

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	HP Shadowbase for Other Servers
Release:	Gravic Version 6.101 HP Shadowbase for Other Servers: T1123-AAA (SB REPL/OSS) T1124-AAA (SB REPL/HP UX) T1125-AAA (SB REPL/SUSE LINUX) T1126-AAA (SB REPL/RHEL LINUX) T1127-AAA (SB REPL/SOLARIS X86) T1128-AAA (SB REPL/SOLARIS) T1129-AAA (SB REPL/MS WINDOWS) T1130-AAA (SB REPL/AIX)
Date:	June 12, 2015
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2015 (www.gravic.com)
File Name:	IPM6101_other_servers.docx

NOTE: This softdoc covers new features and corrected problems for HP Shadowbase for Other Servers, Version 6.101. It is available as an Adobe PDF file (.PDF). Copies of the PDF file reader can be freely downloaded from www.adobe.com

NOTE: Softdoc files are named IPMnnnnx_Other_Servers (where nnnnx is the HP Shadowbase for Other Servers version number).

Table of Contents

Table of Contents	1
Table of Figures	2
Special Notes for Version 6.101	3
New Features (since version 6.100) in version 6.101	5
SBMON EDIT and EDIT! COMMANDS	6
Edit Command	6
EDIT! Command	7
Removal of COLL Command.....	8
Redefined SQL Schema for SHAD_LOG and SHAD_REJECT tables.....	9
SHAD_LOG Table	9
SHAD_REJECT	9
DBS Mapping Support for double-quote (“) and back-tick (`).....	11
Oracle Table Loading for Oracle Log Source Collection.....	12
Introduction.....	12
Setup	12
Installation.....	12
Operation.....	14
Adding more load tables	14
Oracle Log Source Collection Support for Primary Key Updates.....	15
Windows “Shadowbase for SBMON” service startup BAT file	16
DOC File Reference Numbers	17
SBFILE SETRESTARTPOINT Command.....	18
SETRESTARTPOINT Command	18
New/Modified shadparm.ini Parameters	19
SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION.....	19
Problems Corrected (since version 6.100) in version 6.101	19
Known Problems Remaining	23
New User Error Messages (since version 6.100) in version 6.101:.....	29
Installation Instructions (optional "other" servers).....	30

Table of Figures

Figure 1 - SHAD_LOG Table Redefinition.....	9
Figure 2 - SHAD_REJECT Table Redefinition.....	10
Figure 3 - sbfile SETRESTARTPOINT example	18

Special Notes for Version 6.101

1. HP Shadowbase for Other Servers version 6.101 release has been validated in the following environments (all environments are 64 bit unless otherwise noted):
 - Microsoft SQL Server 2005 through 2012 database target replication.
 - Microsoft SQL Server 2005 through 2012 database source replication.
 - SAP Sybase ASE database target replication for the following versions platforms:
 - Windows (32 bit) Sybase 15
 - Solaris 10,11 Sybase 15
 - Solaris x86 Sybase 15
 - Oracle database target replication for the following versions and platforms:
 - Solaris 10, 11 Oracle 10gR2, 11gR2, 12c
 - Solaris x86 10, 11 Oracle 10gR2, 11gR2, 12c
 - AIX (5.3 to 7.1) Oracle 10gR2, 11gR2, 12c
 - HP-UX Itanium 2 (11.23 to 11.31): Oracle 10gR2, 11gR2, 12c¹
 - Red Hat Enterprise Linux² (5 to 6.1): Oracle 10gR2, 11gR2, 12c
 - SUSE Linux Enterprise Server 10: Oracle 10gR2, 11gR2, 12c
 - Oracle database source replication for the following versions and platforms:
 - Solaris 10, 11: Oracle 10gR2 (trigger only)
Oracle 11gR2 (trigger & log-based)
 - Solaris x86 10, 11: Oracle 10gR2 (trigger only)
Oracle 11gR2
 - AIX (5.3 to 7.1): Oracle 10gR2 (trigger only)
Oracle 11gR2²
 - HP-UX Itanium 2 (11.23): Oracle 10gR2¹ (not supported)
(11.31): Oracle 10gR2¹ (trigger only)
(11.23): Oracle 11gR2¹ (log-based only)
(11.31): Oracle 11gR2¹ (trigger & log based)
 - Red Hat Enterprise Linux² (5 to 6.1): Oracle 10gR2 (log-based only)
Oracle 11gR2 (trigger & log based)
 - SUSE Linux Enterprise Server 10: Oracle 10gR2 (log-based only)

¹ Support for Oracle 11gR2 on HP-UX (Itanium2) requires patch PCHO_40381 be installed.

