

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

| | |
|-------------------|---|
| Product: | HPE Shadowbase for Other Servers |
| Release: | Gravic Version 6.600 T1123-AAF (SB REPL/OSS) T1124-AAC (SB REPL/HP UX) T1125-AAE (SB REPL/SUSE LINUX) T1126-AAD (SB REPL/RHEL LINUX) T1127-AAC (SB REPL/SOLARIS X86) T1128-AAC (SB REPL/SOLARIS) T1129-AAD (SB REPL/MS WINDOWS) T1130-AAC (SB REPL/AIX) |
| Date: | April 1, 2020 |
| Copyright Notice: | Copyright Gravic, Inc. 1995 – 2020 (www.gravic.com) |
| File Name: | IPM6600_other_servers.pdf |

NOTE: This softdoc covers new features and corrected problems for HPE Shadowbase for Other Servers, Version 6.600. 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

| | |
|---|-----------|
| Table of Contents | 1 |
| Disclaimer | 2 |
| Mature Features and Deprecated Features | 3 |
| Supported Database and Operating System Versions | 5 |
| Special Notes for Version 6.600 | 6 |
| New Features, Problems Corrected, and Other Changes in Version 6.600 (since Version 6.300) | 8 |
| New Features in Version 6.600 | 8 |
| Problems Corrected in Version 6.600 | 10 |
| New & Modified SBMON Commands | 12 |
| ADD / EDIT / EDIT! Prompt Changes | 12 |
| New & Modified shadparm.ini Parameters | 14 |
| SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS | 14 |
| SHAD_DB2COLLECT_MQ_CHANNEL_NAME | 14 |
| SHAD_DB2COLLECT_MQ_CHANNEL_PORT | 15 |
| SHAD_DB2COLLECT_USE_MQCONN | 16 |
| SHAD_FLUSH_LEVEL | 17 |
| SHAD_FLUSH_SLEEP_ABNORMAL | 17 |
| SHAD_FLUSH_SLEEP_NORMAL | 18 |
| SHAD_OPCOLLECT_FILTER_PARTIAL_TXS | 18 |
| SHAD_OPCOLLECT_LARGEST_TX_SPAN | 19 |
| SHAD_OPCOLLECT_LOG_PARTIAL_TXS | 19 |
| SHAD_OPCOLLECT_PRE_V6600_REPOSITIONING_ALGORITHM | 20 |
| SHAD_OPCOLLECT_RETRY_SQL_MISMATCH | 20 |
| SHAD_OPCOLLECT_SELECT_SB_EVENTS | 21 |
| SHAD_OPCOLLECT_TEST_ONLY_SKIP_START_EVENTS | 21 |
| SHAD_OPCOLLECT_TRACE_SQL_MISMATCH_RETRIES | 22 |
| SHAD_OVERRIDE_IDENTITY | 22 |
| SHAD_RETRY_CTREE_ERRORS | 23 |
| SHAD_SAP_ASE_VERIFY_NUMERIC | 23 |
| SHAD_SHADTRS_DUP_TX_LOGGING | 24 |
| SHAD_SKP_INF_ERROR | 24 |
| SHAD_TRANS_EXPECTED_ENDS | 25 |
| New & Modified User Messages | 27 |
| Known Problems Remaining in Version 6.600 | 31 |
| Installation Instructions | 42 |

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.

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 |
|------------------------------|---|------------------------|
| Oracle Trigger Source | 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 |
|---------|-------------|--------------------|-------------------|
|---------|-------------|--------------------|-------------------|

| | | | |
|----------------------------------|--|--------------------------|-------------|
| Execute Immediate Objects | <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"> 1. sbocons 2. sbmspre 3. sborapre 4. sbsypre 5. sbigdir 6. sbmswrt 7. sborcol 8. sbsywrt 9. sbigwrt 10. sborcol <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 |
| SBDOCRD | <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 |
| SQL92 | <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 |
| UserExit Eclipse projects | <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 for additional information regarding Shadowbase support for databases and operating systems.

Special Notes for Version 6.600

1. Version 6.600 is a “general availability” (non-restricted TCF) release. The previous non-restricted TCF release is version 6.220 (HPE version AAB) for HP-UX and IBM AIX, Oracle Solaris, and Oracle Solaris x86 systems, version 6.240 (HPE version AAE) for HPE OSS systems, and version 6.300 (HPE version AAC or AAD (SUSE)) for Microsoft Windows Server, Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) systems.
2. This document describes the changes since Version 6.300. If you are upgrading from a version of Shadowbase for Other Servers prior to that, please first see the version 6.300 software release document, IPM6300_other_servers, for a list of the changes in that version.
3. Due to licensing changes introduced in Version 6.100, existing installations of Shadowbase prior to that version will require a new password file in order to run after the upgrade.
3. Version 6.600 contains support for SAP Sybase ASE databases as a target using ODBC for cached statements only. It does not support Execute Immediate (EI) statements for Sybase. If you are using EI for Sybase, you cannot upgrade to Version 6.600.
4. 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 6.600 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.

5. 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.

6. 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.

7. 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.

New Features, Problems Corrected, and Other Changes in Version 6.600 (since Version 6.300)

New Features in Version 6.600

This section provides a summary of the features added to HPE Shadowbase for Other Servers products since the previous general availability release (V6.300).

1. The SHAD_SHADTRS_DUP_TX_LOGGING option has been added to control whether or not the TRS logs a message when it is skipping a duplicate transaction when applying to the target database.
2. The SHAD_RETRY_CTREE_ERRORS option has been added to allow for one to retry on any arbitrary CTREE error. This option deprecates the SHAD_RETRY_CTREE_ERROR_527 option.
3. The Microsoft SQL Server TRSs and DWs will now connect to SQL Server using a DSN supplied by the user at configuration time. The newer Microsoft ODBC Driver for SQL Server is now supported by Shadowbase, which allows non-Windows Shadowbase TRSs and DWs to connect to a SQL Server instance running on a Windows server. See DSN Prompt (MSSQL TRS and DW) for details about how to fall back to the old behavior of using the SQL Server Native Client to connect.
4. SAP Hana is now supported as a target database through ODBC TRSs and DWs.
5. Shadowbase DB2 Source replication has been enhanced to provide the user with more configuration options for connecting to IBM MQ. The DB2 OPCOL now connects to MQ via MQCONN rather than MQCONN, which allows the user to specify the MQ channel address, port, and name. See the parameters SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS, SHAD_DB2COLLECT_MQ_CHANNEL_NAME, SHAD_DB2COLLECT_MQ_CHANNEL_PORT for information about how to specify this information.
6. Configuring a DB2 OPCOL in SBMON now allows the user to specify a username and password as the last configuration options. See the parameter SHAD_DB2COLLECT_USE_MQCONN for information about how to fall back to the pre-v6.600 behavior.
7. Sybase target replication is now done entirely through ODBC, and as such, Sybase TRSs and DWs should be configured as sbodbctrs and sbodbcdw. This allows Sybase replication to be kept up to date with the latest features added to ODBC replication (for example, the performance-enhancing Shadowbase statement caching algorithm).

