

## SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	HPE Shadowbase for Other Servers
Release:	Gravic Version 6.220 T1123-AAC (SB REPL/OSS) T1124-AAB (SB REPL/HP UX) T1125-AAB (SB REPL/SUSE LINUX) T1126-AAB (SB REPL/RHEL LINUX) T1127-AAB (SB REPL/SOLARIS X86) T1128-AAB (SB REPL/SOLARIS) T1129-AAB (SB REPL/MS WINDOWS) T1130-AAB (SB REPL/AIX)
Date:	March 31, 2016
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2016 ( <a href="http://www.gravic.com">www.gravic.com</a> )
File Name:	IPM6220_other_servers.pdf

NOTE: This softdoc covers new features and corrected problems for HPE Shadowbase for Other Servers, Version 6.220. It is available as an Adobe PDF file (.PDF). Copies of the PDF file reader can be freely downloaded from [www.adobe.com](http://www.adobe.com)

## Table of Contents

<b>Table of Contents</b>	<b>1</b>
<b>Disclaimer</b>	<b>3</b>
<b>Special Notes for Version 6.220</b>	<b>4</b>
<b>New Features in Version 6.220</b>	<b>5</b>
Password Obfuscation	8
Named Transactions for Oracle Databases	8
Reading from Online, Archived, and Offline Archived REDO Logs	11
Oracle LogMiner Continuous Mining Modes	14
Generating the Oracle Dictionary	15
Enhanced Support for SQL/MX Tables with SYSKEYS	15
Specifying the Starting Position for HPE Shadowbase Log Source for Oracle Collection	16
<b>Problems Corrected in Version 6.220</b>	<b>17</b>
<b>New &amp; Modified shadparm.ini Parameters</b>	<b>27</b>
SHAD_HEX_CONVERT_LEVEL	27
SHAD_OPCOLLECT_ADJUST_END_SCN	27
SHAD_OPCOLLECT_ADJUST_START_SCN	28
SHAD_OPCOLLECT_ARCHIVE_DIR	28
SHAD_OPCOLLECT_ARCHIVE_FILES	30
SHAD_OPCOLLECT_ARCHIVE_CATALOG	31
SHAD_OPCOLLECT_CHECK_LOGMNR_SEQUENCE	31
SHAD_OPCOLLECT_DDL_TRACKING	32
SHAD_OPCOLLECT_END_SCN	33
SHAD_OPCOLLECT_INVALID_SQL_OPTION	33
SHAD_OPCOLLECT_MINING_MODE	34
SHAD_OPCOLLECT_PARSE_ERROR_OPTION	35
SHAD_OPCOLLECT_PATH_SEPARATOR	36
SHAD_OPCOLLECT_QUERY_EXECUTION_LIMIT	36
SHAD_OPCOLLECT_RECYCLE_CONNECTION	37
SHAD_OPCOLLECT_RESET_SCN	37
SHAD_OPCOLLECT_SCAN_FOR_NULL_BYTES	38
SHAD_OPCOLLECT_SKIP_CORRUPT_REDO_BLOCKS	39
SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT	39
SHAD_OPCOLLECT_USE_ORDER_BY	40
SHAD_OPCOLLECT_USE_SEG_NAME	40
SHAD_OVERRIDE_SYSKEY	41
SHAD_STATUS_STATS	42
SHAD_NSL_CONVERT_LEVEL (obsolete)	43
<b>New &amp; Modified User Messages</b>	<b>44</b>
Logged for Null (0) Bytes in Text Fields	44
Logged for Out of Order Events	45
Logged for SHAD_OPCOLLECT_ARCHIVE_DIR Parameter Errors	45
Logged for SHAD_OPCOLLECT_ARCHIVE_FILES Parameter Errors	46
Logged for Unmatched Rollback Events	46
Logged when the End SCN is Reached	47
Logged when Reading from Archived REDO Logs is Complete	48
Logged when Selecting an Invalid Mode for Mining	48

Logged for Cursor Allocation Issues _____	49
Logged for HPE Shadowbase Log Source for Oracle Parsing Issues _____	50
Logged for Invalid HPE Shadowbase Log Source for Oracle Collection Parameters _____	54
Logged for Missing Data _____	55
<b>Known Problems Remaining _____</b>	<b>56</b>
<b>Installation Instructions _____</b>	<b>65</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.

## **Special Notes for Version 6.220**

1. Version 6.220 is a “general availability” (non-restricted TCF) release. The previous non-restricted TCF release is version 6.101 (HPE version AAA) for all servers except OSS, and version 6.101 (HPE version AAB) for HPE OSS systems.

The release includes a number of corrections and enhancements from the following restricted TCD/TCF releases (previously available through Gravic only):

- Versions 6.110-6.110K (Linux)
  - Version 6.205 (OSS)
  - Version 6.206 (OSS)
1. 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.
  2. 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 6.220 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. If you are upgrading from a version prior to 6.220 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 need an obfuscated configuration with a prior release, you will need to drop and re-add the objects using the prior version of SBMON.

3. 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.

## ***New Features in Version 6.220***

This section provides a summary of the features added to HPE Shadowbase for Other Servers products since the previous general availability release (V6.101 for all servers except HPE NonStop OSS; V6.200 for HPE NonStop OSS).

1. This release increases the maximum number of prepared statements that the Direct Writer (DW) and the Transaction Replay Server (TRS) can cache from 50 to 1000 for Oracle, DB2, SQL/MX, MySQL, and SQL Server. The maximum number remains at 50 for Sybase databases. See the SHAD\_MAX\_CURSORS description, below, for more information.

These processes issue a PREPARE statement command to compile the statement prior to applying the insert, update, or delete operation using the EXECUTE statement. The processes save prepared statement in cache and reuse the statement without requiring the prepare step when a similar operation is received. Since the PREPARE statement takes a relatively long time to execute, reusing the statement can cause a dramatic increase in throughput and processing efficiency.

When the number of different statements exceeds the cache limit, the processes select the least recently used (LRU) statement from the cache to close. That statement's resources are freed and the statement is removed from cache. The new statement is then prepared and inserted in the cache.

If there are a large number of tables being replicated (or a large number of different statements being generated due to user exits or audit compression), statements may be swapped in and out of cache frequently, resulting in many prepares and much lower throughput. This change dramatically increases the number of tables that a single DW or TRS can support without swapping.

For more information on Shadowbase's statement caching feature, please see the relevant section in the *HP Shadowbase Other Servers Target Installation and Configuration Manual*.

2. HPE Shadowbase Log Source for Oracle collection now supports Oracle 12c databases. Prior to this release, the HPE Shadowbase Log Source for Oracle collector did not free certain resources while reading from the Oracle log. This caused the Oracle driver process to exhaust PGA memory in Oracle 12c and to shutdown, also stopping the HPE Shadowbase Log Source for Oracle collector. Shadowbase has been enhanced to explicitly free resources.
3. Some of the configuration data stored in the COLLCONFIG files, including passwords entered through SBMON, has been obfuscated and is no longer stored in clear text. **Note:** records in the COLLCONFIG will be automatically obfuscated when they are saved by a Version 6.200 Shadowbase process. Once obfuscated, the

configuration information will no longer be readable by prior releases. If you want to be able to downgrade back to a previous version of Shadowbase, you must keep a copy of the collconfig.dat and collconfig.idx files located in the data subdirectory of your installation directory. See *Password Obfuscation*, below, for more information.

4. Shadowbase now supports named transactions for Oracle. This functionality was primarily added for bi-directional replication, but can also be useful for uni-directional replication. When named transactions are enabled, the Direct Writer and Transaction Replay Servers will generate a name for each transaction and will apply the events to the database under the assigned name. HPE Shadowbase Log Source for Oracle will skip any transactions that match the name, allowing for bi-directional cutoff and preventing ping-ponging of events. Database administrators can also use the name to identify transactions and data applied by the Shadowbase processes. See *Named Transactions for Oracle Databases*, below, for more information.
5. HPE Shadowbase Log Source for Oracle can now use an exported flat-file SQL catalog when reading from REDO logs. While this functionality was added to support reading archive logs on foreign systems, it also can be used for processing from REDO logs residing on the database server. See *Generating the Oracle Dictionary*, below, for more information on how to generate the flat-file catalog.
6. HPE Shadowbase Log Source for Oracle can now be configured to stop at a specific SCN (Oracle's System Change Number). See the SHAD\_OPCOLLECT\_END\_SCN parameter, below, for more information.
7. HPE Shadowbase Log Source for Oracle can optionally force the sequence of events returned by the SELECT statements for Oracle LogMiner data. See the SHAD\_OPCOLLECT\_USE\_ORDER\_BY parameter, below, for more information.  
**Note:** enabling the ordering has a significant impact on performance and is not required for correct collection of data. Only enable this parameter if directed to by Support.
8. HPE Shadowbase Log Source for Oracle can be configured to check the sequence of LOGMNR event data. This functionality is intended for diagnostic purposes only and is not needed for normal operations. Only enable this parameter if directed to by Support. See the SHAD\_OPCOLLECT\_CHECK\_LOGMNR\_SEQUENCE parameter, below, for more information.
9. A new parameter has been added to HPE Shadowbase Log Source for Oracle to enable skipping of corrupted redo blocks. When enabled, this command calls the Oracle procedure DBMS\_LOGMNR.START\_LOGMR with the DBMS\_LOGMNR.SKIP\_CORRUPTION option enabled. See the description under SHAD\_OPCOLLECT\_SKIP\_CORRUPT\_REDO\_BLOCKS for more information.
10. A new parameter, SHAD\_OPCOLLECT\_ADJUST\_END\_SCN, has been added to HPE Shadowbase Log Source for Oracle to limit how close to the end of the Oracle

REDO logs the collector will collect. See the description under `SHAD_OPCOLLECT_ADJUST_END_SCN` for more information.

11. HPE Shadowbase Log Source for Oracle now supports four modes of collecting events from Oracle databases using log mining:
  - a. Continuous collection from the online REDO logs. This is the mode originally supported by Shadowbase and is the default if no mode is specified. Shadowbase will continuously collect from the online REDO logs only.
  - b. Continuous collection from both the online and archived REDO logs. This mode is available only on databases with archiving enabled. It makes use of the Oracle LogMiner continuous mining option to mine data from the online and archived REDO. See *Reading from Online, Archived, and Offline Archived REDO Logs*, below for more information on this new functionality.
  - c. Continuous collection from archived REDO logs only. This mode is similar to collecting from both the online and archived logs in b) above, however, the highest SCN collected is restricted to the highest available SCN in the archived logs. See *Reading from Online, Archived, and Offline Archived REDO Logs*, below for more information on this new functionality.
  - d. Offline collection from archived REDO logs. HPE Shadowbase Log Source for Oracle reads from a specified list of archived REDO files. The files can be located on the database server that created them, or can be copied to another Oracle database server and processed there. Unlike the continuous modes, Shadowbase will stop collection once the archived data has been read/processed. See *Reading from Online, Archived, and Offline Archived REDO Logs*, below for more information on this new functionality.

A new parameter, `SHAD_OPCOLLECT_MINING_MODE` has been added to specify the collection mode. See the description under `SHAD_OPCOLLECT_MINING_MODE` for more information.

12. Shadowbase now supports replicating SQL/MX tables with SYSKEYs while maintaining the same SYSKEY value on the target as on the source. Specifically, this includes file system defined SYSKEYs that result from key-sequenced tables that are defined with “CLUSTERING” keys, as well as (key-sequenced) tables that have no primary key definition. In both cases, the file system creates and populates a SYSKEY column as part of the primary key. See *Enhanced Support for SQL/MX Tables with SYSKEYS*, below, for more information.
13. A new parameter has been added to Shadowbase for Other Servers to better handle text strings with binary data. See the description under `SHAD_HEX_CONVERT_LEVEL` for more information.



14. You can now configure HPE Shadowbase Log Source for Oracle to start collection at the end of the file. See *Specifying the Starting Position for HPE Shadowbase Log Source for Oracle* Collection, below, for more information.
15. You can now configure HPE Shadowbase Log Source for Oracle to start mining with the `DBMS_LOGMNR.DDL_DICT_TRACKING` option. With this option, Oracle log mining will track DDL changes and automatically adjust its internal dictionary when a table's schema changes. However, there are a number of restrictions within Oracle log mining as to when this option can be used, and it is not recommended for normal replication. Please see the `SHAD_OPCOLLECT_DDL_TRACKING` shadparm.ini parameter (below) for more information.

## Password Obfuscation

Previous versions of Shadowbase stored configuration data entered through SBMON, including passwords, in clear text in the COLLCONFIG configuration files. Version 6.220 will obfuscate the password and other configuration data to reduce the security vulnerability associated with storing the information in clear text. Shadowbase has also been modified to eliminate the display of the password in error messages and when displaying configuration data. Note that this only applies to data at rest on the disk.

This version can read and use the clear text COLLCONFIG files created and modified by previous versions of Shadowbase. However, reverse is not true: previous versions of Shadowbase cannot use COLLCONFIG files created or modified by this version. If you are upgrading from a previous version of Shadowbase and want to maintain the capability to roll back to that version, you will need to make a copy of the collconfig.dat and collconfig.idx files in data subdirectory of the installation directory. Otherwise, you will need to delete and re-enter the Shadowbase objects if you downgrade.

The new version automatically obfuscates the data whenever it updates the COLLCONFIG records. This happens when a new object is added through SBMON, or when an existing record is edited with SBMON. It can also happen as Shadowbase runs as the COLLCONFIG records also contain status and restart information and are updated by the running processes.