² RHEL Version 5 requires update 11.

Oracle 11gR2 (trigger & log based)

- IBM DB2 Database target replication for the following versions and platforms:
 - AIX (5.3 to 7.1) DB2 9.x
 - Red Hat Enterprise Linux³ (5 to 6.1) DB2 9.x
- MySQL Database target replication for:
 - Red Hat Enterprise Linux³ (5 to 6.1) MySQL 5.x

Note: If you require a version of HP Shadowbase for Other Servers for an environment not listed above, inquire with your HP NonStop Sales and Support team. Additional platforms and databases not on the HP price book are typically available on a custom Statement of Work (SOW) basis.

2. Due to licensing changes in Version 6.100, existing installations of Shadowbase prior to that version will require a new password file in order to run after the upgrade.

³ RHEL Version 5 requires update 11.

New Features (since version 6.100) in version 6.101

HP Shadowbase for Other Servers, version 6.100 contains the following new/enhanced features and functionality.

- SBMON EDIT and EDIT!
- Removal of COLL command support
- Redefined SQL Schema for SHAD_LOG and SHAD_REJECT tables
- DBS Mapping support for double-quote (") and back-tick (`)
- Oracle Log Source support for loading tables
- Oracle Log Source support for primary key updates
- Windows "Shadowbase for SBMON" service startup BAT file
- DOC File Reference Numbers
- sbfile SETRESTARTPOINT command

SBMON EDIT and EDIT! COMMANDS

Edit Command

The SBMON EDIT command will now check for and display any running objects prior to executing. Changing the configuration parameters while objects are running is dangerous and can lead to unpredictable results. If the check uncovers any running Shadowbase processes, the SBMON will give the user the option to skip editing.

This check is done in the following steps:

- 1) SBMON issues a message indicating that is getting the status of the configured objects, followed by a “Please wait ...” prompt:

```
Acquiring STATUS of all configured objects
Please wait ...
```

- 2) SBMON performs a status check of the objects. This step may take a long time if there are many configured objects that are not running but were not stopped by the SBMON STOP command.
- 3) When all objects have been checked, SBMON will add “Done” to the “Please wait ... “ line”:

```
Acquiring STATUS of all configured objects
Please wait ... Done
```

- 4) If there are any configured objects that are currently running, the user will be warned and the running objects will be listed, followed by a prompt to continued:

```
Warning: Altering configuration attributes
while replication is in operation can negatively
impact replication

<object>   running
<object>   running
.
.
<object>   running

Do you wish to continue? (Y, [N]):
```

If the user answers 'Y', the EDIT command will proceed as in the past. **NOTE:** *Only select 'Y' when directed by Shadowbase support. Selecting 'Y' can lead to unpredictable results.*

If the user answers 'N', the EDIT command will terminate with no changes and the user will be able to stop the running processes.

- 5) If there are no running objects, the EDIT command will continue without prompting.

EDIT! Command

Checking the status of configured objects in large configurations may take a significant amount of time. The EDIT! Command can be used to skip the new check and implements the original behavior of the EDIT command for releases prior to v6.101. Using EDIT! by-passes the configured object status check and goes directly to the configuration editing prompts for the expressed configuration object.

Note that the command is "EDIT!" with no space between the characters. The syntax of the command is:

```
EDIT! <object>
```

```
<object> = Name of the configured object to be modified.
```


Removal of COLL Command

The COLL command is obsolete and has been removed. The SBMON HELP command has been modified to remove the display of COLL info:

```
+help
Syntax: HELP <command>

<command> = Name of command; select command display.
Commands are: ADD          EDIT          ENV          EXIT
FILESTATS  HELP          LICENSE     LIST          LISTSTATUS
LOG         REMOVE       ROLL        START         STATUS
STOP       /

+
```

The SBMON HELP COLL command now reports that command is invalid and not supported:

```
+help coll
Syntax: COLL
Invalid command. Not supported.

+
```

Additionally, the COLL commands themselves are now disabled:

```
+COLL OPEN DOC
Invalid command
+COLL CLOSE DOC
Invalid command
+COLL KEEPALIVE
Invalid command
+COLL BEGIN TRANS
Invalid command
+COLL
Invalid command
+
```

Redefined SQL Schema for SHAD_LOG and SHAD_REJECT tables

The SHAD_LOG and SHAD_REJECT SQL Table schema definitions have been modified to avoid reserved words defined by certain DBMS's by changing the following column names in both tables:

Pre v6.101 Name	New Name
SQLSTATE	STATECODE
SQLCODE	ERRCODE

The following table shows the supported SQL Table schema definitions for both the SHAD_LOG table (Figure 1) and the SHAD_REJECT table (Figure 2) for Version 6.101. Note that the schema change affects only column names. The column data type and column purpose are unchanged. The new column names are italicized for emphasis.