8. DB2 target replication is now done through ODBC Cache objects. TRSs and DWs should now be configured as sbodbctrs and sbodbcdw. This allows DB2 target replication to be more robust and to enjoy a boost in performance that was missing from prior DB2 target replication.
9. Added the SBFILE command DUMPSTMNTBYTX that accepts a DOC name and a DOC sequence number and displays all of the statements of that DOC in transaction order, as a consumer (such as a TFS or TRS) would process it.
10. Added the SBFILE command DUMPSTMNTFORTX that accepts a DOC name, a DOC sequence number, and a transaction ID, and searches the DOC writer's associated DOCs for the specified transaction, displaying the associated events in the order in which they arrived.
11. Leaving the IP address field blank on TRS and DW consumers will now default to localhost, connecting back to the machine they are hosted on.

Problems Corrected in Version 6.600

This section provides a summary of the problems corrected in HPE Shadowbase for Other Servers products since the previous general availability release (V6.300).

1. Oracle Log Source (OLS) SQL mismatch errors: The OLS collector uses Logminer queries to collect the SQL applied against the database by periodically issuing overlapping selects against the log. It includes logic to ensure that it read the last record collected in the previous query. The logic verifies that it read a record at the position of the last record collected, and, as a final check, that the SQL statement collected matches that in the statement read in the new query. Depending on the configuration, the collector may shutdown or continue.

One of our customers has reported seeing infrequent statement mismatches. In all cases, the statement collected was a “COMMIT”, and the statement read at the same address was a “ROLLBACK”. Unfortunately, Gravic was not able to reproduce the issue in-house, nor was our customer able to reliably reproduce the problem. As a result, we were not able to determine a root cause.

One possibility is that there may be multiple records reported at a single address. This release contains changes to handle that multiple records at the same address, should that happen.

Another possibility is that restart position closer to the record in question causes the record to be missed (we have seen similar behavior running Oracle logminer and overlap the queries by a configurable amount to protect against it). OLS has been enhanced to re-issue the query, starting earlier in the log. The first retry will start back at twice the configured value, the second retry, if necessary, will start at the beginning of the previous query.

OLS has also been enhanced to automatically turn on extensive diagnostic logging when the condition is detected and being retried.

2. TRS column mismatch errors were sometimes encountered when using SQL92 or ORACLE SQL formats while replicating data containing ‘TO_DATE (’. Shadowbase SQL processing has been enhanced to avoid this issue.
3. There was an issue in SBFILE where a string value of ‘NULL’ was interpreted as a literal NULL value. The processing in SBFILE now displays the string ‘NULL’ as it would any other string.
4. Shadowbase DB2 target replication previously had issues that led TRSs and DWs to shut down while replicating data from a NonStop source. DB2 Target replication has been made more robust so as to avoid these unplanned shut downs.
5. Shadowbase previously had problems when it attempted to insert rows into a table that the supplied user does not have permission to access. The permission error was

incorrectly interpreted as an Insert Not Found case. The error parsing has been enhanced so that Shadowbase will recognize a permission error and will shut down in order to give the user a chance to fix the permission issue before Shadowbase moves too far ahead.

6. NonStop TFSs (sbnskw) on Solaris SPARC would fail without any error message when replicating to a NonStop Consumer. The sbnskw TFS has been fixed and will now replicate data to a NonStop Consumer as intended.
7. There was an issue with TRSs and DWs where OSB would receive an Audit Compressed update on a row that does not exist on the target and, interpreting it as an insert not found case, it would turn that Audit Compressed update into an insert, which would fail due to the Audit Compress leaving out fields that were required for the insert but not essential for the original update. Shadowbase has been enhanced to understand this circumstance specifically instead of seeing the two errors that it triggered previously. The new behavior in this instance is that the TRS or DW will log a message explaining the situation and shut down. See the new parameter SHAD_SKIP_INF_ERRORS to adjust this behavior.
8. The reload command for Microsoft SQL Server has been corrected. You no longer need to reboot the database in order to reenable collection. This has the side-effect that it is now possible to re-enable collection to an OPCOL that is in an erroneous state causing undesirable behavior or possible data corruption. Ensure you read the error log before re-enabling collection as to not enter an erroneous state.
9. An issue with the LIST command in SBMON has been fixed. Previously, consumers of a DOC or OPCOL had a chance to be missing the column that describes their state as running or stopped, with the rest of the information being shifted to the left in that information's absence. The state of these consumers should now be reflected properly.
10. The "list <DOC name>" command has been corrected to show TRS progress. It will no longer erroneously show that zero transactions have been processed.
11. There was an issue in the TRS that could cause the WHERE clause in some statements in the SHAD_REJECT table to be incorrect due to trailing digits being added when SHAD_REJECT processing was enabled.
12. Oracle Log Source (OLS) Data "Ponging": The algorithm used by Shadowbase to detect and skip replicated events in a bi-directional environment could fail, resulting in some events being replicated and re-applied on the original system incorrectly. Shadowbase uses the transaction name returned by Oracle Log Miner to identify replicated events. In some circumstances, particularly with large batch transactions, Oracle Log Miner does not properly return the associated transaction name with the event, causing Shadowbase to improperly identify the event as locally originated. This issue has been fixed.

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 (V6.300).

ADD / EDIT / EDIT! Prompt Changes

DSN Prompt (MSSQL TRS and DW)

```
Data Sources(ODBC) User DSN [] :
```

```
Change currently defined DSN (Y/[N])?
```

```
You have chosen an empty DSN.  
Is this correct? Choosing a blank DSN will revert back to pre-v6600 behavior of  
connection via the SQL Server Native Client. (y/N) :
```

These determine whether MSSQL DWs and TRSs use a standard, user defined ODBC Data Source Name (DSN) to connect to the specified database. Setting a DSN enables Shadowbase objects to use both the old SQL Server Native Client Drivers or the newer Microsoft ODBC Drivers, provided there is a properly configured ODBC Data Source setup. This change additionally enables replication to SQL Server from both Red Hat Enterprise Linux and SUSE Linux Enterprise Server via the newly released Microsoft ODBC Driver for SQL Server on Linux. Leaving this field blank will result in the TRS or DW falling back to the 6.300 method of connecting to SQL Server via the SQL Server Native Client 10 or 11.