If you want to ensure the passwords are obfuscated after Version 6.220 is installed, edit and each object in the configuration. The COLLCONFIG record will be saved even if there are no changes to its configuration.

## Named Transactions for Oracle Databases

Shadowbase for Other Servers now supports naming transaction for Oracle database. A new parameter, `SHAD_TX_NAME`, has been added to enable this feature and to specify the prefix for the transaction name.

When you specify this parameter for a Direct Writer (DW) or Transaction Replay Server (TRS), the DW or TRS will name the transactions using the value as a prefix. The format of the transaction name is:

```
<shad_tx_name_value>.<object_name>.<counter>
```

Where:

<shad\_tx\_name\_value> is the value specified for SHAD\_TX\_NAME

<object\_name> is the name of the DW or TRS

<counter> is a sequentially increasing counter

For example, if the SHADPARM.INI file contains SHAD\_TX\_NAME=SB and the TRS is named TRS1, the transactions will be named:

```
SB.TRS1.1
```

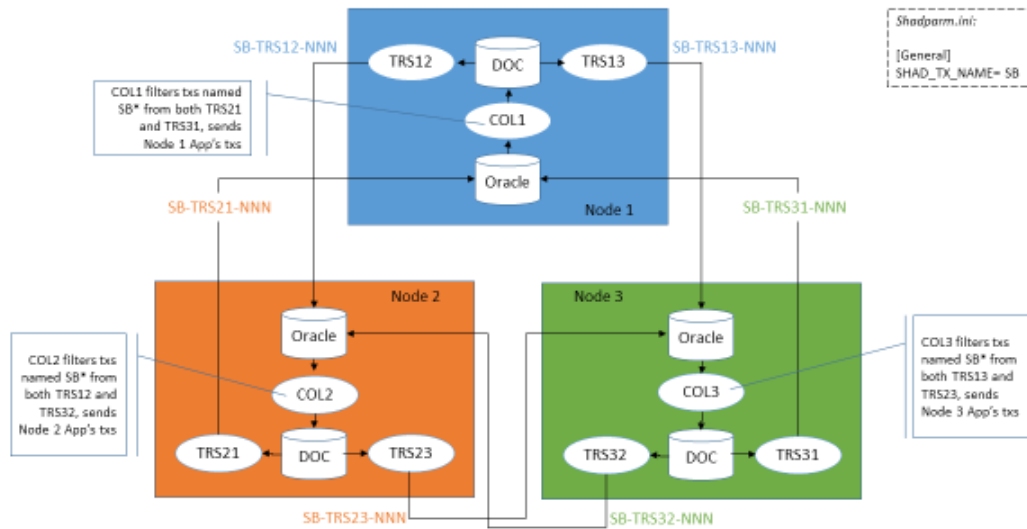
```
SB.TRS1.2
```

```
...
```

You can use the transaction name to determine which transactions have been applied by Shadowbase. For example, you can query the Oracle v\$transaction view to see the active transactions and use the name to identify which transactions belong to which Shadowbase object.

If you specify SHAD\_TX\_NAME for an HPE Shadowbase Log Source for Oracle collector, the collector will not collect any data applied under a transaction name starting with the value specified. This is used in a bi-directional replication environment to prevent data applied by Shadowbase from being replicated back to the original source system (i.e., to avoid data oscillation). Specifying the same prefix as the TRS and DW's applying into the database will prevent replicated data from being replicated back to the originating system.

Figure 1 - Multi-Node Configuration, Round Robin (below) shows one possible configuration of active-active replication in a three node network. In this configuration, each node replicates the data generated by the local application to two other nodes and the collector filters and does not replicate any data generated by replication. The configuration is simple: all transactions applied by replication begin with "SB" and all transactions that begin with "SB" are filtered from replicating. Just adding the SHAD\_TX\_NAME=SB parameter to the general section of shadparm.ini implements this configuration.



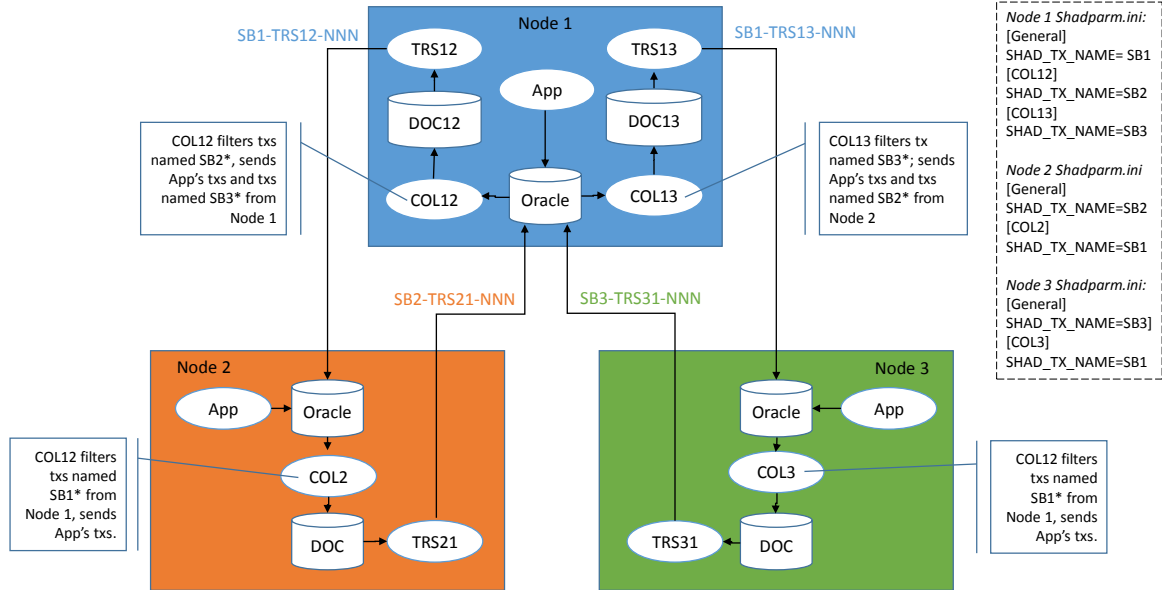
**Figure 1 - Multi-Node Configuration, Round Robin**

Figure 2 - Multi-Node Configuration Pass-Through, below, illustrates a slightly more complicated configuration for bi-directional replication with three nodes. In this configuration, there is no direct replication between Nodes 2 and 3. Instead, both replicate back to Node 1. Node 1 replicates both the application database changes and the Node 2 replication changes to Node 3 and replicates the application database changes and the Node 3 replication changes to Node 2.

When the application makes a database change on Node 2, it is read by the collector (COL2) and placed in the DOC. TRS21 reads it from the DOC applies it to the Node 1 database, using the name SB2-TRS21-nnnn. There are two collectors running on Node 1, COL12, which collects changes to be replicated to Node 2; and COL31, which collects changes to be replicated to Node 3. COL12 is configured to ignore database changes applied under transaction names starting with “SB2”, so the database change applied by replication from Node 2 is filtered and not sent back to Node 2. COL3, however, is configured to ignore changes applied under transaction names start with “SB3” and will collect the database change applied by replication from Node 2 (which starts with “SB2”) for replication to Node 3. TRS13 will read the change and apply it on Node 3.

Note that the configuration in the shadparam.ini files is more complex in this case, as we want to pass through some replicated changes on Node 1 from the other nodes. Node 2 needs to be configured to apply changes to Node 1 using transaction names starting with “SB2”, Node 3 needs to be configured to apply changes to Node 1 using transaction names starting with “SB3”. Both Nodes 2 and 3 are configured to filter the replicated changes from Node 1 by ignoring changes in transaction names starting with “SB1”.

Node 1 is configured to apply changes to Nodes 2 and 3 under transaction names start with “SB1”. Each of the collectors needs to be configured to filter replicated changes from the Node it is replicating to and to pass through replicated changes from the other node.



**Figure 2 - Multi-Node Configuration Pass-Through**

## Reading from Online, Archived, and Offline Archived REDO Logs

HPE Shadowbase Log Source for Oracle now supports reading and replay data from archived Oracle REDO logs. In this release, Shadowbase can be configured to read and process audit information from:

- Online REDO logs (CONTINUOUS\_ONLINE mining mode): Only the online REDO logs are accessed. Archiving does not need to be enabled on the database.
- Both online and archived REDO logs (CONTINUOUS\_BOTH mining mode): In this mode, the “CONTINUOUS MINING” option for log mining is used to read events from both the online and available archived REDO logs. Note that for continuous mining to be used, archiving must be enabled for the database.
- Archived REDO logs (CONTINUOUS\_ARCHIVE mining mode): This mode uses the “CONTINUOUS MINING” option for Oracle log mining as well. However, the range of events retrieved is limited to those that are archived based upon information in the V\$ARCHIVED\_LOG view. The collector will continue

to run and will read additional data when it is archived by Oracle.

- **Offline, archived REDO logs (OFFLINE\_ARCHIVE mining mode):** In this mode, the REDO logs can either be on the same database server that generated them, or they can be transferred to another database server (such as a development server) running the same version of Oracle for processing. This mode can be used to replay old archived data that is no longer available to the database. In this mode, the collector will read events from the logs specified at startup and will stop when it has processed all logs.

**Note:** The Shadowbase collector needs the schema for the replicated tables to generate the replication stream. If the archived logs are on the same Oracle database server that generated them, Shadowbase will use the schema from those source tables. If you have copied the archived logs to a different database server, the source tables will have to be recreated in that Oracle database so that Shadowbase can access the schema for the tables.

There are several shadparm.ini parameters that control processing of archived REDO logs:

- **SHAD\_OPCOLLECT\_MINING\_MODE:** specifies the mining mode as described above.
- **SHAD\_OPCOLLECT\_ARCHIVE\_DIR:** If present, it specifies the directory containing the archived REDO logs. If omitted, HPE Shadowbase Log Source for Oracle uses only the online REDO logs. Must be present to read from the archive logs. Note: unless SHAD\_OPCOLLECT\_ARCHIVE\_FILES is specified, all files in the directory will be added for processing.
- **SHAD\_OPCOLLECT\_ARCHIVE\_FILES:** If present and SHAD\_OPCOLLECT\_ARCHIVE\_DIR is specified, specifies a text file that contains the names of the archived REDO logs to be added for processing. Each REDO log should be on its own line, and the name can either be fully qualified or just the file name within the SHAD\_OPCOLLECT\_ARCHIVE\_DIRECTORY. SHAD\_OPCOLLECT\_ARCHIVE\_FILES should be used if:
  - HPE Shadowbase Log Source for Oracle is running on a different system from the Oracle database server.
  - SHAD\_OPCOLLECT\_ARCHIVE\_DIR contains either subdirectories or files that are not archived REDO logs.
  - HPE Shadowbase Log Source for Oracle should process only a subset of the archived REDO logs in SHAD\_OPCOLLECT\_ARCHIVE\_DIR.
- **SHAD\_OPCOLLECT\_ARCHIVE\_CATALOG:** If present, specifies a flat file containing the archive catalog. If omitted, the online catalog will be used. Use this parameter if you are processing the REDO logs in an Oracle database server that

is different from the one that created the logs.

- **SHAD\_OPCOLLECT\_PATH\_SEPARATOR:** If present, specifies the directory path separator character (‘/’ for Unix and Linux file systems, ‘\’ for Windows) used by the Oracle database server. If absent, the path separator that is used by the system HPE Shadowbase Log Source for Oracle is running on is used to build the log file names.

The following scenarios for reading **offline archived** REDO logs (SHAD\_OPCOLLECT\_MINING\_MODE=ARCHIVE\_OFFLINE) may help clarify which parameters you need to specify when configuring HPE Shadowbase Log Source for Oracle:

- 1) The archive files are on the database server that generated the audit:
  - **SHAD\_OPCOLLECT\_ARCHIVE\_DIR** – required.
  - **SHAD\_OPCOLLECT\_ARCHIVE\_FILES:**
    - Required if Shadowbase is running on a different server.
    - Required if the archive directory contains files that are not to be processed.
    - Required if the archive directory contains subdirectories.
    - Omitted otherwise.
  - **SHAD\_OPCOLLECT\_ARCHIVE\_CATALOG** – omitted.
  - **SHAD\_OPCOLLECT\_PATH\_SEPARATOR:**
    - Omitted if Shadowbase is running on the same server as the database.
    - Omitted if the server Shadowbase is running on has the same directory path separator as the database server, for example, if both are running on Linux servers.
    - Required if the server Shadowbase is running on has a different directory path separator than the database server. For example, this parameter is required if Shadowbase is running on a Windows server and the database is running on a Linux server.
- 2) The archive files have been copied to a different database server from the one that generated the audit for processing:
  - **SHAD\_OPCOLLECT\_ARCHIVE\_DIR** – required.
  - **SHAD\_OPCOLLECT\_ARCHIVE\_FILES:**
    - Required if Shadowbase is running on a different server.
    - Required if the archive directory contains files that are not to be processed.
    - Required if the archive directory contains subdirectories.
    - Omitted otherwise.
  - **SHAD\_OPCOLLECT\_ARCHIVE\_CATALOG** – required.
  - **SHAD\_OPCOLLECT\_PATH\_SEPARATOR:**

- Omitted if Shadowbase is running on the same server as the database.
- Omitted if the server Shadowbase is running on has the same directory path separator as the database server, for example, if both are running on Linux servers.
- Required if the server Shadowbase is running on has a different directory path separator than the database server. For example, this parameter is required if Shadowbase is running on a Windows server and the database is running on a Linux server.