SHAD_LOG Table

Column Name	Data Type	Description
MSGLEVEL	Integer	Message severity level
LOGTIMESTAMP	Char(26)	Logged message generation timestamp
REPOBJTYPE	Char(10)	Type of replication object producing message
REPOBJNAME	Char(5)	Name of replication object producing message
<i>STATECODE</i>	Char(5)	The DBMS <i>sqlstate</i> code associated with message (if available, NULL otherwise).
<i>ERRCODE</i>	Integer	The error code associated with message (if available, NULL otherwise)
MSGTEXT	Varchar(max)	The message text

Figure 1 - SHAD_LOG Table Redefinition

Not all *SHAD_LOG* messages will contain DBMS generated STATECODE and/or ERRCODE column values as Shadowbase generates messages that are not associated with a DBMS event. Therefore, these fields will be populated by varying sources (errno, CTREE error code, ISAM error code, socket error code, etc.) and will use NULL for cases where they are not used.

SHAD_REJECT

Column Name	Data Type	Description
MSGLEVEL	Integer	Message severity level
LOGTIMESTAMP	Char(26)	Logged message generation timestamp
REPOBJTYPE	Char(10)	Type of replication object experiencing fault
REPOBJNAME	Char(5)	Name of replication object experiencing fault
<i>STATECODE</i>	Char(5)	The <i>sqlstate</i> code associated with fault (if available, NULL otherwise)
<i>ERRCODE</i>	Integer	The DBMS error code associated with fault
SQLMSG	Varchar(max)	The DBMS generated message text associated with <i>sqlcode</i> value
SQLSTMT	Varchar(max)	The replication SQL Statement executed which triggered fault

Figure 2 - SHAD_REJECT Table Redefinition

Shadowbase generates the *SHAD_REJECT* messages in direct response to a DBMS event. As a result, all *SHAD_REJECT* messages will contain the DBMS *STATECODE* and or *ERRCODE* values.

DBS Mapping Support for double-quote (“) and back-tick (`)

HP Shadowbase for Other Servers DBS Mapping now supports double-quote (“) and back-tick (`) characters. These are useful, for example, when a column name matches a SQL language reserved word.

The following DBS Mapping Commands now supports double-quote and back-tick characters:

[DBS+]
[NAME+]
[FLD-]
[FLD+]
[FLD++]
[SYSTEM+]
[SYSTEM++]

The following DBS Mapping Commands do not support the use of double-quote or back-tick characters:

[OP-]
[OP+]
[WHERECLAUSE-]

Oracle Table Loading for Oracle Log Source Collection

The Shadowbase Oracle Log Source Loader is used to load target database files and tables in order to match their corresponding source database entities. This section describes the basic functionality of Shadowbase Oracle Log Source Loader, how to configure for loading, and also how to monitor and manage these environments.

Introduction

This outlines how to run the Shadowbase Oracle log source loader procedure in order to run initial loads of Oracle source data to a target database using the running Shadowbase replication engine. This loading procedure is designed to be performed while Shadowbase is up and running, so it is an on-line loading engine. There will be a separate loader procedure for each table which will construct and perform updates on the source data. These updates will be created in such a way as to not make any change to the data itself (a no-operation “no-op” update that sets a column to its existing value). These updates will be picked up by the running Shadowbase OPCOL and replicated to the target database.

IMPORTANT NOTE: Throughout this document, there are references to single quotes in some of the example commands to run. They may appear as tildes (~), but be sure to use single quotes when running these commands.

Setup

Before installing the Shadowbase Oracle log source loader, it is necessary to set the `utl_file_dir` parameter in the source Oracle database. While logged on as the `sysdba`, run the following command to see what it is currently set to:

```
show parameter utl_file_dir
```

If the value is blank, or if it is set to a value other than the Shadowbase data directory, it will need to be altered. In order to alter the parameter, use the following syntax:

```
alter system set utl_file_dir='<fully qualified Shadowbase data
directory>' scope=spfile;
```

For example, if the instance of Shadowbase on this system is installed on `/home/shaduser/shadowbase`, then the command should look like the following:

```
alter system set utl_file_dir='/home/shaduser/shadowbase/data'
scope=spfile;
```

After setting this parameter, the database will need to be stopped and restarted. Once it is restarted, the database is ready to have the Shadowbase Oracle log source loader installed.