Prior to Version 6.600, Shadowbase would use the provided server, username, password, and database name to connect to SQL Server via the SQL Server Native Client. From Version 6.600 and beyond, the TRS and DW behavior has been updated to be more in line with how ODBC TRSs and DWs connect to SQL Server, using just the DSN, username and password. Extra attributes like Server and Database Name may be set in the ODBC Data Source configuration. The pre-6.600 behavior can still be used by not selecting any value for DSN when configuring an MSSQL object via SBMON.

While using this new, more versatile method of replicating to SQL Server is now recommended over the old method of connecting without a DSN via SQL Server Native Client, existing TRS and DW objects configured for MSSQL will retain their old functionality when upgrading from a pre-6.600 version of HPE Shadowbase for Other Servers to 6.600. These objects can be converted to the newer behavior by setting up a Data Source through ODBC, and editing the object in SBMON, responding yes when asked if you would like to change the currently configured DSN. Setting a DSN here will

switch the object to the new behavior. This can be switched back at any time by editing the object again and setting a blank value for DSN.

Notes:

- 1) Setting a blank DSN will trigger the prompt asking if you are sure you want to set a blank DSN and revert to pre-6.600 behavior. This will affect scripted creation of MSSQL objects in that there is now an additional prompt at the end while using the ADD command, with another yes or no prompt if the DSN is left empty
- 2) When editing MSSQL objects, you will be prompted for if you want to change the DSN. This will affect scripted editing of MSSQL objects in that there is now an additional prompt at the end while using the EDIT command, with another prompt if you actually edit the DSN, and yet another prompt if you then set the DSN to nothing.

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 (V6.300).

SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS=<text> |
| Last Change: | -- | | |
| Default Value | -- | | |
| Other Servers | DB2 Log Based Open Collector | | |
| Valid Settings | The address of the IBM MQ server, either as an IP address, or a hostname. | | |

Description

This parameter, when SHAD_DB2COLLECT_USE_MQCONN is set to 1, controls where Shadowbase will look to connect to an IBM MQ channel. This value can be an IP address for remote and local MQ connections, as well as a hostname such as `localhost`.

Notes

This parameter is only used if SHAD_DB2COLLECT_USE_MQCONN is set to 1, and must be used in conjunction with SHAD_DB2COLLECT_MQ_CHANNEL_PORT and SHAD_DB2COLLECT_MQ_CHANNEL_NAME otherwise Shadowbase will fall back to the address found in the MQSERVER environment variable. This parameter must be specified in the shadparm.ini file prior to startup of a DB2 Log Based Open Collector in order for Open Collection to observe the requested setting. Altering this parameter for an executing DB2 Log Based Open Collector object requires a restart of the Open Collector for those alterations to have an effect.

Related Parameters

SHAD_DB2COLLECT_USE_MQCONN
SHAD_DB2COLLECT_MQ_CHANNEL_PORT
SHAD_DB2COLLECT_MQ_CHANNEL_NAME

SHAD_DB2COLLECT_MQ_CHANNEL_NAME

| | | | |
|-------------------------|--|----------------|--|
| Initial Version: | V6.600 | Syntax: | SHAD_DB2COLLECT_MQ_CHANNEL_NAME=<text> |
| Last Change: | -- | | |
| Default Value | -- | | |
| Other Servers | DB2 Log Based Open Collector | | |
| Valid Settings | The name of the MQ channel for Shadowbase to connect to. | | |

Description

This parameter, when SHAD_DB2COLLECT_USE_MQCONN is set to 1, controls the specific name of the IBM MQ channel for Shadowbase to connect to.

This can be used to switch a DB2 Open Collector between channels on the same server.

Notes

This parameter is only used if SHAD_DB2COLLECT_USE_MQCONN is set to 1, and must be used in conjunction with SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS and SHAD_DB2COLLECT_MQ_CHANNEL_PORT otherwise Shadowbase will fall back to the address found in the MQSERVER environment variable.

This parameter must be specified in the shadparm.ini file prior to startup of a DB2 Log Based Open Collector in order for Open Collection to observe the requested setting. Altering this parameter for an executing DB2 Log Based Open Collector object requires a restart of the Open Collector for those alterations to have an effect.

Related Parameters

SHAD_DB2COLLECT_USE_MQCONN
SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS
SHAD_DB2COLLECT_MQ_CHANNEL_PORT

SHAD_DB2COLLECT_MQ_CHANNEL_PORT

| | | | |
|-------------------------|--|----------------|---------------------------------------|
| Initial Version: | V6.600 | Syntax: | SHAD_DB2COLLECT_MQ_CHANNEL_PORT=<num> |
| Last Change: | -- | | |
| Default Value | -- | | |
| Other Servers | DB2 Log Based Open Collector | | |
| Valid Settings | The port of the IBM MQ server that Shadowbase should connect to. | | |

Description

This parameter, when SHAD_DB2COLLECT_USE_MQCONN is set to 1, controls the port on which Shadowbase will look for an IBM MQ server. This value must be five digits or smaller, otherwise the value will be truncated to the first 5 digits.

Notes

This parameter is only used if SHAD_DB2COLLECT_USE_MQCONN is set to 1, and must be used in conjunction with SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS and SHAD_DB2COLLECT_MQ_CHANNEL_NAME otherwise Shadowbase will fall back to the address found in the MQSERVER environment variable.

This parameter must be specified in the shadparm.ini file prior to startup of a DB2 Log Based Open Collector in order for Open Collection to observe the requested setting. Altering this parameter for an executing DB2 Log Based Open Collector object requires a restart of the Open Collector for those alterations to have an effect.