HPE Shadowbase Log Source for Oracle will process the data in the archived REDO logs from the start SCN number, specified by editing the collector's configuration through SBMON, to either the end SCN if one is specified by the SHAD\_OPCOLLECT\_END\_SCN parameter in shadparm.ini or until it reaches the end of the data in the archived REDO logs, whichever comes first. It will then stop.

## Oracle LogMiner Continuous Mining Modes

Shadowbase now supports two collection modes that use the continuous mining feature functionality from Oracle LogMiner. These modes (plus two other modes) are configured using the SHAD\_OPCOLLECT\_MINING\_MODE shadparm.ini parameter. Using the continuous mining modes, Shadowbase can be configured either collect from archived and online REDO logs (SHAD\_OPCOLLECT\_MINING\_MODE=CONTINUOUS\_BOTH), or to collect from archived REDO logs only (SHAD\_OPCOLLECT\_MINING\_MODE=CONTINUOUS\_ARCHIVE). Shadowbase can also be configured to collect from the online REDO logs only, however, this does not use the continuous mining feature of Oracle LogMiner.

The continuous mining feature of Oracle LogMiner requires that archiving be enabled for the database. In addition, Oracle LogMiner uses the information in its control records to determine which REDO logs to load and use. Shadowbase also uses that information to determine the range of SCNs that are available for collection. If archived REDO logs are deleted through operating system commands and not using RMAN, the information in the control records may not match the actual availability of the REDO logs, resulting in collection stopping errors such as:

```
2015-12-03 14:50:40 -[31760] Error: ReadLogMsgs(startMining); Experienced  
Processing Fault; sqlcode[1306]; Err[ORA-01306: dbms_logmnr.start_logmnr() must be  
invoked before selecting from V$LOGMNR_CONTENTS]
```

The Oracle utility rman should be used for archive log maintenance to insure the information contained in the Oracle control files corresponds with the files on disk.

## Generating the Oracle Dictionary

You can generate the flat file dictionary for the archived REDO logs using SQLPLUS with the DBMS\_LOGMNR\_D.BUILD procedure, specifying the STORE\_IN\_FLAT\_FILE option. The following example creates the dictionary in /home/shadtest/dictionary.ora:

```
exec dbms_logmnr_d.build(  
  'dictionary.ora', '/home/shadtest/logminer',  
  OPTIONS => dbms_logmnr_d.store_in_flat_file);
```

Note that the directory /home/shadtest/logminer must have been added in the init.ora file for Oracle as a UTL\_FILE\_DIR parameter:

```
utl_file_dir='/home/shadtest/logminer'
```

See the chapter on Using LogMiner to Analyze Redo Log Files in the Oracle Database Utilities documentation for more information.

## Enhanced Support for SQL/MX Tables with SYSKEYS

Shadowbase now supports replicating SQL/MX tables with SYSKEYs while maintaining the same SYSKEY on the target as on the source. Specifically, this includes file system defined SYSKEYs that result from key-sequenced tables that are defined with “CLUSTERING” keys, as well as (key-sequenced) tables that have no primary key definition. In both cases, the file system creates and populates a SYSKEY column as part of the primary key.

Prior to this release, SQL/MX target tables with SYSKEYs had the SYSKEY generated on the target system and required a mapping file to maintain the correspondence between the source and target systems’ key row values. (For example, for an insert on the source, the source’s file system might assign a ‘5’ for the source SYSKEY value, whereas the file system on the target might assign a ‘7’ for the target SYSKEY value. Previously, Shadowbase had to maintain a “mapping file” on the target to translate any incoming DML for row ‘5’ into DML for row ‘7’ during replication.)

With SYSKEY replication enabled, Shadowbase replication now applies the source SYSKEY value to the target table’s SYSKEY value. This support includes both tables defined with a SYSKEY as the primary key and tables that use the SYSKEY as a clustering key.

**Important Note:** The Shadowbase SQL/MX replication system uses ODBC Data Source Names (DSN) to connect to the database and apply changes. Replication modifies the DSN connection to enable SYSKEY replication, this modification persists even after the DSN connection has been returned to the pool of connections. Other processes that use the same connection may encounter unexpected SQL failures.



**We strongly recommend that you create a dedicated pool of DSN connections for replication of SQL tables with SYSKEYs when using Shadowbase.**

This release has a new shadparm.ini parameter, SHAD\_OVERRIDE\_SYSKEY, to enable or disable exact SYSKEY replication. See the SHAD\_OVERRIDE\_SYSKEY entry, below, for details.

**Specifying the Starting Position for HPE Shadowbase Log Source for Oracle Collection**

You can reset the starting position for HPE Shadowbase Log Source for Oracle collection by editing your collector using SBMON and specifying the starting SCN. SBMON supports starting at the first available event by specifying a starting SCN of 0, starting at end of file by specifying a starting SCN of -1, or starting at a specified SCN by specifying a starting SCN > 0. A sample session to start at the end of file is below:

```
+edit! SMILY
Type of server: 1. DOC, 2. Direct, 3. Transaction Replay, 4. Open Collector,
                5. Listener, 6. Transaction Forward, 7. Remote NSK,
                8. Custom (User) 9. DOC Cleaner 10. Log Server : [4]
Executable Name [sborlog]
sbmscol = MS SQL Server source,
sborcol = Oracle source,
sborlog = Oracle LOG source,
sbsycol = Sybase source.
(sbmscol,sborcol,sborlog,sbsycol) :
Event record format [3]
1 = SQL Statement
2 = Comma Delimited
3 = Cached SQL Statement
(1, 2, 3) :
Transaction Processing ([Y]/N)?
Write transaction boundaries to event file ([Y]/N)?
Current Sequence Number [410] :
Source database format [ORACLE]
(MSSQL, ORACLE, SYBASE, SQL92) :
Database server name
[(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (Host=10.1.60.149) (Port=1521)) (CONNECT_DATA=(
SID=ORAl1)))] :
Source database [qasource] :
User name [qasource] :
Password must be re-entered each time
Password : qas
Starting SCN (Use -1 for EOF, 0 for BOF) [5817736] : -1
Starting RBASQN [253] : -
Invalid Entry: RBASQN Number must be all digits!

Starting RBASQN [253] : 0
Starting RBABLK [61494] : 0
Starting RBABYTE [444] : 0
Starting SSN [444] : 0
+
```

## ***Problems Corrected in Version 6.220***

This section provides a summary of the problems corrected in HPE Shadowbase for Other Servers products since the previous general availability release (V6.101 for all servers except HPE NonStop OSS; V6.200 for HPE NonStop OSS).

1. HPE Shadowbase Log Source for Oracle contains two checks for missing data. The behavior of both checks was supposed to be controlled by the parameter SHAD\_OPCOLLECT\_RESET\_SCN. The default value for this parameter is 0 and indicates that HPE Shadowbase Log Source for Oracle should stop when missing data. However, the second check only logged the following messages and continued, instead of stopping:

```
Potential Data Loss Detected
Last Event Processed in Previous Query Not Read
Current Event SCN [8870042950]
Previous Query SCN [8870034453]
Previous Query RBASQN [49384]
Previous Query RBABLK [26509]
Previous Query RBABYTE [28]
SHAD_OPCOLLECT_RESET_SCN is DISABLED, stopping
Enable SHAD_OPCOLLECT_RESET_SCN to continue processing
prior to restart
```

HPE Shadowbase Log Source for Oracle will also issue the message Performing Shutdown and will stop.

2. A code inspection uncovered a potential bug in HPE Shadowbase Log Source for Oracle that could cause erroneous data loss messages similar to those in item 1, above. When querying the Oracle REDO logs using V\$LOGMNR, Shadowbase uses overlapping queries. The SELECT statement for each query is designed to include the last row read from the previous query. Shadowbase does not restart collecting events until it sees the last row from the previous query. If it does not find the row, it issues the messages listed above.

If Oracle was not able to generate SQL for the last event (which can happen if the table's DDL changes) and the parameter SHAD\_OPCOLLECT\_INVALID\_SQL\_OPTION is set to SKIP or SKIP\_NOLOG, Shadowbase would skip the event prior to the check to see if it was the last row. As a result, Shadowbase would not find the row and would issue the message incorrectly.

3. HPE Shadowbase Log Source for Oracle sometimes incorrectly truncated or translated fields with embedded semicolons, NULLs (binary zeros), and NEWLINES ('\n') at the first instance of the problem character. As a result, Shadowbase did not replicate those fields properly. For example, if the field in the source table contained *“Old version truncated the data incorrectly before; new version replicates it properly”*, previous versions of Shadowbase replicated the data to the target as *“Old*

*version truncated the data incorrectly before*". Shadowbase will now replicated fields with embedded semicolons, NULLs, and NEWLINEs correctly.

4. Several issues with text columns containing single quote characters (‘) or semicolons (;) have been fixed:
  - a. HPE Shadowbase Log Source for Oracle incorrectly dropped single quote characters (‘) if they were at the beginning of the field. For example, previous versions of Shadowbase replicated the field “‘Quote Unquote’” as “Quote Unquote’”. HPE Shadowbase Log Source for Oracle will now correctly capture columns with leading and trailing quotes.
  - b. HPE Shadowbase Log Source for Oracle incorrectly truncated text fields containing semicolons at the semicolon. This release fixes this issue.
  - c. The Transaction Forwarding Server (TFS) incorrectly stripped leading and trailing single quotes (‘) when forwarding events from the DOC. Continuing with the example in item 2, the field “‘Quote Unquote’” would have the quotes removed and would be sent as “Quote Unquote”.
  - d. The Transaction Forwarding Server (TFS), Transaction Replay Server (TRS), and Direct Writer (DW) all incorrectly stripped leading and trailing single quotes when DBS Mapping or User Exits were enabled. Without DBS Mapping or User Exits enabled, the leading and trailing quotes were properly replicated.
6. HPE Shadowbase Log Source for Oracle did not always properly match events when resynching after issuing a SELECT statement for the LOGMNR data. Each SELECT statement overlaps with the previous SELECT. HPE Shadowbase Log Source for Oracle filters events duplicated from the previous SELECT by waiting for the last event in the previous SELECT to be seen before starting event collection from the current SELECT. The matching was sometimes flawed and caused events to be incorrectly collected multiple times.

This issue caused issues replicating on the target system – generating both duplicate key errors and constraint violations. For example, consider a transaction that inserted a parent row and three child rows that was retrieved across two SELECT queries:

**First SELECT:**

SCN	Operation	Comments
< ... >		
10	INSERT PARENT	
10	INSERT CHILD 1	
10	INSERT CHILD 2	Last record of first select.

**Second SELECT (starts with SCN 10):**

SCN	Operation	Comments
-----	-----------	----------

```

10    INSERT PARENT    Overlapping record, already collected
10    INSERT CHILD 1   Overlapping record, already collected
10    INSERT CHILD 2   Overlapping record, already collected
11    INSERT CHILD 3   New record
12    COMMIT
<...>

```

Due to the bug, the collector in rare cases matched the CHILD 1 and CHILD 2 records. In this case, CHILD 2 would get collected twice incorrectly:

#### As Incorrectly Collected

SCN	Operation	Comments
10	INSERT PARENT	Collected from first select
10	INSERT CHILD 1	Collected from first select, incorrectly matched with CHILD 2 on second select
10	INSERT CHILD 2	Collected from first select
10	INSERT CHILD 2	Incorrectly collected from second select
11	INSERT CHILD 3	New record
12	COMMIT	

<...>

This sequence, when replayed in SQL/MX, resulted in the following:

SCN	Operation	Comments
10	INSERT PARENT	Successful (but eventually rolled back)
10	INSERT CHILD 1	Successful (but eventually rolled back)
10	INSERT CHILD 2	Successful (but eventually rolled back)
10	INSERT CHILD 2	Failed, transaction aborted by SQL/MX (see note below)
11	INSERT CHILD 3	New transaction, failed, constraint violation (no PARENT record)
12	COMMIT	Successful, but nothing to commit.

**Note:** The rollback of the transaction by SQL/MX for a duplicate row does not occur on other databases, nor does SQL/MX rollback the transaction for a duplicate row if there are no constraints on the table. Shadowbase, by default, ignores duplicate row errors as they are benign (assuming the database is not aborting the transaction) and as they can occur while replaying data during a restart under certain restart scenarios. Without the transaction abort, all data would be properly replicated to the target even with the duplicate insert. Unfortunately, this is a documented feature of SQL/MX and Gravic has raised it as a ‘feature’ to be reviewed as it does not match the processing performed by other popular commercial databases.

7. HPE Shadowbase Log Source for Oracle did not properly handle out of order SCN numbers in the Redo logs. If the SCN number of the last event returned by the

SELECT statement was less than other SCN numbers in the select (which can happen, for example, at times in Oracle RAC environments), Shadowbase could incorrectly log a missing data error message and stop. This issue has been fixed.

8. HPE Shadowbase Log Source for Oracle collection did not properly handle reversing events in the Oracle Redo logs. A reversing event occurs when a constraint error causes an SQL operation to fail, and the transaction associated with the event is *not* rolled back (the normal case). In this case, Oracle Redo logs will contain the original event that failed, followed by an SQL event that reverses the effect of the failed event. For example, inserting a duplicate record into a table will result in the insert operation appearing in the log, followed by a delete operation for the row.