Installation

In order to install the Shadowbase Oracle log source loader, first log into the system as the user that runs Shadowbase. Make sure `$SHAD_BASE` is set correctly, then change into that directory. From there, log into a sqlplus session as the same user as the

Shadowbase OPCOL that is connecting into the database. While logged in as this user, be sure the session has “set serveroutput on” to see any necessary output messages. Run the following command to install the Shadowbase Oracle package:

```
@bin/shadorainstall.unix.sql <fully-qualified SHAD_BASE location>
```

For example:

```
@bin/shadoraisntall.unix.sql /home/shaduser/shadowbase
```

This will install the package required in order to configure the log source loader. After this script finishes, run the following command for each table that needs to be configured to have its data loaded:

```
exec Shadpack.ShadDMLBuild('<schema name>', '<table name>');
```

For example:

```
exec Shadpack.ShadDMLBuild('SOURCEUSER', 'TABLE1');
```

This command collects information about the table and its columns and saves it into tables that Shadowbase will use to generate the load scripts.

IMPORTANT NOTE: The Shadowbase Oracle log source loader does not support tables that only have key columns. If a table does not have any non-key columns, do not attempt to configure the log source loader for this table.

Once the ShadDMLBuild command has been run for all tables, run the following command for the same set of tables:

```
exec Shadpack.ShadLogLoadBuild('<schema name>', '<table name>', '<fully qualified Shadowbase data directory>');
```

For example:

```
exec  
Shadpack.ShadLogLoadBuild('SOURCEUSER', 'TABLE1', '/home/shaduser/shadowbase/data');
```

This command creates a procedure create script that it saves to the \$SHAD_BASE/data directory.

Once this command has been run for each table, run the following command for each table:

```
@data/SHAD_LOG_LOAD_<schema name>_<table name>.sql
```

For example:

```
@data/SHAD_LOG_LOAD_SOURCEUSER_TABLE1.sql
```

This will create the log load procedure for each table. Once it has been run for each table, Shadowbase Oracle log source loader is now installed and ready for use.

Operation

In order to load data from an Oracle source table into Shadowbase, first be sure the source OPCOL is running and collecting data from the database. Log into sqlplus as the OPCOL user and run the following command to load data from a source table into the OPCOL:

```
exec shadpack.shadlogload('<table name>',<batch level>,<isolation level>,<wait time>);
```

where <table name> is the name of the table, or 'LOADALL' in order to load data for all configured tables

<batch level> is the number of events to put into each transaction, e.g. 100

<isolation level> is either 1 for READ COMMITTED or 2 for SERIALIZABLE

<wait time> is the number of seconds to delay between each transaction, 0 for no delay

For example:

```
exec shadpack.shadlogload('TABLE1',100,1,0);
```

The result of this procedure will load all rows as update statements into the OPCOL DOC which can then be replicated to the target database.

At any time, it is possible to view the current loader status of all of the configured source tables. By running the following command, it will display a summary of all tables:

```
exec shadpack.shadstatus('LOADER');
```

This will display the loader name as well as the number of rows currently loaded for each source table. If a load is in progress, it will display the number of rows that have been loaded during the current load. If a load is not in progress, it will display the number of rows that were loaded during the most recent load.

Adding more load tables

It may be desirable to add the ability to load more tables after initial configuration. If so, simply run the ShadDMLBuild and ShadLogLoadBuild scripts for the given table, then obey the new procedure create script by running the @data/SHAD_LOG_LOAD_<schema name>_<table name>.sql script. This will add the ability to load this new table to the replication environment.

Oracle Log Source Collection Support for Primary Key Updates

Oracle Log Source Collection now supports processing of update events that modify a table's primary key index. A new parameter, *SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION*, has been added to the *shadparam.ini* to control how these updates are handled. Please refer to the corresponding section below for specific details.

Windows “Shadowbase for SBMON” service startup BAT file

The NTSBSERV.exe process used on the Windows version of Shadowbase was modified to isolate startup issues such as encountered in cluster failover situations. Specifically, when installed as the “Shadowbase for SBMON” Windows service, it will run *sbstart_sbmon.bat* instead of *sbstart.bat* when the SHAD_SERVICE_RUN_SBSTART=1 parameter is set. So, it is much easier to provide startup commands that are only executed once at startup such as needed in a cluster failover to another node. In this case, the *sbstart_sbmon.bat* file would be edited to kill or stop (via SBMON) previously running objects.

DOC File Reference Numbers

Each DOC file sequence number may contain, in part, transaction data that started in an earlier sequence number. For instance, as shown below in the SBMON filestats output for an object called “DOC1”, a transaction might start in DOC #221, have some DML operations in #222 and finally finish (i.e. commit) in #223. In this case, the reference sequence number for all three DOCs is #221.