Related Parameters

SHAD_DB2COLLECT_USE_MQCONN

SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS
SHAD_DB2COLLECT_MQ_CHANNEL_NAME

SHAD_DB2COLLECT_USE_MQCONN

| | | | |
|-------------------------|---|----------------|----------------------------------|
| Initial Version: | V6.600 | Syntax: | SHAD_DB2COLLECT_USE_MQCONN=<num> |
| Last Change: | -- | | |
| Default Value | 1 (ON) | | |
| Other Servers | DB2 Log Based Open Collector | | |
| Valid Settings | 0 = OFF (OPCOL uses newer MQCONNX to connect to IBM MQ) 1 = ON (OPCOL falls back to pre-6.600 method of using MQCONN to connect to IBM MQ) | | |

Description

This parameter sets whether the DB2 OPCOL uses the newer MQCONNX, which allows for more flexible connections, to connect to IBM MQ or falls back to the old method of using MQCONN combined with values derived from the MQSERVER environment variable.

Notes

This parameter being turned on allows the related parameters SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS, SHAD_DB2COLLECT_MQ_CHANNEL_PORT, and SHAD_DB2COLLECT_MQ_CHANNEL_NAME to be set, allowing for remote connections and for multiple MQ instances on an MQ server to be differentiated. This parameter must be specified in the shadparm.ini file prior to startup of a DB2 Log Based Open Collector in order for Open Collection to observe the requested setting. Altering this parameter for an executing DB2 Log Based Open Collector object requires a restart of the Open Collector for those alterations to have an effect.

Related Parameters

SHAD_DB2COLLECT_MQ_CHANNEL_ADDRESS
SHAD_DB2COLLECT_MQ_CHANNEL_PORT
SHAD_DB2COLLECT_MQ_CHANNEL_NAME

SHAD_FLUSH_LEVEL

| | | | |
|-------------------------|---|----------------|------------------------|
| Initial Version: | v3.910 | Syntax: | SHAD_FLUSH_LEVEL=<num> |
| Last Change: | v6.600 | | |
| Default Value | 1 | | |
| Other Servers | DOC OPCOL Custom | | |
| Valid Settings | 0: No forced DOC disk cache flushing, SHAD_FLUSH_SLEEP_NORMAL and SHAD_FLUSH_SLEEP_ABNORMAL in effect. 1: Level 0 and OS controlled DOC disk cache flushing via fflush(); flushall() on Windows 2: Level 0 and OS controlled disk cache flushing via sync() 3: Level 0 and OS controlled DOC disk cache flushing via fflush() and sync() | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter controls how an HPE Shadowbase object handles its DOC's disk cache

Notes

Setting the flush level to 3 may cause severe DOC slowdowns.

Related Parameters

SHAD_FLUSH_SLEEP_NORMAL
SHAD_FLUSH_SLEEP_ABNORMAL
SHAD_DOC_WRITETHRU_LEVEL

SHAD_FLUSH_SLEEP_ABNORMAL

| | | | |
|-------------------------|--|----------------|---------------------------------|
| Initial Version: | Pre v6.100 | Syntax: | SHAD_FLUSH_SLEEP_ABNORMAL=<num> |
| Last Change: | v6.600 | | |
| Default Value | 0 | | |
| Other Servers | DOC OPCOL | | |
| Valid Settings | 0 - 2147483647 milliseconds. | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter specifies the number of milliseconds the process will sleep to allow buffers to be flushed when an error or abnormal situation is detected. The default is 0 milliseconds.

Related Parameters

SHAD_FLUSH_SLEEP_NORMAL
SHAD_FLUSH_LEVEL
SHAD_DOC_WRITETHRU_LEVEL

SHAD_FLUSH_SLEEP_NORMAL

| | | | |
|-------------------------|--|----------------|-------------------------------|
| Initial Version: | Pre v6.100 | Syntax: | SHAD_FLUSH_SLEEP_NORMAL=<num> |
| Last Change: | v6.600 | | |
| Default Value | 0 | | |
| Other Servers | DOC OPCOL | | |
| Valid Settings | 0 - 2147483647 milliseconds. | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter specifies the number of milliseconds the process will sleep to allow buffers to be flushed in normal situations. The default is 0 milliseconds.

Notes

This parameter should be used in the rare case of Shadowbase objects failing to record transactions due to a filesystem failure to flush the transaction from memory to disk. If you encounter this problem, setting this parameter to a value of 1 or 2 milliseconds could potentially give the Shadowbase objects enough time to flush transactions to disk, thereby fixing the problem. Keep in mind, every millisecond added will result in significant decrease in transaction throughput.

Related Parameters

SHAD_FLUSH_SLEEP_ABNORMAL
SHAD_FLUSH_LEVEL
SHAD_DOC_WRITETHRU_LEVEL

SHAD_OPCCOLLECT_FILTER_PARTIAL_TXS

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCCOLLECT_FILTER_PARTIAL_TXS = <number> |
| Last Change: | -- | | |
| Default Value | 1 (ON) | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0 =OFF (Partial transactions are not filtered) 1 =ON (Partial transactions are filtered) | | |

Description

This parameter specifies whether partial transactions, transactions that have insert, update, or delete events but no start events should be included for replication or skipped. The default, 1, skips the transactions.

Since events are read from Oracle Log Miner in event order, not in transaction order, partial transactions are expected when OLS is started. If OLS is restarted, the partial transactions will have been completely collected during the previous run and can be safely ignored with no data loss. On a cold start at the beginning or the end of the log, the default setting will only collect events from a complete transaction and will skip in-progress transactions.

Gravic strongly recommends in using the default setting for this parameter and filtering partial transactions.

SHAD_OPCOLLECT_LARGEST_TX_SPAN

| | | | |
|-------------------------|-----------------------------|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_LARGEST_TX_SPAN = <number> |
| Last Change: | -- | | |
| Default Value | 2000 | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0 - 2147483647 | | |

Description

This parameter specifies the number of Oracle System Change Number (SCNs) the largest transaction could span. The default is 2000 SCNs.

This parameter is used in two ways. First, if SHAD_OPCOLLECT_LOG_PARTIAL_TXS is set to ON, partial transactions will only be logged after the OLS collector has read and processed at least SHAD_OPCOLLECT_LARGEST_TX_SPAN SCNs.

The second use only occurs when:

- the OLS collector is in a bi-direction configuration,
- the OLS is started at the beginning or ending of the log, and
- SHAD_OPCOLLECT_FILTER_PARTIAL_TRANSACTIONS is 0 (off).

To prevent ponging, the collector reads SHAD_OPCOLLECT_LARGEST_TX_SPAN SCNs prior to the start of event collection to insure the Shadowbase transaction names are properly returned by Oracle.