Previous versions of HPE Shadowbase Log Source for Oracle ignored the reversing event, which resulted in Shadowbase applying the original event to the target database without a corresponding event to undo the original event. Depending on target database and constraints, the event may generate the same error on the target (resulting in the two databases continuing to match), it may be applied successfully (resulting in extra data on the target), or it may cause the transaction to be rolled back and Shadowbase to fail (resulting in missing data on the target).

Shadowbase will now look for a reversing event before writing an insert, update, or delete (IUD) event to the DOC. If the following event is another IUD event or a commit, the event is written to the DOC. If the next event is a reversing event for the same row, the IUD event is discarded (not written to the DOC) and hence is not replicated.

HPE Shadowbase Log Source for Oracle has a new parameter, SHAD\_OPACCOLLECT\_UNMATCHED\_ROLLBACK\_EVENT, to handle the situation if Shadowbase reads a reversing event that does not match the previous event in a transaction that is not rolled back.

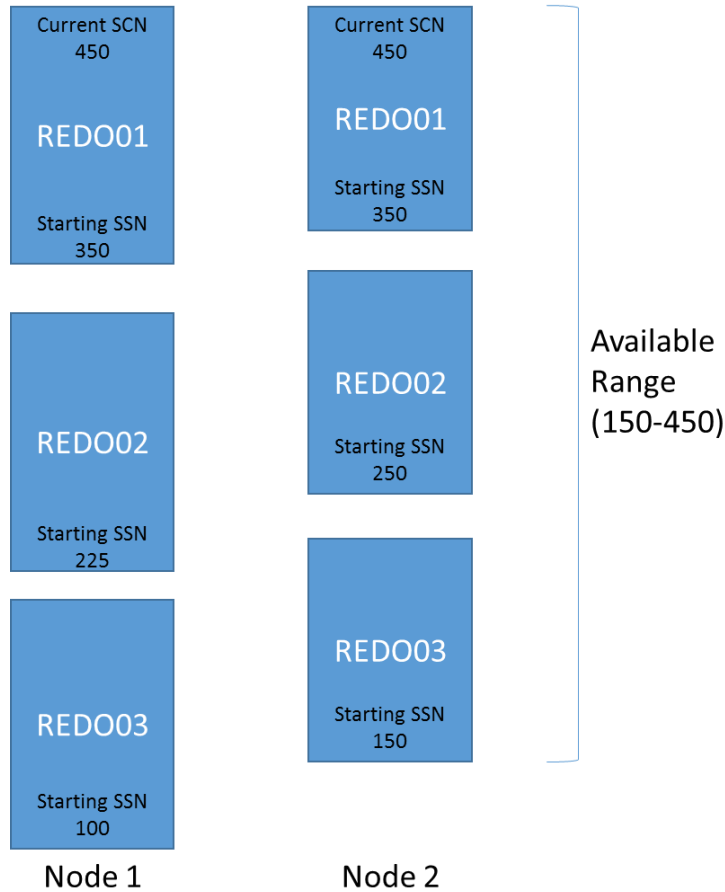
9. HPE Shadowbase Log Source for Oracle was not fully determining the range of data present in RAC environments when determining if unread data has rolled from the online REDO logs and is no longer available for replication. The collector checks that the last SCN read in the previous iteration is still available when reading from the online logs. The lowest available SCNs was determined by using the minimum starting SCN across the REDO logs. This calculation works for single instance Oracle databases, but does not necessarily work in RAC environments where each instance maintains and rolls its own set of REDO logs, particularly if the instances are not well balanced.

The range is now computed by getting the minimum starting SCN of the REDO logs for each instance, and then taking the maximum of those values. The figure below illustrates the change for a 2 node RAC system. Previously, HPE SHADOWBASE LOG SOURCE FOR ORACLE calculated the range of available SCNs for the

example as 100 to 450 when the actual range is 150 to 450.

- 10. A problem that was introduced in Version 6.110H has been fixed. If SHAD\_OPCOLLECT\_RECYCLE\_CONNECTION was enabled, the HPE Shadowbase Log Source for Oracle collector failed while trying to shut down the connection. As part of the connection shutdown process, the collector issued a duplicate call to the Oracle procedure DBMS\_ORACLE.END\_LOGMNR() function, which failed since there was no session active. The collector shutdown on the failure.

This problem has been corrected. END\_LOGMNR() will no longer be called as part of the connection recycling process unless there is an outstanding mining session active. If it is called and fails, an error message will be logged but the collector will continue. It is safe to continue since the collector will close the connection anyway, which will cause the LOGMNR resources to be released anyway.



- 11. The error messages logged for an invalid combination of the SHAD\_OPCOLLECT\_ARCHIVE\_DIR and SHAD\_OPCOLLECT\_MINING\_MODE parameters have been corrected to use the correct parameter name.

12. HPE Shadowbase Log Source for Oracle was incorrectly parsing fields from an insert statement that had one of the following keywords embedded in its text: 'NULL', 'TO\_DATE', 'TO\_TIMESTAMP' or 'TO\_TIMESTAMP\_TZ'. The error could occur as long as the field did not begin with the keyword.

This parsing error only affected INSERT statements captured by HPE Shadowbase Log Source for Oracle. Other sources, including the HPE Nonstop collector, do not use the parsing logic. The HPE Shadowbase Log Source for Oracle collector parsed the UPDATE and DELETE statements correctly.

The behavior of the previous versions varied depending on the embedded keyword. Shadowbase would either stop or generate incorrect data for replication.

For example, if NULL was embedded in the text, Shadowbase would assign NULL to the next column, and shift the remaining data over one column. If the original statement was:

```
INSERT INTO TEST_TABLE (C1, C2, C3) VALUES ('C1 is not NULL', 'C2', 'C3')
```

Shadowbase would generate the following incorrect statement:

```
INSERT INTO TEST_TABLE (C1,C2, C3) VALUES ('C1 is not NULL, NULL, 'C2')
```

Similar behavior was observed for 'TO\_DATE' and 'TO\_TIMESTAMP' text strings. Embedded 'TO\_TIMESTAMP\_TZ' strings caused the Shadowbase collector to stop with an unsupported function message.

13. This release includes changes to the SQL92 version of the SQL dialect to remove the TO\_DATE special processing. The TO\_DATE special processing created an issue that resulted in the improper generation of an SQL statement when one of the column values began with the string TO\_DATE and was more than 7 characters long.

**Note:** The TO\_DATE special processing is a legacy feature that is used by several of our customers in their Shadowbase user exits to convert non-date/time fields to date/time fields. In previous versions, both the SQL92 and ORACLE version of our SQL dialect supported this feature. It has been removed from the SQL92 dialect only and is still supported in the ORACLE version. If you are using the ORACLE dialect, you should switch to SQL92 to disable the processing.

14. HPE Shadowbase Log Source for Oracle collection was referencing the system table X\$KTUXE view. This view is not needed and has been removed.

15. HPE Shadowbase Log Source for Oracle collection did not properly handle collecting data from tables that had the same table name but where in different schemas. This has been corrected.
16. HPE Shadowbase Log Source for Oracle collection generated incorrect update and delete statements for tables where the order of the columns when created is different from the order of the columns in the primary key. The generated statements could result in corrupted and incorrect data in the target tables. For example, HPE SHADOWBASE LOG SOURCE FOR ORACLE generated incorrect update and delete statements for the following table:

```
SQL> desc AMOUNT_TYPE_XREF;
Name                                         Null?    Type
-----
EXTERNAL_AMOUNT_TYPE_ID                    NOT NULL VARCHAR2 (50)
SOURCE_SYSTEM                               NOT NULL VARCHAR2 (30)
CREATED_ON                                  DATE
CREATED_BY                                 NOT NULL VARCHAR2 (30)
MODIFIED_ON                                 DATE
MODIFIED_BY                                VARCHAR2 (30)
AMOUNT_TYPE_ID                              NUMBER (38)
```

With the primary key defined as (SOURCE\_SYSTEM, EXTERNAL\_AMOUNT\_TYPE\_ID).

17. The HPE Shadowbase Log Source for Oracle collector did not properly pick up data from all partitions of a partitioned table. For partitioned tables, the collector generated error messages similar to:

```
2015-07-29 17:48:13 -[19336] Error: OCIDescribeSetCols(Descriptor): fault
detected; rc:=100 (Not Found); Table[QASOURCE.MARGIN_BALANCE,P2]
2015-07-29 17:48:13 -[19336] Error: shadparm.ini parameters
SHAD_OCI_INCOMPLETE_SCHEMA and SHAD_OCI_INCOMPLETE_SCHEMA_LIMIT control the
response to this condition
2015-07-29 17:48:13 -[19336] Error: SHAD_OCI_INCOMPLETE_SCHEMA=[SHUTDOWN]
(default); SHAD_OCI_INCOMPLETE_SCHEMA_LIMIT=[-1]
```

Changing the SHAD\_OCI\_INCOMPLETE\_SCHEMA settings allows the data to be skipped, but does not allow it to be collected. This problem has been corrected.

18. When an event is read for a table that has been dropped, Oracle will generate an invalid SQL statement, which caused the Collector to stop. A new parameter has been added to optionally allow those events to be skipped. See the description for , below.
19. When an event is read by HPE Shadowbase Log Source for Oracle with a field value that is not supported, previously the Collector stopped. The restart SCN had to be manually adjusted to skip the event to allow the Collector to continue. A new parameter has been added to optionally skip those events. See the description for , below.
20. The TRS could lose data when it was caught up while replaying from a DOC generated by the HPE Shadowbase Log Source for Oracle collector. For proper



sequencing of transactions, the TRS expects the commit time to be unique and ascending. Under certain conditions, the HPE Shadowbase Log Source for Oracle collector was not correctly generating a unique, ascending timestamp.

21. The Direct Writer could abend trying to print out an SQL statement when the parameter SHAD\_DISPLAY\_STATEMENTS was enabled (set to 1) in the shadparm.ini. This problem has been fixed.
22. Previously, editing an ODBC object such as a Direct Writer did not allow the Data Sources (ODBC) User DNS parameter to be changed. Editing the object allowed you to specify a new DNS name, however, the new name was not saved when the edit was complete. As a result, the DNS name remained unchanged. This issue has been corrected.
23. With diagnostic messages enabled, the HPE Shadowbase Log Source for Oracle collector writes the statements executed against the database to the log. These statements can then be run manually to help diagnose issues. Previously, the collector incorrectly formatted the “EXEC DBMS\_LOGMNR.START” statement. As a result, the statement failed when run manually.
24. Version 6.220 contains a work-around to an intermittent Oracle LogMiner issue related to insert, updates, and deletes of tables with long rows. This issue causes Oracle LogMiner to return incorrect results for an event if its SCN is at the beginning of the range of SCNs specified when the Oracle LogMiner session is initialized by calling the Oracle procedure DBMS\_LOGMNR.START\_MINING(). This behavior was uncovered in a highly stressed Oracle environment with large rows and very small REDO logs rolling extremely frequently for events whose SCN was the same as that specified in the START\_MINING() call. This issue seems to have two manifestations:
  - a. Oracle LogMiner returned syntactically correct REDO information for the event, but with the actual values for some of the table’s columns replaced by NULL values, or,
  - b. Oracle LogMiner returned one or more partial events with “UNSUPPORTED” for the REDO information.

Depending on the columns that were incorrectly set to NULL, the first manifestation could cause Shadowbase to:

- c. Incorrectly detect that a key column update occurred, if Oracle incorrectly sets the key column to NULL. Depending on how key column updates are handled, this could cause the collector to skip the event, to stop, or to convert it to a delete and insert. The insert would fail when applied to the target database as one of the key columns would be NULL.

- d. Fail when the event is replicated to the target database, if one or more of the incorrect columns are not nullable.
- e. Apply incorrect data, if all of the incorrect columns are nullable on the target.

When Oracle LogMiner returned partial events with “UNSUPPORTED”, Shadowbase would detect and log potential data loss messages, and, by default, stop.

Gravic has added functionality to avoid this issue by starting the collection of the data earlier than is needed by offsetting the starting SCN for collection by a configurable amount. See the SHAD\_OPCOLLECT\_ADJUST\_START\_SCN for more information.

- 24. Gravic has corrected an issue introduced in Version 6.110E that could result in data loss when the HPE Shadowbase Log Source for Oracle collector is stopped with multiple active transactions. In this case, the collector is supposed to restart at the oldest event in the first transaction started. However, the collector was not properly determining the first transaction started, and, as a result, did not always restart far enough back in the REDO log to recollect all of the data for the first transaction.
- 25. The CONTINUOUS\_ARCHIVE mode introduced in version 6.110H used the continuous mining mode of Oracle LogMiner. It correctly restricted the data being retrieved to data that had been archived, however, when the data was available in both archived and online REDO logs, Oracle LogMiner accessed the online logs. The CONTINUOUS\_ARCHIVE mode now explicitly loads the archive files instead of using the continuous mining mode, restricting access to only the archived REDO logs.
- 26. The Visual Studio solution for building user exits on Windows included with the installation did not have the USRX\_DLL defined in the Preprocessor definitions for the release build configurations, causing the builds to fail.
- 27. In certain very rare cases, the TFS would incorrectly attempt to process a string as hex-encoded. This could occur if hex encoding of strings with binary data was enabled (SHAD-NLS\_CONVERT\_LEVEL=1 in prior releases) and a row was replicated that a column with an ‘x’ character in the second byte immediately followed by a column with length 1 that only had ‘0’ character in the first byte. The ‘0’ from the second column and the ‘x’ from the first column were improperly overlaid to form a ‘0x’, which indicates that the second column is hex-encoded when SHAD\_NSL\_CONVERT\_LEVEL=1.
- 28. If HPE Shadowbase Log Source for Oracle is configured to use CONTINUOUS\_ONLINE monitoring and additional REDO logs are configured in Oracle to expand the redo capacity, HPE Shadowbase Log Source for Oracle could fail with an Oracle error “Err[ORA-01289]: cannot add duplicate logfile <log file

name>. HPE Shadowbase Log Source for Oracle was incorrectly adding empty REDO logs to the query.