SBMON Filestats command:

```
+filestats DOC1
Starting sequence number [1] :
Ending sequence number [999] :
```

Seq Num	Begins	Commit	Aborts	Inserts	Updates	Deletes	Multi OPs	Total Events	DataFile Size (K)	Ref Seq
221	1	0	0	1	0	0	0	1	63	221
222	0	0	0	1	0	0	0	1	63	221
223	0	1	0	2	0	0	0	2	63	221
<after edit of sequence number in SBMON>										
500	0	0	0	0	0	0	0	0	63	***

Certain circumstances, such as editing a DOC sequence number in the above case to #500, or temporarily during a DOC roll, can leave the new DOC sequence number in an unknown reference sequence number state. We reflect that situation with three stars “***” in the Filestats command output. Previously, the reference sequence number was shown as the DOC’s sequence number leading to situations where a DOC Cleaner or manual deletion of old DOCs might clean up too many generations.

SBFILE SETRESTARTPOINT Command

The SBFILE utility of HP Shadowbase for Other Servers, version 6.101 now supports setting/resetting the restart point manually via a new command, SETRESTARTPOINT.

SETRESTARTPOINT Command

The sbfile SETRESTARTPOINT command modifies the configuration record of the select Shadowbase replication object. The command replaces the Transaction Key field contents of the configuration record with the value entered. A validation that the replication object's configuration record exists and that the Transaction Key entry is 16 bytes are performed.

Note: The SETRESTARTPOINT command should only be used in extreme circumstances as data loss or target database corruption may occur. It is recommend that you call HP Shadowbase Support prior to utilizing the SETRESTARTPOINT command.

To run, enter:

```
sbfile SETRESTARTPOINT <Object Name> <Transaction Key>
```

<Object name>

is the name of an existing replication object

<Transaction Key>

is a 16 character hexadecimal number.

Sample SETRESTARTPOINT command use output:

```
$ sbfile SETRESTARTPOINT TRS1 0123456789ABCDEF  
Setting Restart Point from [0000000000000000] to [0123456789ABCDEF]
```

Figure 3 - sbfile SETRESTARTPOINT example

New/Modified shadparm.ini Parameters

SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION

Initial	Syntax: SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION=option
Version: V6.101	
Last Change: --	
Default Value	STOP
Open Servers	Oracle Log Source Collector (sborlog)
Valid Settings	One of <i>STOP</i> , <i>SKIP</i> , <i>SKIP-NOLOG</i> , or <i>CONVERT</i>

Description

This parameter specifies the behavior of the Oracle Log Source Collector (sborlog) when it reads an event that updates the primary index of a table (the primary index typically has the 'unique' attribute for the key values and is often used for direct positioning in WHERE clauses). The options are:

- *STOP* (default, previous behavior): collection will stop with an error message when a primary key update event occurs on a replicated table.
- *SKIP*: collection will log but not collect primary key update events on replicated tables.
- *SKIP-NOLOG*: collection will skip without logging primary key update events on replicated tables.
- *CONVERT*: primary key updates will be converted into a delete of the original row on the target, followed by an insert of the new values.

Notes

Most users will want to set the option to *CONVERT*, to apply the events to keep the data consistent between source and target. However, the target tables should be examined for potential issues with converting the update into a delete and insert operation. These issues may include:

- Triggers on the target table (especially those that implement cascading deletes), and
- Relational constraints on the target table.

Problems Corrected (since version 6.100) in version 6.101

1. The License Feature validation failed to catch an invalid License file for the License Feature BIDIR. This issue has been corrected.
2. The License Feature validation failed to distinguish OS_RHEL from OS_SUSE. This issue has been corrected.