SHAD_OPCOLLECT_LOG_PARTIAL_TXS

| | | | |
|-------------------------|--|----------------|--|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_LOG_PARTIAL_TXS= {ON OFF ALL} |
| Last Change: | -- | | |
| Default Value | ON | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | ON = Any partial transaction (a transaction without a START event) will be logged after the initial events at startup are read. OFF = Partial transactions are not logged. ALL = All partial transactions are logged, even at startup. | | |

Description

This parameter specifies whether partial transactions (transactions without a START event) are logged. The default setting, ON, will start logging partial transaction after the initial startup of the OLS collector. Partial transactions are expected at startup due to the nature of reading of reading from Oracle Log Miner.

SHAD_OPCOLLECT_PRE_V6600_REPOSITIONING_ALGORITHM

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_PRE_V6600_REPOSITIONING_ALGORITHM = <NUM> |
| Last Change: | -- | | |
| Default Value | 0 | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0: Do not use the repositioning algorithm in versions prior to V6.600. 1: Use repositioning algorithm in versions prior to V6.600. | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter enables the use of the repositioning algorithm used in pre V6.600 versions of the OLS collector. It should not be used except under the direction of Shadowbase support.

Related Parameters

[SHAD_OPCOLLECT_RETRY_SQL_MISMATCH](#)

SHAD_OPCOLLECT_RETRY_SQL_MISMATCH

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_RETRY_SQL_MISMATCH=<num> |
| Last Change: | -- | | |
| Default Value | 1 | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0: Do not retry SQL mismatch errors when repositioning in the Oracle log. 1: Retry the SQL mismatch errors when reposition in the Oracle log | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter controls the behavior of OLS when the last SQL statement collected in the previous query does not match the SQL statement collected at the same address in the new query. The default value, 1, will cause the OLS collector to retry the query several times, starting further back in the log. Setting the value 0 will inhibit the logic.

Related Parameters

[SHAD_OPCOLLECT_TRACE_SQL_MISMATCH_RETRIES](#)

SHAD_OPCOLLECT_SELECT_SB_EVENTS

| | | | |
|-------------------------|--|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_SELECT_SB_EVENTS= <number> |
| Last Change: | -- | | |
| Default Value | 0 (OFF) | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0 =OFF (Events generated by Shadowbase in a bi-directional environment are selected) | | |
| | 1 =ON (Events generated by Shadowbase are not selected) | | |

Description

This parameter specifies whether insert, update, and delete events generated by Shadowbase in a bi-directional environment are filtered by the SELECT statement used to read Oracle Log Miner events. If off (0, the default), the SELECT statement includes a sub-clause in the WHERE clause to exclude the Shadowbase transactions. If on, that sub-clause is omitted and all IUD events are collected.

If SHAD_OPCOLLECT_FILTER_PARTIAL_TXS is off (0), this parameter is ignored and Shadowbase events are always collected.

Shadowbase generated events will be skipped regardless of the setting of this parameter in a bi-directional environment.

SHAD_OPCOLLECT_TEST_ONLY_SKIP_START_EVENTS

| | | | |
|-------------------------|--------------------------------|----------------|--|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_TEST_ONLY_SKIP_START_EVENTS = ENABLE_TESTING |
| Last Change: | -- | | |
| Default Value | N/A (testing not enabled) | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | Omitted, or ENABLE_TESTING. | | |

Description

This parameter causes the Oracle Log Source collector to skip START transaction events for testing only. It is used internally by Shadowbase development and quality assurance to test the parameters related to collecting partial transactions – it essentially makes every transaction a partial transaction.

WARNING! Never set this parameter in a production environment. Setting this parameter can cause data loss and excessive logging:

- If SHAD_OPCOLLECT_FILTER_PARTIAL_TXS is set, all events will be skipped.
- If SHAD_OPCOLLECT_LOG_PARTIAL_TXS is enter ON or ALL, all transactions will be logged.

SHAD_OPCOLLECT_TRACE_SQL_MISMATCH_RETRIES

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_OPCOLLECT_TRACE_SQL_MISMATCH_RETRIES=<num> |
| Last Change: | -- | | |
| Default Value | 0 | | |
| Other Servers | Oracle Log Source Collector | | |
| Valid Settings | 0: Do not log detailed diagnostic information when retrying an SQL statement mismatch while reposition in the Oracle log. 1: Log detailed diagnostic information during the retry. | | |
| Basic/Advanced | Advanced – See Basic vs. Advanced Parameters section | | |

Description

This parameter enables or disables detailed tracing by the OLS collector when the last SQL statement collected in the previous query does not match the SQL statement collected at the same address in the new query. The default value, 0, will cause the OLS collector not to issue diagnostic traces, which can be voluminous, during the retry.

Related Parameters

[SHAD_OPCOLLECT_RETRY_SQL_MISMATCH](#)

SHAD_OVERRIDE_IDENTITY

| | | | |
|-------------------------|---|----------------|------------------------------|
| Initial Version: | V6.600 | Syntax: | SHAD_OVERRIDE_IDENTITY=<num> |
| Last Change: | -- | | |
| Default Value | 2 (NONE) | | |
| Other Servers | Cached DW Cached TRS (ODBC – SQL/MX only) | | |
| Valid Settings | 0: OFF. Conditions the ODBC/MX connection to disable exact IDENTITY column replication; system generated values will be used. 1: ON. Conditions the ODBC/MX connection to enable exact IDENTITY column replication. IDENTITY column values must be specified and the supplied values will be used. 2: NONE. No connection conditioning will be performed. The default behavior for the connection will be used. | | |
| Basic/Advanced | Basic | | |

Description

This parameter enables or disables exact IDENTITY column 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.

Notes

Prior to Version 6.600, the default way to replicate tables with IDENTITY columns was to maintain a mapping table and user exit to ensure updates and deletes were replicated to the proper row. The new 6.600 behavior eliminates the need for said mapping table and user exit, and allows OSB to replicate IDENTITY columns without error.

SHAD_RETRY_CTREE_ERRORS

| | | | |
|-------------------------|---|----------------|---|
| Initial Version: | V6.600 | Syntax: | SHAD_RETRY_CTREE_ERRORS=<num> <num> ... |
| Last Change: | -- | | |
| Default Value | (NONE) | | |
| Other Servers | TRS, TFS | | |
| Valid Settings | A space separated list of any integer between 1 and 65534 | | |

Description

This parameter specifies a list of space separated c-Tree errors to retry while reading from the STMT DOC. By default, none are retried and the TRS or TFS will stop. We have found that in most cases, these errors are transient. Setting this parameter to on allows the TRS or TFS to recover faster than the normal method of stopping and restarting.