29. A number of performance issues have been addressed in this release:

- a) The SHAD\_CHECKPOINT\_EVENTS parameter should control how frequently the statistics and restart point are updated in the COLLCONFIG file. The parameter value, however, was ignored in prior releases and the default value of 1 was used, resulting in frequent updates to the COLLCONFIG file. This release fixes this issue so that increasing the value of SHAD\_CHECKPOINT\_EVENTS will significantly reduce the number of updates to the COLLCONFIG file.
- b) The SBMON or SEM STATUS command will also cause the Shadowbase objects to write statistics to the COLLCONFIG file. A signal handler receives the request and sets a flag to have the COLLCONFIG updated at the next available time, typically when processing the next event. Previous versions did not clear this flag after updating the COLLCONFIG with the statistics. As a result, after the user requested one status command, the Shadowbase object updated the COLLCONFIG with every event, resulting in a significant performance impact. The current release corrects this issue.
- c) The shadparm.ini SHAD\_STATUS\_STATS = 0 parameter setting did not work on Linux, Solaris, HP-UX, AIX, and OSS systems. This setting uses a different methodology to see if the object is running and inhibits the writing of statistics to the COLLCONFIG file while processing a status command. The current release corrects this issue. The documentation for this command in the *HPE Shadowbase Other Servers Parameter Reference Guide* incorrectly showed the default as 0, it has been and remains 1. Corrected documentation for this parameter is provided in SHAD\_STATUS\_STATS below.
- d) Issuing a STATUS or other command could cause an internal deadlock within the Shadowbase object, causing it to hang. The deadlock occurred within the signal handler processing the command when the process was interrupted while issuing a malloc or other memory heap command. The current release removes the potential for the deadlock.

## ***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 (V6.101 for all servers except HPE NonStop OSS; V6.200 for HPE NonStop OSS).

### **SHAD\_HEX\_CONVERT\_LEVEL**

<b>Initial Version:</b>	V6.110G
<b>Syntax:</b>	SHAD_HEX_CONVERT_LEVEL=value
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0 or 1

#### **Description**

Specifies how the Shadowbase handles text strings that begin with '0x'. If set to 0 (the default), text strings that begin with 0x are applied to the database as is. If set to 1, the string is converted from hex. For example, if the source string is '0x4142' and SHAD\_HEX\_CONVERT\_LEVEL is set to 0, the string '0x4142' will be stored on the target. If SHAD\_HEX\_CONVERT\_LEVEL is set to 1, the string '0x4142' will be stored as 'AB' ('A' = hex 41; 'B' = hex 42).

### **SHAD\_OPCOLLECT\_ADJUST\_END\_SCN**

<b>Initial Version:</b>	V6.110G
<b>Syntax:</b>	SHAD_OPCOLLECT_ADJUST_END_SCN=value
<b>Last Change:</b>	--
<b>Default Value</b>	2
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0 - 2147483647

#### **Description**

Specifies how close to the end of the REDO logs the Shadowbase Oracle Log collector reads. Each time the collector starts reading the log, it determines the minimum current SCN across all database instances and uses that as the highest available SCN that can be selected in the current iteration. This parameter's value is subtracted from the highest available SCN prior to the selection to further restrict the search. If SHAD\_OPCOLLECT\_ADJUST\_END\_SCN is set to 0, the highest available SCN will be used in the SELECT statement.

For example, consider a two node RAC database, with the current SCN for Node 1 set to 249 and the current SCN for Node 2 set to 250. The highest available SCN is 249 (Node 1's current SCN). If SHAD\_OPCOLLECT\_ADJUST\_END\_SCN is set to 2 (the default), the collector will issue a SELECT statement to collect up to and including SCN 247.

**Note:** SHAD\_OPCOLLECT\_ADJUST\_SCN will be used when performing collection using the CONTINUOUS\_ONLINE, CONTINUOUS\_ARCHIVE, and CONTINUOUS\_BOTH modes.

## SHAD\_OPCOLLECT\_ADJUST\_START\_SCN

<b>Initial Version:</b>	V6.110L
<b>Syntax:</b>	SHAD_OPCOLLECT_ADJUST_START_SCN=value
<b>Last Change:</b>	--
<b>Default Value</b>	2
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	1 to 281474976710655

### Description

This parameter provides a work around to an issue in Oracle Log Mining by adjusting the start position of each collection cycle. Testing has indicated that Oracle will not reliably return records at the beginning of the collection range. This issue is infrequent, but it can cause Oracle to return syntactically correct data, but with some of the actual column values incorrectly set to NULL. It can also cause records in the beginning of the collection range to be missed, or to have partial records return with no usable information.

The value specified by this parameter is used to start the collection cycle at an earlier point by subtracting the value from the actual start SCN to determine the start of collection. Events up to and including the one last collected in the previous cycle are skipped, thus insuring that the data that is needed for the collection cycle is correct.

The default value for this parameter is 2. However, based upon additional testing, Gravic recommends a higher value, such as 10, to insure reliable collection.

## SHAD\_OPCOLLECT\_ARCHIVE\_DIR

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_ARCHIVE_DIR=<directory path>
<b>Last Change:</b>	--
<b>Default Value</b>	none
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	Must be a fully qualified path name accessible by the Oracle Database or omitted.

### Description

This parameter enables processing of archived REDO logs and specifies the directory where the REDO logs are located. Note that the directory must be accessible by the Oracle database server as the logs are read through the Oracle V\$LOGMNR\_CONTENTS view.

If the directory is accessible to Shadowbase and if the

SHAD\_OPCOLLECT\_ARCHIVE\_FILES parameter is not set, Shadowbase will add all of the files located in the directory to the V\$LOGMNR\_CONTENTS view. If SHAD\_OPCOLLECT\_ARCHIVE\_FILES is set, Shadowbase will combine the directory with the files specified by that parameter to form the list of redo logs to be processed.

**Notes:**

- 1) The log for the collector will list the both the directory and the REDO logs that will when this parameter is specified.
- 2) If the SHAD\_OPCOLLECT\_ARCHIVE\_FILES is not set and either there are no files in the directory specified or Shadowbase cannot access the files in the directory, the collector will stop.

**Examples:**

Assume the directory /Oracle/ArchiveLogs contains the following 10 archive files:

Archive1.arc	Archive4.arc	Archive7.arc	Archive10.arc
Archive2.arc	Archive5.arc	Archive8.arc	
Archive3.arc	Archive6.arc	Archive9.arc	

Example 1

If the shadparam.ini only contains the parameter:

```
SHAD_OPCOLLECT_ARCHIVE_DIR=/Oracle/ArchiveLogs
```

all 10 of the archive files will be added to the V\$LOGMNR\_CONTENTS for processing.

Example 2

If the shadparam.ini contained the parameters:

```
SHAD_OPCOLLECT_ARCHIVE_DIR=/Oracle/ArchiveLogs
SHAD_OPCOLLECT_ARCHIVE_FILES=/SB/ArchiveLogs.txt
```

and the /SB/ArchiveLogs.txt file contained:

```
Archive1.arc
Archive2.arc
Archive3.arc
```

only the events from the first three REDO log files would be processed.

***Example 3***

You can also use fully qualified filenames in file specified for by the SHAD\_OPCOLLECT\_ARCHIVE\_FILES. If the /SB/ArchiveLogs.txt from Example 2 contained:

```
/Oracle/ArchiveLogs/Archive1.arc
/Oracle/ArchiveLogs/Archive2.arc
/Oracle/ArchiveLogs/Archive3.arc
```

the results would be the same, only the first three REDO log files would be processed. Note that the directory specified by SHAD\_OPCOLLECT\_ARCHIVE\_DIR must match the beginning of the file names in the SHAD\_OPCOLLECT\_ARCHIVE\_FILES file exactly.

**SHAD\_OPCOLLECT\_ARCHIVE\_FILES**

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_ARCHIVE_FILES=<local file name>
<b>Last Change:</b>	--
<b>Default Value</b>	none
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	Must be a filename that can be read by the collector or omitted.

**Description**

This parameter specifies a text file that contains the list of archived Oracle REDO logs to be used when reading and processing events from Oracle. The text file must be readable by the Shadowbase collector, and must have each REDO log name on a separate line. The REDO log names can be either fully qualified filenames, including the directory name, or they can be just filenames relative to the SHAD\_OPCOLLECT\_ARCHIVE\_FILES directory. Use this parameter if:

- The Shadowbase collector is running on a different server from the Oracle database server.
- Not all of the archived REDO logs in the directory specified by the SHAD\_OPCOLLECT\_ARCHIVE\_DIR parameter are to be used in by the Shadowbase collector.

See the examples under the description for SHAD\_OPCOLLECT\_ARCHIVE\_DIR for usage examples.

**Notes:**

- 1) If the specified file is not readable by the collector or contains no lines, the collector will log an error message and stop.

- 2) The collector uses the following algorithm to build the filename for each REDO archive log for Oracle: First, the filename read from the file is checked to see if it starts with the directory specified by the SHAD\_OPCOLLECT\_ARCHIVE\_DIR. If it does, the filename as read from the file is used. If it does not, the directory is prepended to the filename, and the combined name is used.

If the full path name is specified, the beginning of the path name must match the directory parameter exactly, including the character case. Otherwise, the collector will incorrectly prepend the directory again, resulting in an invalid file.

- 3) The generated file names will be logged in the collector's error log during initialization.
- 4) This collector only uses this parameter if the SHAD\_OPCOLLECT\_ARCHIVE\_DIR parameter is also specified.

## SHAD\_OPCOLLECT\_ARCHIVE\_CATALOG

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_ARCHIVE_CATALOG=<filename>
<b>Last Change:</b>	--
<b>Default Value</b>	none
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	Filename of exported flat-file catalog, or omitted

### Description

This parameter provides the name of the exported flat-file catalog used by V\$LOGMNR\_CONTENTS to read archived REDO logs. If omitted, the Shadowbase collector will use the online directory. Use this parameter if Shadowbase is reading the archived REDO logs through a different Oracle database instance than the one that generated the logs. See the chapter on Using LogMiner to Analyze Redo Log Files in the Oracle Database Utilities documentation for more information.

## SHAD\_OPCOLLECT\_CHECK\_LOGMNR\_SEQUENCE

<b>Initial Version:</b>	V6.110D
<b>Syntax:</b>	SHAD_OPCOLLECT_CHECK_LOGMNR_SEQUENCE =option
<b>Last Change:</b>	--
<b>Default Value</b>	OFF
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	One of <i>STOP</i> , <i>LOG</i> , or <i>OFF</i>

### Description

This parameter specifies if the HPE Shadowbase Log Source for Oracle Collector checks the sequence of events collected from Oracle LOGMNR. The order of events



is no longer critical; this parameter should be left at its default value (OFF). We will likely remove this parameter in a future release.

- *OFF* (default): The collector does not check the event order.
- *LOG*: The collector checks the order of events to see if it is in SCN / RBASQN / RBABLK / RBABYTE / SSN ascending order. If it is not, the collector logs the out-of-order condition and continues.
- *STOP*: The collector checks the order of events to see if it is in SCN / RBASQN / RBABLK / RBABYTE / SSN ascending order. If it is not, the collector logs the out-of-order condition and stops.

## SHAD\_OPCOLLECT\_DDL\_TRACKING

<b>Initial Version:</b>	V6.110D
<b>Syntax:</b>	SHAD_OPCOLLECT_DDL_TRACKING=option
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0: Oracle log mining does not use the dictionary tracking option 1: Oracle log mining uses the dictionary tracking option

### Description

This option, when enabled, will cause Oracle log mining to track DDL changes in its internal dictionary while returning records from the event log. This allows Oracle to return log records with valid REDO\_SQL (necessary for replication) after DDL events such as adding or dropping columns. However, this option has significant limitations that reduce its utility:

- 1) You cannot use this option with the online catalog. You must export and specify a flat dictionary (see the **SHAD\_OPCOLLECT\_ARCHIVE\_CATALOG** option for more information).
- 2) The SCN associated with the export operation must be in logs being processed. Log mining starts reading at the point where the catalog was exported to insure that all of the DDL events are captured in its internal database.
- 3) In the continuous modes of mining, Shadowbase will read to the end of the data, wait a period, and then restart reading where it left off (with a little overlap). If **SHAD\_OPCOLLECT\_DDL\_TRACKING** is enabled,

mining will actually start at the point where the dictionary was exported, resulting increasing processing and delays as the logs grow.

Gravic recommends only enabling this option for the ARCHIVE\_ONLY mode of operation (see **SHAD\_OPCOLLECT\_MINING\_MODE** for more information).

Enabling this parameter causes the Oracle procedure DBMS\_LOGMNR.START\_LOGMNR to be called with the DBMS\_LOGMNR.DDL\_DICT\_TRACKING option enabled. See the Oracle documentation for more information on the use of this option.

## SHAD\_OPCOLLECT\_END\_SCN

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_END_SCN=value
<b>Last Change:</b>	--
<b>Default Value</b>	none
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	1 to 281474976710655

### Description

This parameter limits the selection of Oracle System Change Numbers (SCNs) to those less or equal to the valued specified. By changing the starting SCN through SBMON and specifying SHAD\_OPCOLLECT\_END\_SCN, processing can be limited to a specific range of SCNs.