3. Shadowbase Oracle target replication was not properly closing SQL cursors when attempting to get primary key constraint information for tables with no primary key constraints defined. This failure caused an *ORA-01000: maximum open cursors exceeded* error during target replication, causing the TRS to fail. The issue has been corrected and an *ORA-01000: maximum open cursors exceeded* should not occur regardless of whether primary key constraint information is available or not.
4. The cached Shadowbase Transaction Replay Server (TRS) and Direct Writer (DW) trapped if replicating to a table whose name is a duplicate of a table owned by a different user in the database, such as the case where replicating to USER1.MYTABLE in a database where USER2.We reMYTABLE also exists. The TRS and DW now correctly replicate in this case.
5. The HP Shadowbase Other Servers MSSQL Source Collector failed to identify NULL column values correctly, resulting in failures on the target when the data replicated. This issue has been corrected.
6. The SHAD_LOG_FILE_SIZE and SHAD_LOG_FILE_SPACE parameters did not work correctly and had little or no effect in controlling errlog log file size and space usage. This issue has been corrected.
7. The HP Shadowbase Other Servers Oracle Log-based Source Collector fail to validate Supplemental Logging settings within an Oracle RAC configuration. This issue has been corrected.
8. In Oracle RAC environments, the Oracle Log-based Source Collector did not properly load all of the Oracle log files during event collection, resulting in lost data. In previous versions, the collector limited the collection to one log file based upon the Oracle SCN number associated with the log files. This worked well in non-RAC environments; however, in RAC environments, the log files for each instance overlap and multiple log files need to be loaded. The collector now loads the correct log files and works properly in RAC environments.
9. Issuing two DOC ROLL commands through SBMON caused a CTREE error 26 with the following sequence of errors:

```
Starting DOC roll for |DOC | - DO NOT INTERRUPT!  
2015-04-12 07:44:23 -[10352] Info: Starting roll. Collector name  
[DOC], Created new DOCs= 1, Current sequence #70, New sequence will be  
#71.  
2015-04-12 07:44:24 -[10352] Info: Successfully copied incomplete  
transaction IDs only to next DOC [opType 32].  
2015-04-12 07:44:24 -[10352] Info: Successfully copied incomplete  
transaction IDs only to next DOC [opType 33].  
2015-04-12 07:44:25 -[10352] Error: CloseCTISAM() : Error 26 closing  
c-tree data file  
2015-04-12 07:44:25 -[10352] Error: CloseCTISAM() : Error 26 closing  
c-tree data file  
2015-04-12 07:44:25 -[10352] Info: Successful roll. Currently  
operating on sequence #71.  
DOC roll successful.
```

This issue has been corrected.

10. Shadowbase used 2-byte integers for the generic error codes specified by the SHAD_ODBC_GENERIC_CODES. Sybase IQ uses 4-byte error codes, resulting in truncation. Shadowbase now uses 4-byte integers for the generic error codes. In addition, the parsed values for the string of error codes is now displayed in the log:

```
SHAD_ODBC_GENERIC_CODES set to <code string>  
ODBC_SQLTYPE = <ODBC_SQLTYPE value>  
ODBC_ERR_ROW_NOT_FOUND = <ODBC_ERR_ROW_NOT_FOUND value>  
ODBC_ERR_COL_NULL = <ODBC_ERR_COL_NULL value>  
ODBC_ERR_INS_NULL = <ODBC_ERR_INS_NULL value>  
ODBC_ERR_DUP_KEY = <ODBC_ERR_DUP_KEY value>  
ODBC_ERR_INS_DUP = <ODBC_ERR_INS_DUP value>  
ODBC_ERR_DUP_OBJ = <ODBC_ERR_DUP_OBJ value>  
ODBC_ERR_DUP_IDX = <ODBC_ERR_DUP_IDX value>
```

11. Prior to this release, the cached version of the Transaction Forwarding Server (TFS) did not support user exits. The TFS now supports user exits.
12. Under certain circumstances, the TFS incorrectly advanced the restart point on a communication failure, resulting in the possible loss of the data in transit (sent but not acknowledged by the target system). The TFS now correctly resends the data that was in transit during the communication failure when restarted.
13. Under certain circumstances, the DOC Cleaner would stop deleting DOCs. In certain circumstances, the DOC Cleaner would not be able to read a reference sequence number from the DOC file and it would pause indefinitely. This issue has been fixed.

14. SBMON incorrectly started multiple instances of the DOC Cleaner for a multi-ported DOC. This issue was largely benign, as all but the first instance would exit with an “Error: Consumer number out of range for <DOC>” message in the log. SBMON now correctly only starts one instance of the DOC Cleaner.
15. The SBMON list command for DOC Cleaners showed the default number of DOC generations to save, not the number actually configured. It now correctly shows the configured number.
16. The Oracle Trigger-based Source Collectors (OPCOL) all incorrectly tried to roll the DOC file. Only once instance should roll the DOC files. Under normal processing, this error is benign, however, if one or more OPCOL processes fell sufficiently behind, DOC files could be incorrectly deleted and data lost. The OPCOL process will now correctly roll the DOC files.
17. Shadowbase processes configured as Windows services did not start and log messages properly. The processes created the log file, however, the processes did not write any entries and the log file remained empty, even though the service was started and the processes were running. The processes will now correctly log messages when configured as Windows services.
18. The cached Direct Writer (DW) and Transaction Replay Server (TRS) leaked memory when the number of statements exceeded the configured cache size. The DW and TRS processes will no longer leak memory in this case.
19. Under certain circumstance, the statements recorded in the Log Server DOC for rejected SQL statement were incorrect. This issue has been corrected.
20. The DOC reader program failed when reading DOCs in Version 6.100 and beyond. This issue has been corrected.
21. This release has several improvements in parameter validation and logging for the Oracle log-based collector parameters. The collector now correctly identifies more types of invalid parameters.
22. This release corrects an issue replicating columns that have the ASCII value ‘NULL’. In previous releases, if DBS mapping or user exit processing is enabled, and if the column’s defined length is 4 characters, and if the value for the column is ‘NULL’, the TRS or DW may incorrectly convert the text value and attempt to set the target column to the SQL value NULL.
23. Previous releases did not correctly track the error log size, causing delayed log file rolls and log file sizes to grow significantly larger than configured. This behavior was particularly acute in verbose mode (diagnostic level set to 3 or higher) and when statistics were logged. Error logs will now roll correctly on size.