Some errors, most notably 527, seem more prevalent on systems with busy disks and network based file systems. Systems that have the DOC data on local drives do not encounter these errors as frequently.

This setting deprecates SHAD_RETRY_CTREE_ERROR_527.

SHAD_SAP_ASE_VERIFY_NUMERIC

| | | | |
|-------------------------|--|----------------|-----------------------------------|
| Initial Version: | v6.600 | Syntax: | SHAD_SAP_ASE_VERIFY_NUMERIC=<num> |
| Default Value | 1 (ON) | | |
| Other Servers | Cached DW Cached TRS (ODBC – SAP ASE only) | | |
| Valid Settings | 0 = OFF 1 = ON | | |
| Basic/Advanced | Basic | | |

Description

This parameter controls whether or not a Shadowbase Cached TRS or DW will parse target values to be inserted into numeric columns to ensure that the values will actually be accepted into the numeric columns. If a column value is determined to be non-numeric, this parameter will ensure that the proper SAP ASE error is returned, which allows Invalid Num Processing to work properly as well as ensuring data integrity.

Notes

This parameter is necessary because SAP ASE (formerly Sybase) automatically converts invalid numeric data into a `0` instead of throwing the expected error. This parameter is a workaround to ensure that inserting data into numeric columns works the same with SAP ASE as it does with other supported ODBC databases.

SHAD_SHADTRS_DUP_TX_LOGGING

| | | | |
|-------------------------|-------------------|----------------|-----------------------------------|
| Initial Version: | V6.600 | Syntax: | SHAD_SHADTRS_DUP_TX_LOGGING=<NUM> |
| Last Change: | -- | | |
| Default Value | 1 (ON) | | |
| Other Servers | TRS | | |
| Valid Settings | 0 = OFF 1 = ON | | |

Description

This parameter controls whether or not a Shadowbase TRS produces a message when it notices a duplicated transaction. The default setting is on, the TRS will produce the skipping duplicated transaction message.

Notes

Prior to Version 6.600, the default behavior was to always log this message. It has been changed to allow for setups that may have duplicated transactions as a side-effect due to how the availability setup is configured in a clustered environment.

SHAD_SKP_INF_ERROR

| | | | |
|-------------------------|-------------------|----------------|---------------------------|
| Initial Version: | V6.600 | Syntax: | SHAD_SKIP_INF_ERROR=<num> |
| Last Change: | -- | | |
| Default Value | 0 (OFF) | | |
| Other Servers | TRS, DW | | |
| Valid Settings | 0 = OFF 1 = ON | | |

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.

Notes

Prior to Version 6.600, the default behavior was to log each of the problems encountered, and skip these types of transactions. 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_TRANS_EXPECTED_ENDS

| | | | |
|-------------------------|--|----------------|--------------------------------|
| Initial Version: | pre v2.600 | Syntax: | SHAD_TRANS_EXPECTED_ENDS=<num> |
| Last Change: | V6.600 | | |
| Default Value | 1 | | |
| Other Servers | DOC | | |
| Valid Settings | -1: Do not explicitly set the expected ends. 1 - 32767: Explicitly set the expected ends. | | |

Description

* Note, the current change is modifying the default from ‘-1’ to ‘1’. This will reduce replication overhead in most situations.

This parameter allows the expected ends to be explicitly set. Explicitly setting or overriding the expected ends in certain situations permits the user to manage the COMMIT level processing to be completed before a message is marked ready to be replayed out of a DOC by a TRS process. The situations described below provide examples of when the user might want to or need to override the default expected ends behavior.

When replicating from a single NonStop collector through multiple consumer threads to a single Shadowbase Other Server platform, the SHAD_TRANS_EXPECTED_ENDS may need to be set. This will occur when one or more consumers under the same HPE NonStop Shadowbase collector are configured for replicating to the same HPE Shadowbase Other Server target (the NonStop CONSUMER NETADDRESS and NETBACKUPADDRESS parameters are the same for more than one consumer):

- If the processes on the Shadowbase Other Server target are all configured as DIRECT WRITERS, do not alter the Shadowbase Other Server target parameter SHAD_TRANS_EXPECTED_ENDS for these corresponding objects.
- If the processes on the Shadowbase Other Server target are configured as single port DOC WRITERS, or are a combination of single port DOC WRITERS and DIRECT WRITERS, do not alter the other server target parameter SHAD_TRANS_EXPECTED_ENDS for these corresponding objects.
- If there is only one process configured on the Shadowbase Other Server target server and it is a multi-port DOC WRITER (a single DOC WRITER is configured with multiple PORTS), do not alter the other server target parameter SHAD_TRANS_EXPECTED_ENDS for this DOC WRITER.
- If the processes on the Shadowbase Other Server target are configured as a combination of DIRECT WRITERS and single port DOC WRITERS, along with a multi-ported DOC WRITER, the

SHAD_TRANS_EXPECTED_ENDS for the single-ported DOC WRITER and the DIRECT WRITER(s) must not be altered. The SHAD_TRANS_EXPECTED_ENDS for the multi-ported DOC WRITER must be set to equal the number of PORTS for this object (i.e. if DOC WRITER has 4 ports, SHAD_TRANS_EXPECTED_ENDS=4 should be configured).

Multi-ported DOC writers require special processing, and typically require the SHAD_TRANS_EXPECTED_ENDS parameter to be set to the number of ports writing into the common DOC.

Please consult the *HPE NonStop Shadowbase Operations Manual* for more information related to this parameter (specifically the parameter SBCOLLMULTIENDS write up).

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 (V6.300).

Previous query did not return all records at the last location (SCN <scn> RBASQN <rbasqn> RBABLK <rbablock> RBABYTE <rbabyte> SNN [<scn>])
This may have a caused a SQL mismatch error message in previous versions.

Cause: HPE Shadowbase for Oracle Log Collection detected a rare condition that might have a caused an SQL mismatch error message in versions before this version.

Note: this message confirms that the OLS collector is properly handling a rare condition.

Effect: None – informational only.

Recovery: None required.

"Unable to obtain nullable status of columns from DB2; Will attempt to replace invalid value with NULL."

Cause: If Invalid Num Processing is enabled for TRSs and DWs replicating into DB2 target tables, this message will appear during the replacing of the invalid numeric value with a NULL value. It is just to explain that OSB cannot know if inserting a NULL will be allowed before attempting the insert.