### Notes:

- 1) If the parameter is invalid (for example, if it contains non-numeric characters or if it is too large), the collector will stop.

## SHAD\_OPCOLLECT\_INVALID\_SQL\_OPTION

<b>Initial Version:</b>	V6.110
<b>Syntax:</b>	SHAD_OPCOLLECT_INVALID_SQL_OPTION=option
<b>Last Change:</b>	--
<b>Default Value</b>	STOP
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	One of <i>STOP</i> , <i>SKIP</i> , or <i>SKIP-NOLOG</i>

### Description

Oracle uses schema information (such as column names and types) when generating events for the V\$LOGMNR\_CONTENTS view. If the table has been altered prior to the events being read, the current schema information used by Oracle does not match

the event. In this case, Oracle will generate an invalid SQL statement that just uses column numbers for the column name and hex strings for the column values. The HPE Shadowbase Log Source for Oracle collector cannot parse or use these events. This parameter specifies the behavior of the HPE Shadowbase Log Source for Oracle Collector (sborlog) when it reads such an event. The options are:

- *STOP* (default, previous behavior): collection will stop with an error message when it reads an event with an invalid SQL statement.
- *SKIP*: collection will log but not collect the event when it reads an event with an invalid SQL statement. Collection continues with the next event.
- *SKIP-NOLOG*: collection will skip without logging when it reads an event with an invalid SQL statement. Collection continues with the next event.

### Notes

Gravic recommends that you stop application updates, drain the collector, and stop the collector prior to altering replicated tables. On restart of the application and Shadowbase, the new events will have the new scheme format and Shadowbase will load the correct schema for the events. Otherwise, there may be some events that cannot be parsed and replicated, resulting in data loss.

## SHAD\_OPCOLLECT\_MINING\_MODE

<b>Initial Version:</b>	V6.110H
<b>Syntax:</b>	SHAD_OPCOLLECT_MINING_MODE=value
<b>Last Change:</b>	--
<b>Default Value</b>	CONTINUOUS_ONLINE
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	CONTINUOUS_ONLINE, CONTINUOUS_ARCHIVE, CONTINUOUS_BOTH, OFFLINE_ARCHIVE.

### Description

Specifies how the Shadowbase Log Source Collector for Oracle (SBORLOG) collects events from Oracle LogMiner. The collector supports four modes of operation, specified by the four valid settings:

- *CONTINUOUS\_ONLINE*: Shadowbase continuously collects events from the online Oracle REDO logs only. REDO log archiving does not have to be enabled for the database. Shadowbase explicitly adds the REDO logs to be read using the `DBMS_LOGMNR.ADD_LOGFILE()` procedure call and the collection is limited to the available range of SCNs in the online REDO logs.
- *CONTINUOUS\_BOTH*: Shadowbase continuously collects events from both the online and archived REDO logs. REDO log archiving must be enabled for the database. In this mode, Shadowbase uses the continuous mining feature of Oracle LogMiner. For more information on continuous mining, please see the Oracle

documentation under [Using LogMiner to Analyze Redo Log Files](#).

- **CONTINUOUS\_ARCHIVE**: Shadowbase continuously collects events from the archived REDO logs only. REDO log archiving must be enabled for the database. In this mode, Shadowbase uses the continuous mining feature of Oracle LogMiner. For more information on continuous mining, please see the Oracle documentation under [Using LogMiner to Analyze Redo Log Files](#).

This mode is essentially the same as the **CONTINUOUS\_BOTH** mode, except Shadowbase uses the SCN number to limit the data collection to data that is available in the Archive files only.

- **OFFLINE\_ARCHIVE**: Shadowbase collects events from a specific set of archived REDO logs and stops once it has processed those events. You must specify the directory containing the archive files using the **SHAD\_OPACCOLLECT\_ARCHIVE\_DIR** parameter. See the description under **SHAD\_OPACCOLLECT\_ARCHIVE\_DIR** for more information.

Note: If you need to archive an online REDO logs, you can use the Oracle command `'ALTER SYSTEM ARCHIVE LOG CURRENT'` or `'ALTER SYSTEM SWITCH LOGFILE'` to cause the logs to be archived. `ALTER SYSTEM ARCHIVE LOG CURRENT` will archive the REDO logs from all nodes of a RAC system and is synchronous, whereas `ALTER SYSTEM SWITCH LOGFILE` will only archive the REDO logs from the attached node and is asynchronous.

## SHAD\_OPACCOLLECT\_PARSE\_ERROR\_OPTION

<b>Initial Version:</b>	V6.110
<b>Syntax:</b>	SHAD_OPACCOLLECT_PARSE_ERROR_OPTION=option
<b>Last Change:</b>	--
<b>Default Value</b>	STOP
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	One of <i>STOP</i> , <i>SKIP</i> , or <i>SKIP-NOLOG</i>

### Description

This parameter specifies the behavior of the HPE Shadowbase Log Source for Oracle Collector (sborlog) when it reads an event that it cannot parse, most likely because of a change to the table's schema. The options are:

- **STOP** (default, previous behavior): collection will stop with an error message when a parsing error occurs on a replicated table.
- **SKIP**: collection will log and continue, but will not collect the event when a parsing error occurs.
- **SKIP-NOLOG**: collection will skip without logging parsing errors.

## Notes

The collector reads and saves the schema for a table when it encounters the first event to replicate associated with the table. Restarting the collector after it encounters this error may correct the issue as the changed schema will be loaded.

## SHAD\_OPCOLLECT\_PATH\_SEPARATOR

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_PATH_SEPARATOR=value
<b>Last Change:</b>	--
<b>Default Value</b>	See description
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	\, /, or omitted

## Description

Specifies the separator character in directory paths for the Oracle database server when building archived REDO log file names. This parameter is only used for reading archived log files and only needs to be specified when the Shadowbase collector is running on a system that has a different path specifier from the Oracle database server, such as when the Shadowbase collector is running on a Windows server and the Oracle database is running on a Solaris system.

The parameter default is based upon the type of system the collector is running on. For Windows, the default is the backslash (\) character; for Linux, Solaris, HPE-UX, and AIX, the default is the slash (/) character.

## SHAD\_OPCOLLECT\_QUERY\_EXECUTION\_LIMIT

<b>Initial Version:</b>	V6.110C
<b>Syntax:</b>	SHAD_OPCOLLECT_QUERY_EXECUTION_LIMIT=<value>
<b>Last Change:</b>	--
<b>Default Value</b>	250
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	1 to 2147483646

## Description

This parameter specifies how many SELECT iterations the HPE Shadowbase Log Source for Oracle collector will perform before resetting the connection to Oracle. This parameter is only used if the SHAD\_OPCOLLECT\_RECYCLE\_CONNECTION parameter is enabled, otherwise, the database connections to Oracle are not reset.

Gravic has seen instances of resource leaks associated with connections with Oracle that grow over time. Periodically closing and re-establishing the database connection causes Oracle to release the resources, ameliorating the effects of the resource leak.

Note that closing and opening database connections is relatively resource intensive. Increasing the rate (by decreasing this parameter) will reduce the amount of resources leaked by Oracle at the cost of decreasing the performance of Shadowbase.

### Related Parameters

Also see SHAD\_OPCOLLECT\_RECYCLE\_CONNECTION.

## SHAD\_OPCOLLECT\_RECYCLE\_CONNECTION

<b>Initial Version:</b>	V6.110C
<b>Syntax:</b>	SHAD_OPCOLLECT_RECYCLE_CONNECTION=<option>
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0: Don't recycle database connections to Oracle 1: Recycle database connections to Oracle.

### Description

This parameter specifies if connection recycling is enabled for the HPE Shadowbase Log Source for Oracle collector. Gravic has seen instances of resource leaks associated with connections with Oracle that grow over time. Periodically closing and re-establishing the database connection causes Oracle to release the resources, ameliorating the effects of the resource leak.

Note that closing and opening database connections is relatively resource intensive. This parameter determines if the collector recycles connections. If this parameter is enabled (set to 1), SHAD\_OPCOLLECT\_QUERY\_EXECUTION\_LIMIT will determine how frequently the connection is recycled.

### Related Parameters

See SHAD\_OPCOLLECT\_QUERY\_EXECUTION\_LIMIT, which controls the number of SELECT statement iterations prior to the connection being recycled.

## SHAD\_OPCOLLECT\_RESET\_SCN

<b>Initial Version:</b>	V6.100
<b>Syntax:</b>	SHAD_OPCOLLECT_RESET_SCN=<option>
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0: Stop when missing data is detected. 1: Reset the starting or ending SCN when missing data is detected.

### Description

This parameter specifies how the collector behaves when it detects that it may miss data in a collection cycle. If disabled (the default), the collector will stop after logging information about the potentially missing data. If

enabled, the collector will continue, potentially missing data that should be replicated.

The collector uses SELECT statements against the Oracle LogMiner view V\$LOGMNR\_CONTENTS to read the data from the Oracle logs. In the continuous mining modes (CONTINUOUS\_ONLINE, CONTINUOUS\_OFFLINE, and CONTINUOUS\_BOTH), the collector will repeatedly issue the SELECT statements, with the next SELECT statement starting at approximately the SCN with the last event from the previous SELECT statement (with some overlap as specified by the **SHAD\_OPCOLLECT\_ADJUST\_START\_SCN** parameter). Prior to issuing the SELECT, it checks to see that the new starting SCN is still contained within the REDO logs.

## SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES

<b>Initial Version:</b>	V6.110C
<b>Syntax:</b>	SHAD_OPCOLLECT_SCAN_FOR_NULL_BYTES=option
<b>Last Change:</b>	--
<b>Default Value</b>	STOP
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	One of <i>STOP</i> , <i>KEEP</i> , <i>OFF</i> , <i>SKIP</i> , <i>SKIP-NOLOG</i> or <i>&lt;number&gt;</i>

### Description

This parameter specifies the behavior of the HPE Shadowbase Log Source for Oracle Collector (sborlog) when it reads an event that has text fields with embedded NULL (0) bytes. The options are:

- *STOP* (default): collection will stop with an error message when a column with embedded NULL bytes is read.
- *KEEP*: The collector will keep the NULL bytes.
- *OFF*: The collector will not scan for NULL bytes. **Note:** only use this option if you are absolutely sure there are no NULL bytes in the text fields of the tables being replicated. If this option is in effect and there are NULL bytes in the text field, replication may fail or corrupt data.
- *SKIP*: The collector will log and continue, but will not collect the event when a parsing error occurs.
- *SKIP-NOLOG*: collection will skip without logging parsing errors.
- *<number>*: if a number is specified, the NULL byte will be replaced with the specified number. Numbers can either be specified in decimal (any number starting with digits '1' through '9'), hexadecimal (any number starting with '0x' or '0X'), or octal any number starting with just '0'. See the example below.

## Replacement Example:

The following settings can all be used to replace NULL bytes with ascii spaces (‘ ‘, decimal value 32):

```
SHAD_OPCOLLECT_SCAN_FOR_NULL_BYTES=32
SHAD_OPCOLLECT_SCAN_FOR_NULL_BYTES=0x20
SHAD_OPCOLLECT_SCAN_FOR_NULL_BYTES=040
```

## Note:

- 1) In a future release, the default value for this parameter will likely change to ‘KEEP’.

## SHAD\_OPCOLLECT\_SKIP\_CORRUPT\_REDO\_BLOCKS

<b>Initial Version:</b>	V6.110G
<b>Syntax:</b>	SHAD_OPCOLLECT_SKIP_CORRUPT_REDO_BLOCKS=value
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0,1

## Description

Specifies if the Shadowbase collector starts Oracle LogMiner with the SKIP\_CORRUPTION option. If 0 (the default), the option is not enabled. Oracle LogMiner will stop collecting data and return with an error, which will cause the collector to stop and log the error. If set to 1, the SKIP\_CORRUPTION option will be enabled and any corrupted blocks will be *silently* skipped. Gravic recommends only enabling this error after the collector has stopped due to corrupted REDO logs to allow collection past the corruption.

## SHAD\_OPCOLLECT\_UNMATCHED\_ROLLBACK\_EVENT

<b>Initial Version:</b>	V6.110E
<b>Syntax:</b>	SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT=value
<b>Last Change:</b>	--
<b>Default Value</b>	STOP
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	STOP, LOG, OFF, TRACK-ALL

## Description

Specifies how the Shadowbase collector handles an unmatched rollback event. When the collector reads a rollback event from the Oracle REDO logs, it checks to see if that event matches the row id for the previous event **for that transaction**. If it does, the previous event is removed from the replication stream (as it was rolled back) and replication continues.



If the events do not match, the collector flags the transaction as it expects the entire transaction to be aborted. If the transaction is committed instead, the collector performs the action specified by this parameter as the transaction contained unmatched rollback events. The values are:

- STOP (default value) – the collector logs the event and stops.
- LOG – the collector logs the event and continues.
- OFF – the collector continues without logging.
- TRACK-ALL –all events are tracked in memory and checked for rollback events. Use this option only if directed by Shadowbase Support.

## SHAD\_OPCOLLECT\_USE\_ORDER\_BY

<b>Initial Version:</b>	V6.110D
<b>Syntax:</b>	SHAD_OPCOLLECT_USE_ORDER_BY=option
<b>Last Change:</b>	--
<b>Default Value</b>	0 (disabled)
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0 or 1

### Description

This parameter enables ordering of LOGMNR data events by the SCN, RBASQN, RBABLK, RBABYTE, and SSN columns by including an *ORDER BY* clause in the *SELECT* statement. This is no longer required and should not be changed from the default (disabled). A significant decrease in performance will occur with the parameter enabled.

