 1: READ ONLY. 2: LOCK. 3: ROWVER. 4: VALUES		
<b>Basic/Advanced</b>	Advanced – See Basic vs. Advanced Parameters section		

### Description

This parameter specifies the cursor concurrency to be implemented. Although HPE Shadowbase does not specifically use cursors to perform replication, performance of some RDBMS can be modified through the setting of this ODBC attribute.

For a setting of -1 the concurrency mode is established by the Driver or DBMS. When set to 1 the cursor is read-only. No updates are allowed. 2 sets the cursor to use the lowest level of locking sufficient to ensure that the row can be updated. 3 sets the cursor to use optimistic concurrency control, comparing row versions such as SQLBase ROWID or Sybase TIMESTAMP. 4 sets the cursor to use optimistic concurrency control comparing values.

### Notes

Options 3 and 4 are not supported by every DBMS. In these cases the DBMS will substitute those options with option 2 and warnings will be logged.

## SHAD\_OPCOLLECT\_TIMESTAMP\_CONVERSION\_LEVEL

<b>Initial Version:</b>	V6.700	<b>Syntax:</b>	SHAD_OPCOLLECT_TIMESTAMP_CONVERSION_LEVEL=<num>
<b>Last Change:</b>			
<b>Default Value</b>	0		
<b>Other Servers</b>	OPCOL		
<b>Valid Settings</b>	0 – 1000		
<b>Basic/Advanced</b>	Basic		

### Description

This parameter specifies the number of period characters “.” that are to be replaced with colons “:” in collected TIMESTAMP WITH TIMEZONE column values.

Setting this parameter to 0 will disable timestamp conversion altogether, while any values between 1 and 1,000 will attempt to replace up to that many periods.

Setting this to a value larger than the number of periods present will cause periods to be replaced in the value until there are none left.

**Notes**

Default formatted `TIMESTAMP WITH TIMEZONE` columns will have 2 periods in the time portion to replace, while the third period denotes fractional seconds and likely should not be replaced to maximize compatibility.

## ***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 (V6.600).

Performing daily license check for system [<system-name>]  
Daily Check: License for system [<system-name>] is valid.  
SHUTTING DOWN: daily license validation failure for system [<system-name>]  
– contact Shadowbase support

Cause: Licenses are now validated daily to check their expiration

Effect: If the license has expired the system notifies the user and shuts down. Otherwise it logs the success of the check and continues

Recovery: Contact Shadowbase Support for a valid license

SHAD\_OP COLLECT\_COLUMN\_WRAP set to start with '<StartWrapChar>' and end with '<EndWrapChar>'

Cause: Documents the characters that will be used to wrap Oracle column names defined by SHAD\_OP COLLECT\_COLUMN\_WRAP parameter.

Changing Update to Insert

Cause: When SHAD\_DISPLAY\_INF is set to 1 this will be logged with the corresponding statement when an update is changed to be an insert because an insert is not found

## Known Problems Remaining in Version 6.700

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\_OPACCOLLECT\_LOG\_STATS=3). If you specify SHAD\_OPACCOLLECT\_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 SHADPARAM.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 SHADPARAM.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\_LIBRARY\_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 '% \';
```

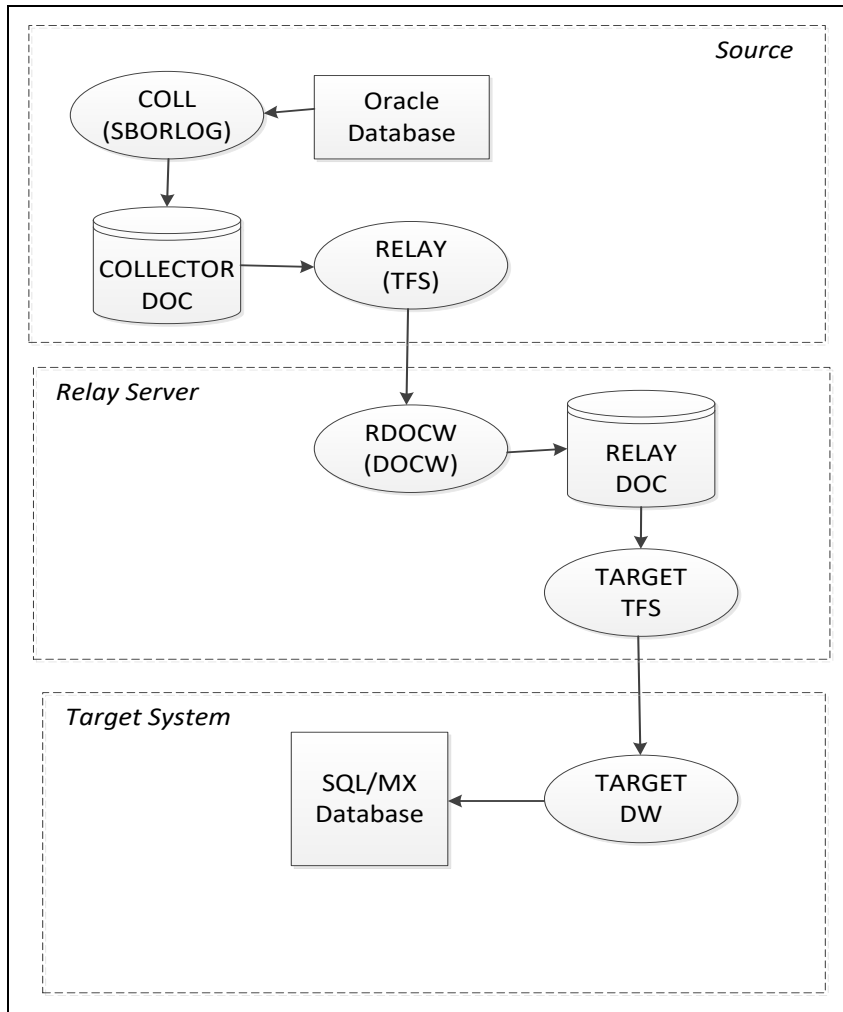
```
UPDATE SHAD_USERS_EXCLUDE
SET USERNAME=RTRIM(USERNAME)
WHERE USERNAME LIKE '% \';
```

## 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. There remains an issue with target replication into SAP HANA columns with the BOOLEAN data type. SAP HANA itself can accept literal string “TRUE” and “FALSE” values which it converts and stores as TINYINT, however HPE Shadowbase target replication currently only supports replicating numeric values into HANA BOOLEAN columns.
24. 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.
25. 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.

26. 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.
27. 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.
28. 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.
29. 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
30. 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.
31. 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.

32. There is an issue when running from an OPLOG to a TRS on the same system with high cache invalidation. This causes the TRS to abend once it has processed around 32 thousand transactions. The workarounds for this are to set SHAD\_MAX\_CURSORS to a higher value that prevents the need to remove and re-add cached statements (i.e. enough to hold all of the statements used by the TRS in cache), or to go from a source into a DOC then from the DOC to a TRS.
33. 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.
34. Replication to DB2 columns with data types of BINARY, VARBINARY, and BLOB is not currently supported.
35. 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.
36. 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.
37. 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.
38. 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.
39. 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.
40. The Microsoft SQL Server 'SQL\_VARIANT' data type has limited support and has not been fully tested.
41. 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.



## Installation Instructions

To install Shadowbase, new steps have been added at the beginning to ensure that users acknowledge the license format change since the previous version. 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).

\*\*\*\*\* End of Document \*\*\*\*\*