Effect: None – informational only.

Recovery: None required.

SHAD_OPCOLLECT_FILTER_PARTIAL_TXS set to [0]; Transactions without a BEGIN event will be processed
SHAD_OPCOLLECT_FILTER_PARTIAL_TXS set to [1]; Transactions without a BEGIN event will not be processed

Cause: User specified settings accepted for parameter. (informational)

Effect: The process reports the condition, and continues.

Recovery: None, unless parameter setting unintended. A restart will be required if this is the case.

SHAD_OPCOLLECT_FILTER_PARTIAL_TXS setting [<value>] invalid;
Transactions without a BEGIN event will not be processed

Cause: User specified settings are incorrect for parameter, default value is used.

Effect: The process reports the condition, and continues.

Recovery: Correct the error and restart.

SHAD_OPCOLLECT_SELECT_SB_EVENTS set to [0]; Shadowbase events will be filtered by the SELECT statement.

SHAD_OPCOLLECT_SELECT_SB_EVENTS set to [1]; Shadowbase events will be selected in the SELECT statement

Cause: User specified settings accepted for parameter. (informational)

Effect: The process reports the condition, and continues.

Recovery: None, unless parameter setting unintended. A restart will be required if this is the case.

SHAD_OPCOLLECT_SELECT_SB_EVENTS setting [<value>] invalid;
Shadowbase events will be filtered in the SELECT statement

Cause: User specified settings are incorrect for parameter, default value is used.

Effect: The process reports the condition, and continues.

Recovery: Correct the error and restart.

SHAD_OPCOLLECT_LOG_PARTIAL_TXS set to [ON]; Transactions without BEGIN events will be logged after initial startup.

SHAD_OPCOLLECT_LOG_PARTIAL_TXS set to [ALL]; All transactions without BEGIN events will be logged.

SHAD_OPCOLLECT_LOG_PARTIAL_TXS set to [OFF].

Cause: User specified settings accepted for parameter. (informational)

Effect: The process reports the condition, and continues.

Recovery: None, unless parameter setting unintended. A restart will be required if this is the case.

SHAD_OPCOLLECT_LOG_PARTIAL_TXS setting [<value>] invalid: Expected 'ON', 'OFF', or 'ALL'

Transactions without BEGIN events will be logged after initial startup.

Cause: User specified settings are incorrect for parameter, default value is used.

Effect: The process reports the condition, and continues.

Recovery: Correct the error and restart.

SHAD_OPCOLLECT_LARGEST_TX_SPAN set to [<value>]

Cause: User specified settings accepted for parameter. (informational)

Effect: The process reports the condition, and continues.

Recovery: None, unless parameter setting unintended. A restart will be required if this is the case.

SHAD_OPCOLLECT_LARGEST_TX_SPAN setting [<value>] invalid; using default value of 2000 SCNs

Cause: User specified settings are incorrect for parameter, default value is used.

Effect: The process reports the condition, and continues.

Recovery: Correct the error and restart.

SHAD_OPCOLLECT_TEST_ONLY_SKIP_START_EVENTS enabled, start events will be skipped to allow testing of partial transactions.

Cause: User specified settings accepted for parameter. (informational)

Effect: The process reports the condition, and continues.

Recovery: None, unless parameter setting unintended. A restart will be required if this is the case.

Partial transaction found: <event type> event read for transaction [<transaction-id>] at SCN [<scn>] RBASQN <rbaqn> RBABLK <rbablk> RBABYTE<rbabyte> SSN [<ssn>,<instance#>]; START event for transaction has not been read.

Cause: Partial transaction logging is enabled and a partial transaction (transaction without a start event) was found.

Effect: The process reports the condition, and continues.

Recovery: Information only.

Known Problems Remaining in Version 6.600

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. This issue will be addressed in an upcoming release.
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 (*sborrow*) has the following limitations:
- a. HPE Shadowbase Log Source for Oracle collection does not support replication of tables with columns defined as `TIMESTAMP WITH TIMEZONE` or `TIMESTAMP WITH LOCAL TIMEZONE`.
 - b. HPE Shadowbase Log Source for Oracle does not support replication of tables with the following column types:
 - BLOB
 - CLOB
 - NCLOB
 - BFILE
 - RAW
 - c. 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.
 - d. Shadowbase Log Source for Oracle currently only collects DML events (Inserts, Updates, and Deletes).
 - e. There is an issue handling delete operations that have a date as part of the primary key field. Contact Shadowbase Support if any of the tables to be replicated include a data field in the primary key.
 - f. HPE Shadowbase Log Source for Oracle does not support collecting tables with table or column names that are reserved words or that require quoting to be resolved.
 - g. 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.
 - h. HPE Shadowbase Log Source for Oracle does not support collecting tables returned by Oracle LogMiner using the `UNISTR` function string.

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. When configuring HPE Shadowbase Log Source for Oracle collection with an Oracle password that is set to expire soon, the configuration fails. Oracle will issue a message when the initial connection is completed, which causes the script to fail. The sample below illustrates the issue:

```

* * * * *
*                               Shadowbase Open Server                               *
*                               Oracle Log Based Source Collection                     *
*                               Configuration and Maintenance                         *
*                                                                                   *
* Copyright 2014 Gravic, Inc.                                                       *
* All Rights Reserved.                                                             *
* SBSUPPORT@GRAVIC.COM                                                            WWW.GRAVIC.COM *
* * * * *
                                M A I N - M E N U

1. Install Open Collection Filtering Tables
2. Remove Open Collection Filtering Tables
3. Install Shadlogpack Oracle Package
4. Remove Shadlogpack Oracle Package
5. Add to Filtering Tables
6. Update Filtering Tables
7. Delete from Filtering Tables
8. Display Configuration Settings
9. Check an Oracle User's Permissions
10. V&V Filtering Tables
11. V&V Supplemental Logging

Enter your choice [1-10, Q to quit]:1

A connection to an Oracle DBMS Instance has not yet been
established
Please specify the following Oracle Connection Information

User Name: qasource
Password:
Instance: ORA11R2
ERROR:
ORA-28002: the password will expire within 7 days
Press any key to continue

```

The corrective action is to remove the condition that indicated in the warning message – in this case, to set a new password.