- 0 (default): no ordering of event data is performed by LOGMNR.
- 1: event data is specifically ordered, resulting in longer query times due to data sorting.

**Note:** LOGMNR ordering is no longer required. We will likely remove this option in a future release. Only enable under the direction of Shadowbase Support.

## SHAD\_OPCOLLECT\_USE\_SEG\_NAME

<b>Initial Version:</b>	V6.110
<b>Syntax:</b>	SHAD_OPCOLLECT_USE_SEG_NAME=option
<b>Last Change:</b>	--
<b>Default Value</b>	0
<b>Open Servers</b>	HPE Shadowbase Log Source for Oracle Collector (sborlog)
<b>Valid Settings</b>	0 (disabled) or 1 (enabled)

## Description

This parameter specifies which column of the V\$LOGMNR\_CONTENTS view is used by the HPE Shadowbase Log Source for Oracle Collector (sborlog) to determine the table name. In previous versions of Shadowbase, the collector used the SEG\_NAME column. This column, however, also has the partition information embedded with the name for partitioned tables. The collector now uses the TABLE\_NAME column to determine the table name by default. Enabling this parameter will cause the collector to revert to its old behavior and use the SEG\_NAME column. Contact Support if you believe you need to enable the old/prior behavior.

The options are:

- *0* (default): The TABLE\_NAME column is used instead of the SEG\_NAME (new behavior, works with partitioned files).
- *1* (old behavior): The SEG\_NAME column is used.

## Notes

Enabling this option will not work if the table is partitioned. Most users will want to use the default value by omitting the parameter from the shadparm.ini file. Do not enable this parameter unless directed by Shadowbase Support.

## SHAD\_OVERRIDE\_SYSKEY

<b>Initial Version:</b>	V6.205	<b>Syntax:</b> SHAD_OVERRIDE_SYSKEY=<num>
<b>Last Change:</b>	--	
<b>Default Value</b>	2	
<b>Other Servers</b>	Cached Direct Writer   Cached TRS (ODBC – SQL/MX only)	
<b>Valid Settings</b>	0 = OFF. Conditions the ODBC/MX connection to disable exact SYSKEY replication; system generated values will be used. 1 = ON. Conditions the ODBC/MX connection to enable exact SYSKEY replication. SYSKEY values must be specified and the supplied values will be used (i.e., Shadowbase will use the source's (or incoming) SYSKEY value for INSERTs). 2 = NONE. No connection conditioning will be performed. The default behavior for the connection will be used (ie the target file system assigns the SYSKEY value on INSERT).	

## Description

This parameter enables or disables exact SYSKEY replication for cached TRS and Direct Writer objects replicating to 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

The ODBC/MX connection is specially conditioned to allow for specification of SYSKEY column values when exact SYSKEY processing is enabled. This conditioning affects the behavior of certain operations on the connection. This conditioning remains in effect even after the connection has been disconnected and will remain in effect if the connection is reused from the connection pool before being torn down. **As a result, it is highly recommended to use a dedicated DSN for Shadowbase when SHAD\_OVERRIDE\_SYSKEY is enabled (set to 1).** All TRS and Direct Writer objects using the same DSN must have the same setting for the SHAD\_OVERRIDE\_SYSKEY.

See the HP NonStop Shadowbase SQLMX Manual for information on adding and configuring a dedicated DSN for Shadowbase replication.


## SHAD\_STATUS\_STATS

<b>Initial Version:</b>	V3.930
<b>Syntax:</b>	SHAD_STATUS_STATS = <option>
<b>Last Change:</b>	--
<b>Default Value</b>	1
<b>Open Servers</b>	DOC Writer   OPCOL   OPLOG
<b>Valid Settings</b>	0 =Disable check pointing of current statistics with the SBMON / SEM status command. 1 =Enable check pointing of current statistics with the SBMON / SEM status command

## Description

This parameter controls how the object's status is determined when a SBMON or SEM command requests it and whether or not the object writes its current statistics to the COLLCONFIG file as part of the command processing.

The options are:

- **SHAD\_STATUS\_STATS = 0:** SBMON or the listener uses the list of running processes from the operating system to determine if the Shadowbase object is running. The command is not sent to the Shadowbase object and the statistics are

not written to the COLLCONFIG file.

- **SHAD\_STATUS\_STATS = 1 (default):** SBMON or the listener signals the Shadowbase object and waits for a response. The Shadowbase object will write its current statistics to the COLLCONFIG file.

**Notes:**

- 1) This parameter is not used when Shadowbase is running on a Microsoft Windows based platform. On Windows based platforms, the default value (1) is always used.
- 2) When using the default setting (SHAD\_STATUS\_STATS), SBMON or the listener will wait a configurable period of time per object for a response. This can result in significant delays when checking the status of many stopped Shadowbase objects.

**SHAD\_NSL\_CONVERT\_LEVEL (obsolete)**

As of this release, SHAD\_NSL\_CONVERT\_LEVEL is obsolete and will be ignored. It has been replaced by SHAD\_HEX\_CONVERT\_LEVEL.

## ***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 (V6.101 for all servers except HPE NonStop OSS; V6.200 for HPE NonStop OSS).

### **Logged for Null (0) Bytes in Text Fields**

```
ProcessLogEvent(): V$LOGMNR SQL statement includes NUL (0) characters:
  at scn=<number>; rbasqn=<number>; rbablk=<number>; rbabyte=<number>;
  sql:
  <sql statement>
```

*Only logged if STOP option is in effect:*

shadparm.ini parameter SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES controls the response to this error:

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = STOP (default, current value): collector stops.

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = SKIP: collector logs the SQL and continues, skipping the event.

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = SKIP-NOLOG: collector skips the event without logging.

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = KEEP: collector keeps the event with the NUL byte for downstream processing.

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = 0xNN: collector will replace the NULL byte with character specified.

SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES = OFF: collector does not scan for NULL bytes. **IMPORTANT:** this setting may cause data corruption if your data contains NULL bytes.

Performing Shutdown

- Cause:** The SQL statement generated by Oracle's V\$LOGMNR\_CONTENTS view is has NULL (0) bytes in text fields, and the SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES parameter is set to either STOP (default setting) or SKIP.
- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_SCAN\_FOR\_NULL\_BYTES. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first three lines will be logged and the event will be skipped.
- Recovery:** Restart using one of the other options to keep the NULL bytes, translate them to another value, or to skip without logging the message.

## Logged for Out of Order Events

```
ProcessBatch: Records received from the Oracle LogMiner query are out
of sequence
```

```
This indicates that there may be a loss of data on a restart of the
collector, or data may be applied out of order
```

```
You can avoid these sequencing issues by setting the
```

```
SHAD_OPACOLLECT_USE_ORDER_BY to 1
```

```
ProcessBatch(previous record): SCN[%d], RBASQN[%d], RBABLK[%d],
RBABYTE[%d], SSN[%d]
```

```
ProcessBatch(current record): SCN[%d], RBASQN[%d], RBABLK[%d],
RBABYTE[%d], SSN[%d]
```

*Only logged if STOP option is in effect:*

Performing shutdown

You can set SHAD\_OPACOLLECT\_CHECK\_LOGMNR\_SEQUENCE to:

```
'STOP'    to stop when out of order events occur (current setting)
'LOG'     to log and continue when out of order events occur
'OFF'     to inhibit the check
```

- Cause:** Event order checking is enabled and the events are not expected order (SCN, RBASQNC, RBABLK, RBABYTE, SSN).
- Effect:** Depends on the setting of SHAD\_OPACOLLECT\_CHECK\_LOGMNR\_SEQUENCE. If set to STOP, the collector will stop, logging the possible options as well as the other messages. If set to LOG, the information will be logged and the collector will continue.
- Recovery:** Remove the parameter from the SHADPARAM.INI (or set it to OFF) as the event order checking is no longer needed.

## Logged for SHAD\_OPACOLLECT\_ARCHIVE\_DIR Parameter Errors

```
Unable to list the Oracle REDO file archive list in directory
[<directory name>] specified by the SHAD_OPACOLLECT_ARCHIVE_DIR param
The output of the directory listing is in <file name>
Performing shutdown
```

*Or*

```
No Oracle REDO files were found in the directory [<directory name>]
specified by the SHAD_OPACOLLECT_ARCHIVE_DIR param
The output of the directory listing is in <file name>
Performing shutdown
```

- Cause:** The Collector was not able to read or find any files in the directory specified by the SHAD\_OPCOLLECTED\_ARCHIVE\_DIR parameter. It generates a listing of the files either by using the ls command (Linux & Unix systems) or the dir command (Windows), the output of the command is in the file specified.
- Effect:** Collector stops without processing any REDO logs.
- Recovery:** Correct the parameter setting if it is incorrect. Add the SHAD\_OPCOLLECT\_ARCHIVE\_FILES parameter if the directory is inaccessible to the collector. Remove the parameter if you are reading from the online logs.

### Logged for SHAD\_OPCOLLECT\_ARCHIVE\_FILES Parameter Errors

```
Unable to load the Oracle REDO file archive list from [<file name>]
specified by the SHAD_OPCOLLECT_ARCHIVE_FILES param
Performing shutdown
```

*Or*

```
No Oracle REDO files were found in file [<file name>] specified by the
SHAD_OPCOLLECT_ARCHIVE_FILES param
Performing shutdown
```

- Cause:** The Collector was not able to read the file specified by the SHAD\_OPCOLLECT\_ARCHIVE\_FILES parameter, or the file was empty.
- Effect:** Collector stops without processing any REDO logs.
- Recovery:** Correct the parameter setting if it is incorrect. Insure the file can be read by the Collector, and that it contains the list of REDO log file names, one name per line.

### Logged for Unmatched Rollback Events

```
ProcessLogMsg(commit): COMMIT event received for transaction with
unmatched ROLLBACK IUD events, ROLLBACK(ABORT) event was expected
SB Tx Id[<number>]; Oracle Tx Id (XIDUSN.XIDSLT.XIDSQN) [<value>]
First Event Collected: SCN[<value>],RBASQN[<number>],
```

```

RBABLK[<number>],RBABYTE[<number>],SSN[<number>]
ROLLBACK IUD Event: SCN[<value>],RBASQN[<number>],RBABLK[<number>],
RBABYTE[<number>],SSN[<number>]
COMMIT Event: SCN[<value>],RBASQN[<number>],RBABLK[<number>],
RBABYTE[<number>],SSN[<number>]

```

*Only logged for the STOP option:*

```

shadparm.ini parameter SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT
parameter controls the response to this condition:
SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT=STOP (current setting,
default): the collector stops.
SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT=OFF: No checking is done,
the collector continues without logging.
SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT=LOG: the collector logs
the event and continues. Some events that should
have been rolled back may be improperly replicated.
SHAD_OPCOLLECT_UNMATCHED_ROLLBACK_EVENT=TRACK-ALL (experimental,
memory intensive): the collector will maintain
and check all events associated with the transaction in
memory. This option may eliminate the issue, but requires
maintaining all events in memory.

```

- Cause:** The Collector read one or more unmatched rollback events
- Effect:** Depends on the SHAD\_OPCOLLECT\_UNMATCHED\_ROLLBACK\_EVENT parameter. If set to STOP, the collector logs the issue and stops. If set to LOG, the Collector continues.
- Recovery:** Review the parameter and set appropriately. Use V\$LOGMNR\_CONTENTS to look at the events associated with the error.

## Logged when the End SCN is Reached

```

ReadLogMsg()->Oracle Redo Log reading complete, reached SCN specified
by SHAD_OPCOLLECT_END_SCN. Current SCN[<value>], End SCN[<value>]
Shutting down

```

- Cause:** The Collector processed all audit up to the ending SCN specified by the SHAD\_OPCOLLECT\_END\_SCN parameter.
- Effect:** Collector stops normally.
- Recovery:** None required.



## Logged when Reading from Archived REDO Logs is Complete

```
ReadLogMsg()->Oracle Redo Log reading complete, completed selection
  from archived log files at SCN[<valued>]
Shutting down
```

**Cause:** The Collector processed all audit in the archived REDO logs.

**Effect:** Collector stops normally.

**Recovery:** None required.

## Logged when Selecting an Invalid Mode for Mining

```
Continuous collection from archive logs is configured, but one or more
  database instances does not have archiving enabled
Either enable archiving for all database instances, or use
  SHAD_OPACollect_MINING_MODE=CONTINUOUS_ONLINE to collect from online
  REDO logs only.
Shutting down
```

**Cause:** There is a configuration error. You have specified either CONTINUOUS\_BOTH or CONTINUOUS\_ARCHIVE for the collection mode for a database that is not enabled for REDO log archiving.

**Effect:** Collector stops.

**Recovery:** Either enable archiving of REDO logs for the database, or change the collection mode to CONTINUOUS\_ONLINE to collect from the online REDO logs only.

```
Error: SHAD_OPACollect_ARCHIVE_DIR must be specified for
  SHAD_OPACollect_MINING_MODE=OFFLINE_ARCHIVE
Shutting down
```

**Cause:** There is a configuration error. You have specified OFFLINE\_ARCHIVE mode for collection without specifying a location for the archive files.