24. In previous releases, if Shadowbase got an error attempting to read the key column information for Oracle tables, it did not correctly close the Oracle cursor associated with the query to retrieve the key columns. This could result in an “ORA-1000: maximum number of cursor exceeded” error and Shadowbase replication failing. Shadowbase will now properly close the cursor preventing the resource leak.
25. The FLD+ mapping to combine two columns into a new column did not work properly and the new column was not added. This issue has been corrected.
26. For multi-byte character sets, Shadowbase could replace unprintable characters with periods (‘.’) prior to applying to the database. This occurred during invalid number processing, when the original attempt to apply the data to the database failed due to invalid data. The second attempt replaced the unprintable characters with the periods prior to applying the event to the database. This issue has been fixed.
27. The Oracle log-based collector did not differentiate between column values that were NULL and values that had the string ‘NULL’; the collector treated both as SQL NULL values. It now handles both values properly.
28. The Microsoft SQL Server trigger-based Collector did not properly detect and record NULL columns. Shadowbase now properly detects and collects NULL values for downstream application.
29. Internal testing uncovered issues using DOC blocking with the Oracle log-based collector, resulting in lost or corrupted data. These issues have been corrected and the collector’s DOC blocking works correctly.
30. Previous releases did not correctly produce the TBLPRIKEY column data for SHAD_REJECT messages logged to the Log Server DOC. The release corrects this issue.
31. The HP Shadowbase Other Servers for Windows Loader Stored Procedure count hang the query session if a zero (0) wait time was specified when invoking the stored procedure. Using a value of wait time of one (1) or more is now the recommended correction.

For example:

```
EXECUTE SPBL_<owner>_<database>_<table> 'U' 100 2 0; -- May cause hang condition.
```

```
EXECUTE SPBL_<owner>_<database>_<table> 'U' 100 2 1; -- Should not hang.
```

Note: The wait time argument is the fourth and last argument to the Loader stored procedure.

Known Problems Remaining

- 1) The cached SQL statement TRS will halt replication under the following condition.

- The SSQLD000.dat/.idx files are deleted and recreated with SSTMNCxxx.dat/.idx data file(s) remaining from a previous replication session.

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 file becomes unusable or is accidentally removed, contact support for assistance and resolution to this issue.

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

2) Audit Log: SHAD_EVENT_TIMESTAMP column contents.

- 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 HP 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 corrected in an upcoming Shadowbase NonStop release, such that the SHAD_EVENT_TIMESTAMP column will contain the time the event was recorded in the HP 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 of the TRS/TFS instead of leaving it in a stopped state. This issue will be addressed in an upcoming release.

4) Restricted SBMON ROLL command usage.

- 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 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.
- This potential only exists when the system time is not drifted backwards, but rather when a hard reset system clock is done to set the clock back.. Normal system clock changes, which slowly correct the time by drifting the clock backwards, do not affect replication. To avoid these situations, drift your clocks to arrive at the correct time instead of forcing them to specific times in the past.

- The potential exists for data loss when the following conditions occur sub-second.

Oracle Source Collection:

- The system clock is set back to an earlier time; and
- Within *less than the clock adjustment period*: 1) an Oracle source collector restart happens due to a shutdown, or 2) the adjustment is done precisely before a DOC roll occurs; and
- The source collection DOC data is being forwarded to another DOC database via the Shadowbase Transaction Forwarding Server (TFS) and the TFS is caught up.

NonStop to DOC Writer:

- Negative NonStop system clock adjustment is performed; and
- A NonStop Shadowbase process restart in less than the clock adjustment period due to a processing or system fault; and
- The DOC database is being forwarded to another DOC database via the Shadowbase Transaction Forwarding Server (TFS).