13. 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.
14. A configuration using a consumptive Direct Writer (a Direct Writer that is not connected to a database) connected to an HPE NonStop system as a source is not supported. If you need to use a consumptive process, you must either setup a configuration that replicates from the NonStop to a Doc Writer, and then uses a consumptive TRS; or use a Direct Writer that does connect to the database with your consumptive user exit.
15. 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.
16. 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.

17. 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.

18. 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`.

19. HPE Shadowbase Log Source for Oracle will generate an incorrect `UPDATE` statement if an update is performed against a table with only key columns defined and the update does not actually change any of the key columns. For example, consider the following SQL commands:

```
create table test3 (kcol1 int not null,
                  kcol2 int not null,
                  primary key (kcol1,kcol2));

insert into test3 values (9,2);

update test3 set kcol1=9;
```

Shadowbase will generate an invalid update statement for the third update – there will be no set values:

```
UPDATE TEST3 SET WHERE KCOL1='9' AND KCOL2='2';
```

If the update statement modified one of the columns, Shadowbase would have processed the event as specified by the `SHAD_OPACCOLLECT_PRIKEY_UPDATE_OPTION`.

If this occurs in your environment, there is a work-around. You can safely ignore `UPDATE` statements using DBS mapping on the target for tables with only key columns. The collector will either stop, skip, or convert updates that change column values into a delete followed by an insert, depending on the setting of the `SHAD_OPACCOLLECT_PRIKEY_UPDATE_OPTION`. The only updates that will be stored in the DOC for processing by the TRS or TFS are ones that do not change any data, which do not need to be replicated.

20. 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.

21. 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.
22. 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.
23. 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 '% ';
```

24. 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.

25. Special considerations are required when using a TFS as a relay server. 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.

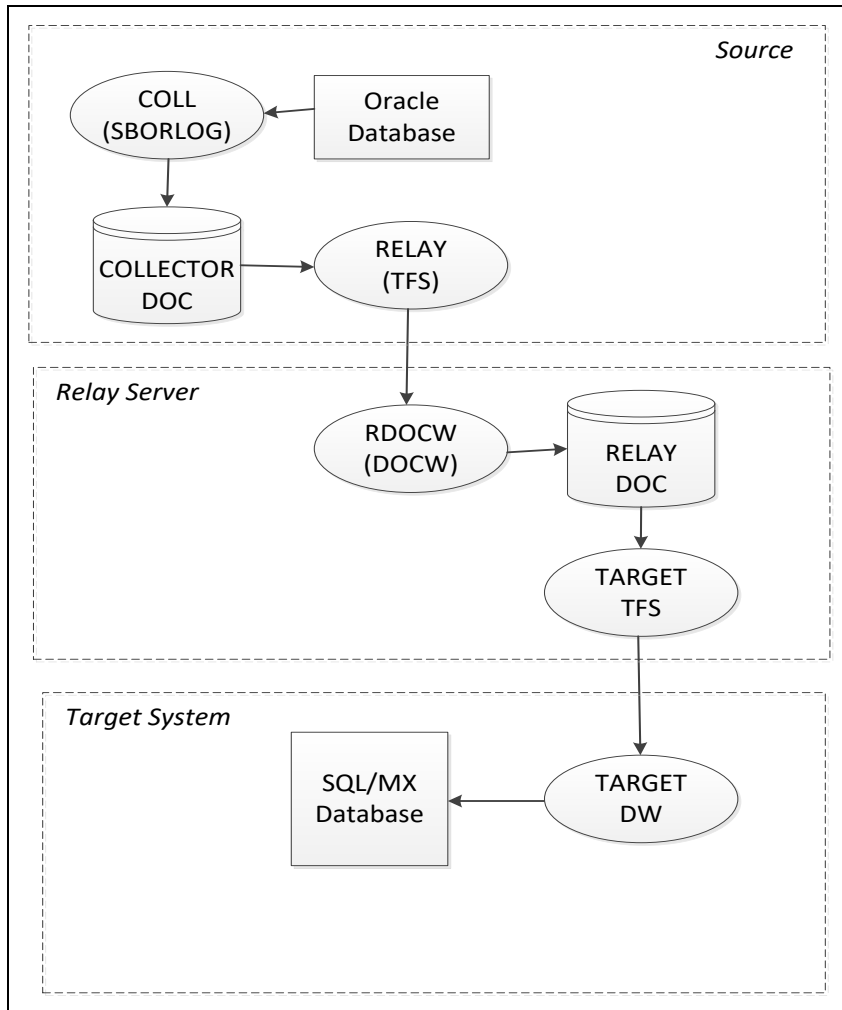


Figure 1 - Sample Relay Server Configuration

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 duplicated data from being replayed.

26. 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.
27. When replicating to and from 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.

28. 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 another source SQL type such as SQLSBV1.
29. 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.
30. There is an issue related to DB2 DWs and TRSs running on RHEL and SuSE Linux where if the data being replicated is diverse enough to cause the OSB Statement Cache to frequently churn, memory use of the target DWs and TRSs will grow until the system is out of memory. The workaround for replicating into DB2 is to configure unixODBC Data Source Names (DSNs) for the relevant DB2 servers and to configure the target DWs and TRSs as ODBC objects using those DSNs. The memory growth found in the DB2 objects is not present in the ODBC objects.
31. 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.
32. There is an issue with OSB's INVALID NUM On-Error processing with the two specific Oracle SQL types `BINARY_FLOAT` and `BINARY_DOUBLE`. Due to current limitations of the Oracle Call Interface, these types are reported to Shadowbase as being standard character data types and thus cannot be replaced by On-Error processing in the event invalid data is inserted into a column of either type.
33. 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

34. 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.
35. There are certain cases where when running a TRS or other HPE Shadowbase object from a service on Microsoft Windows Server 2019 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.
36. There is an issue with invalid num processing on statements going into Microsoft SQL Server where if the field containing an invalid numeric value is a `bit` type, OSB invalid numeric processing is not currently able to replace it.
37. 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.
38. There is an issue with the Windows DB2 ODBC driver version 11.05.00.1077 which prevents ODBC objects from connecting to a DB2 instance as usual. The current workaround is to use a DB2 ODBC driver with a version less than 11.05.
39. There is an issue when editing a DB2 collector configuration. The prompt for IBM MQ username does not register the previously entered username and is followed by two consecutive prompts for the IBM MQ password. The response to the username prompt is overridden by the response to the first password prompt. The second password prompt behaves correctly as the password.
40. 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.

Installation Instructions

Please follow the installation instructions included in the README.<platform>.<version>.TXT file that accompanies this release.

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