**Effect:** Collector stops.

**Recovery:** Set the SHAD\_OPCOLLECT\_ARCHIVE\_DIR parameter to specify the archive files.

```
Error: SHAD_OPCOLLECT_ARCHIVE_DIR cannot be specified with
SHAD_OPCOLLECT_MINING_MODE=<mode>
Shutting down
```

**Cause:** There is a configuration error. You have specified an archive directory with a mode that uses continuous collection.

**Effect:** Collector stops.

**Recovery:** Either set the SHAD\_OPCOLLECT\_MINING\_MODE to OFFLINE\_ARCHIVE to use the files in the archive directory, or remove the SHAD\_OPCOLLECT\_ARCHIVE\_DIR parameter to use continuous collection.

## Logged for Cursor Allocation Issues

```
ParseStatement(): realloc() of cursor for length <length> for statement
<SQL statement> failed
```

```
ShadODBCCache::ProcessUpdateToInsert(): realloc() of cursor for length
<length> for statement <SQL statement> failed
```

```
ShadSqlldaCache::ProcessUpdateToInsert(): realloc() of cursor for length
<length> for statement <SQL statement> failed
```

```
UserExitCacheDB::ProcessCache(): out of memory: realloc(<length>) for
stmt_text failed
```

```
UserExitCacheDB::AssignCursor(): out of memory: realloc(<length>) for
stmt_text failed
```

```
UserExitCacheDB::FindInCache(): out of memory: realloc(<length>) for
stmt_text failed
```

**Cause:** These messages indicate that the TRS or DW have run out of memory while trying to allocate memory for statements in the statement cache.

- Effect:** The process stops.
- Recovery:** Reduce the statement cache size using the SHAD\_MAX\_CURSORS parameter in the shadparm.ini. If the problem continues, contact Shadowbase Support as the message may indicate a memory leak or other programming bug.

## Logged for HPE Shadowbase Log Source for Oracle Parsing Issues

```
OracleLogDLL::ProcessBatch(): ORACLE failed to generate valid SQL for
table <table>
This can occur if the table has been altered between when the event was
generated and when it was read.
Generated SQL: <sql statement>
```

*Only logged if STOP option is in effect:*

```
shadparm.ini parameter SHAD_OPCOLLECT_INVALID_SQL_OPTION controls the
response to this error:
SHAD_OPCOLLECT_INVALID_SQL_OPTION = STOP (default, current value):
collector stops
SHAD_OPCOLLECT_INVALID_SQL_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPCOLLECT_INVALID_SQL_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown
```

- Cause:** The SQL statement generated by Oracle's V\$LOGMNR\_CONTENTS view is invalid, most likely due to a schema change or the table being dropped prior to the events being generated. The table name is provided in the message, as is the generated SQL.
- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_INVALID\_SQL\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first three messages will be logged and the event will be skipped.
- Recovery:** Use one of the 'skip' options to skip the event.

```
OracleLogDLL::ProcessBatch(): ORACLE failed to generate valid SQL for
table <table>
This can occur if the table has been altered between when the event was
generated and when it was read.
```

```
Generated SQL: <sql statement>
```

*Only logged if STOP option is in effect:*

```
shadparm.ini parameter SHAD_OPACCOLLECT_INVALID_SQL_OPTION controls the
response to this error:
SHAD_OPACCOLLECT_INVALID_SQL_OPTION = STOP (default, current value):
collector stops
SHAD_OPACCOLLECT_INVALID_SQL_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPACCOLLECT_INVALID_SQL_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown
```

- Cause:** The SQL statement generated by Oracle's V\$LOGMNR\_CONTENTS view is invalid, most likely due to a schema change or the table being dropped prior to the events being generated. The table name is provided in the message, as is the generated SQL statement.
- Effect:** Depends on the setting of SHAD\_OPACCOLLECT\_INVALID\_SQL\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first three messages will be logged and the event will be skipped.
- Recovery:** Use one of the 'skip' options to skip the event.

```
OracleLogDLL::ExtractColNameAndValue(): Column <name> in table <table>
is missing from table definition
```

*Only logged if STOP option is in effect:*

```
shadparm.ini parameter SHAD_OPACCOLLECT_PARSE_ERROR_OPTION controls the
response to this error:
SHAD_OPACCOLLECT_PARSE_ERROR_OPTION = STOP (default, current value):
collector stops.
SHAD_OPACCOLLECT_PARSE_ERROR_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPACCOLLECT_PARSE_ERROR_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown
```

- Cause:** The SQL statement generated by Oracle's V\$LOGMNR\_CONTENTS view had a column that is not in the schema read by the collector, possibly because the table has been altered.

- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_PARSE\_ERROR\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first message will be logged and the event will be skipped.
- Recovery:** Restarting the collector will cause the new schema to be loaded and may correct the issue. Otherwise, use one of the SKIP options to skip these events.

```
OracleLogDDL::ExtractColNameAndValue(): Column <name> in table <table>
has no value associated with it
```

*Only logged if STOP option is in effect:*

```
shadparm.ini parameter SHAD_OPCOLLECT_PARSE_ERROR_OPTION controls the
response to this error:
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = STOP (default, current value):
collector stops.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown
```

- Cause:** The SQL statement generated by Oracle's V\$LOGMNR\_CONTENTS view did not have a value for a column in the schema read by the collector, possibly because the table was altered.
- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_PARSE\_ERROR\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first message will be logged and the event will be skipped.
- Recovery:** Restarting the collector will cause the new schema to be loaded and may correct the issue. Otherwise, use one of the SKIP options to skip these events.

```
ProcessLogMsg: OPERATION_CODE[<op code>-<expected op>] returned by
V$LOGMNR_CONTENTS does not match SQL_REDO contents

SB TxId[<txid>], SCN[<scn>],RBASQN[<rbasqn>],RBABLK[<rbablk>],
RBABYTE[<rbabyte>],SSN[<ssn>]
SQL_REDO[insert into "QASOURCE"."JRH"("PK","COL") values ('0','zz');]
```

*Only logged if STOP option is in effect:*

```

shadparm.ini parameter SHAD_OPCOLLECT_PARSE_ERROR_OPTION controls the
response to this error:
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = STOP (default, current value):
collector stops.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown

```

- Cause:** The HPE Shadowbase Log Source for Oracle collector validates that the OPERATION\_CODE column of the V\$LOGMNR\_CONTENTS query matches the statement in the SQL\_REDO column and logs this error if the columns do not match.
- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_PARSE\_ERROR\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first message will be logged and the event will be skipped.
- Recovery:** The validation that produces this error message is a sanity check. If this message occurs, it indicates either an issue with the collector or the V\$LOGMNR\_CONTENTS view and should be reported to Support. To continue replication, use SKIP or SKIP-NOLOG to skip these events.

```

ProcessLogMsg: Unexpected OPERATION_CODE[<operation code>] returned by
V$LOGMNR_CONTENTS query

```

```

SB TxId[<txid>], SCN[<scn>], RBASQN[<rbasqn>], RBABLK[<rbablk>],
RBABYTE[<rbabyte>], SSN[<ssn>]
SQL_REDO[insert into "QASOURCE"."JRH"("PK","COL") values ('0','zz');]

```

*Only logged if STOP option is in effect:*

```

shadparm.ini parameter SHAD_OPCOLLECT_PARSE_ERROR_OPTION controls the
response to this error:
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = STOP (default, current value):
collector stops.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP: collector logs errors and
continues, skipping the event.
SHAD_OPCOLLECT_PARSE_ERROR_OPTION = SKIP-NOLOG: collector skips the
event without logging.
Performing Shutdown

```

- Cause:** The HPE Shadowbase Log Source for Oracle collector issues a SELECT statement that returns a limited number of operation types (INSERT, UPDATE, DELETE, COMMIT, ROLLBACK). It received an operation code that should not have been returned by the SELECT statement.
- Effect:** Depends on the setting of SHAD\_OPCOLLECT\_PARSE\_ERROR\_OPTION. If set to STOP or unspecified, the collector will stop, logging the possible options as well as the other messages. If set to SKIP, only the first message will be logged and the event will be skipped.
- Recovery:** The validation that produces this error message is a sanity check. If this message occurs, it indicates either an issue with the collector or the V\$LOGMNR\_CONTENTS view and should be reported to Support. To continue replication, use SKIP or SKIP-NOLOG to skip these events.

## Logged for Invalid HPE Shadowbase Log Source for Oracle Collection Parameters

```
SHAD_OPCOLLECT_INSERTS, SHAD_OPCOLLECT_UPDATES, and
SHAD_OPCOLLECT_DELETES are all disabled. At least one must be enabled
for replication
Performing shutdown
```

- Cause:** Collection of insert, update, and delete events for HPE Shadowbase Log Source for Oracle collection was disabled. As this means no events will be collected, the collector stops.
- Effect:** The collector stops.
- Recovery:** Modify the shadparam.ini to allow collection of at least one type of event.

```
SHAD_OPCOLLECT_USE_INCLUDE_USERS, SHAD_OPCOLLECT_USE_EXCLUDE_USERS,
SHAD_OPCOLLECT_USE_EXCLUDE_TABLES, and
SHAD_OPCOLLECT_USE_INCLUDE_TABLES are all disabled.. At least one must
be enabled for replication.
Performing shutdown
```

- Cause:** None of the configuration tables used to determine the events to be included in replication are specified.

**Effect:** The collector stops.

**Recovery:** Modify the shadparm.ini to use at least one of the criteria configuration tables.

### **Logged for Missing Data**

This release has several additional checks for missing data for HPE Shadowbase Log Source for Oracle collection. The following messages have been added for the checks.

```
SHAD_OPCOLLECT_USE_INCLUDE_USERS, SHAD_OPCOLLECT_USE_EXCLUDE_USERS,  
SHAD_OPCOLLECT_USE_EXCLUDE_TABLES, and  
SHAD_OPCOLLECT_USE_INCLUDE_TABLES are all disabled.. At least one must  
be enabled for replication.  
Performing shutdown
```

**Cause:** None of the configuration tables used to determine the events to be included in replication are specified.

**Effect:** The collector stops.

**Recovery:** Modify the shadparm.ini to use at least one of the criteria configuration tables.



## ***Known Problems Remaining***

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 (*sbolog*) has the following limitations:
- a. We are not supporting source collection Oracle 10g Release 2. An issue configuring proper permissions for Oracle user accounts other than SYSDBA prevents proper operation. We are continuing to investigate this issue. If you have need of log based source collection for Oracle 10gR2, please contact your HPE NonStop Sales and Support team or check the Shadowbase knowledge base to see if a solution has been posted. Trigger based source collection is available for Oracle 10gR2 as an alternative method.
  - b. 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`.
  - c. HPE Shadowbase Log Source for Oracle does not support replication of tables with `CLOB` or `BLOB` columns.
  - d. 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.
  - e. Shadowbase Log Source for Oracle currently only collects DML events (Inserts, Updates, and Deletes).
  - f. 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.
  - g. 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.
  - h. 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.

- i. 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 SHADPARM.INI
```

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

```
SHAD_REMOTE_MAX_EVENTS=200
```

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

10. The sample DOC reader code (SBDOCRD) for reading non-cached DOCS fails. Both a sample program that reads and prints the DOC, as well as source code to allow the user to process the data in the DOC. Both the sample program and code fail.

This code has limited functionality and is deprecated. However, if you need a working copy of SBDOCRD, please contact Support.

11. 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.
12. 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.

13. 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.

14. 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.
15. 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.
16. 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.

17. 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.

18. SBFILe does not display text fields containing the string 'N', 'NU', 'NUL', or 'NULL' correctly – the enclosing single quotes are left off. For example, if TEST\_TABLE has four varchar fields, SBFILe will display the statement:

```
INSERT INTO TEST_TABLE (C1,C2,C3,C4) VALUES ('N', 'NU', 'NUL', 'NULL')
```

incorrectly as:

```
INSERT INTO TEST_TABLE (C1,C2,C3,C4) VALUES (N, NU, NUL, NULL)
```

**Note:** This is a display issue only. The data will be correctly applied to the database.

19. 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.

20. The example user exit make file shipped with the OSS version of HPE Shadowbase for Other Servers references the release version of USRXDLLINIT.o library, which is not shipped with OSS. You should modify the make file to use the debug version of the libraries instead.
21. 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.

22. Shadowbase determines the target database for ODBC databases using the ODBC driver name. If it does not recognize the driver, it requires that the generic error codes parameter SHAD\_ODBC\_GENERIC\_CODES be set to specify the error codes return by the database for common issues such as deleting a non-existent record. If it recognizes the target database, the SHAD\_ODBC\_GENERIC\_CODES parameter does not have to be specified as Shadowbase already has the codes. In general, Shadowbase can determine supported database versions, eliminating the need for the parameter in most cases. However, for MySQL 5.6, it does not recognize the driver and, instead, logs the following message and stops:

```
2015-05-27 11:13:27 -[29645] Error: set_odbc_errcodes():Non-supported ODBC
Driver detected: LIBMYODBC5A.SO; terminating.
2015-05-27 11:13:27 -[29645] Critical Error: Client Connection to target
DBMS[] failure
2015-05-27 11:13:27 -[29645] Critical Error: Performing Shutdown
```

In this case, you must use the SHAD\_ODBC\_GENERIC\_CODES parameter to specify the error codes.

23. 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_OPCOLLECT_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_OPCOLLECT_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.

24. 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 updated. Added columns may not be picked up, and deleted columns may result in Shadowbase stopping.
25. The TRS fails without warning when replicating a column to Oracle that starts with "TO\_DATE" if DBS Mapping is enabled. A workaround is to change the SQL type to SQL92.
26. 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.
27. 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.
28. 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 '% \';
```

## ***Installation Instructions***

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

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