The above conditions must all be true for the potential of data loss to exist. *No data loss will occur if this clock is drifted, only if it is hard reset backwards to a specific time or by a specific amount.*

To avoid this situation, use the following commands:

UNIX: ntpdate with “-B” option. Network Time Protocol with “-B” option always drifts the clock, rather than using settimeofday when the system time is >128ms out of sync. Also, set your crontab job to frequently call ntpdate, say once per hour, to reduce the odds of a large time correction.

WINDOWS: Set the Windows Time Service “MaxAllowedPhaseOffset” to a large value to guard against immediate clock resets. Use the following link for more specific information on the Windows Time Service.

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

- AND
- WHERE
- VALUES

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

- 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 HP 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.
- We are not supporting source collection for Oracle 12c at this time. There is a memory leak in Oracle 12c related to PGA memory that prevents replication from running reliably. This issue is under investigation; please contact your HP NonStop Sales and Support team or check the Shadowbase knowledge base to see if a solution has been posted if you need Oracle 12c source replication.
- Oracle log-based source collection does not support replication of tables with columns defined as `TIMESTAMP WITH TIMEZONE` or `TIMESTAMP WITH LOCAL TIMEZONE`.
- 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.
- SBORLOG currently only collects DML events (Insert, Update, Deletes).

8) 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
```

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

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

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

```
SHAD_REMOTE_MAX_EVENTS=200
```

11) On Windows, the DOC will not always roll at the correct size and can exceed the specified roll size significantly.

12) The statement generated for the SHADREJECT table may be incorrect when there is binary data embedded in a character field.

13) 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.

14) 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.

15) 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.

16) When configuring Oracle log-source 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.

17) SHAD_OCI_INCOMPLETE_SCHEMA and SHAD_SQL_ERROR_* parameters usage: SHAD_OCI_INCOMPLETE_SCHEMA specifies the action the Shadowbase process will take if it cannot read the table's definition (schema). The SHAD_SQL_ERROR_* parameters specifies the action the Shadowbase process will take if it encounters the error while applying data to the target database. They do not interact. 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. Continuing with the example, if SHAD_OCI_INCOMPLETE_SCHEMA is set to SHUTDOWN, the process will stop even though the error is specified to be skipped.

New User Error Messages (since version 6.100) in version 6.101:

```
BreakoutStatement():ParseInsert(INSERT); fault detected; Statement[<text>]
```

- Cause:** A parse error occurred against the Statement[<text>]
- Effect:** HP Shadowbase for Other Servers replication object experiencing the error will not process the statement 'as-is'.
- Recovery:** Correct the statement or skip the transaction containing the statement. Contact your HP Sales and Support team to obtain assistance.

```
*** OracleLogDLL::ProcessIUD(PriKeyUpdate): Stopping on UPDATE event for
Table <table name> UPDATE event modifies primary key Column data
```

UPDATE event info:

```
scn=<scn>; rbablk=<rbablk>; rbsqn=<rbsqn>; rbabyte=<rbabyte>; sql:
<sql statement>
```

You can set SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION to:

```
'STOP'      to stop on these events (current setting, default)
'SKIP'      to log and skip these events
'SKIP-NOLOG' to skip these events without logging
'CONVERT'   to convert these events into an insert and a delete
```

- Cause:** The collector read an update statement that modifies the primary key for the file and the collector was configured to stop on primary key updates. The message includes the table name (<table name>), information about where the event is located in the Oracle log files (<scn>, <rbablk>, <rbsqn>, <rbabyte>), and the statement from the log (<sql statement>).
- Effect:** The collector will stop.
- Recovery:** Set the *SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION* to SKIP or SKIP-NOLOG to skip any primary key update events; or

set the parameter to CONVERT to convert the event into an insert and delete.

OracleLogDLL::ProcessIUD(PriKeyUpdate): Skipping UPDATE event for Table <table name> UPDATE event modifies primary key Column data.

UPDATE event info:

scn=<scn>; rbablk=<rbablk>; rbaqn=<rbaqn>; rbabyte=<rbabyte>; sql:
<sql statement>

Cause: The collector read an update statement that modifies the primary key for the file and the collector was configured to skip and log the event. The message includes the table name (<table name>), information about where the event is located in the Oracle log files (<scn>, <rbablk>, <rbaqn>, <rbabyte>), and the statement from the log (<sql statement>).

Effect: The event is skipped.

Recovery: None required. If there are many events, you may want to change the *SHAD_OPCOLLECT_PRIKEY_UPDATE_OPTION* to SKIP-NOLOG to skip the events without logging. You can also set the parameter to CONVERT to convert the event into an insert and delete so that it can be processed.

Installation Instructions (optional "other" servers)

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

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