

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	Shadowbase
Release:	Gravic Version 6.600 HP NonStop Shadowbase: T1122-AAN (SB Repl/Guardian)
Release Date:	October 15, 2019
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2019 (www.gravic.com)
File Name:	IPM6600.pdf

NOTE: If this is a TCD delivery, please see [NOTE FOR TCDs](#) for TCD delivery information.

NOTE: This softdoc covers new features and corrected problems for Shadowbase for H, J, and L series Guardian NonStop operating system versions. This softdoc is available in an Adobe PDF file (.PDF). Softdoc files are named IPMnnnn (where nnnn is the Shadowbase version number).

NOTE: This softdoc applies to the HPE NonStop Shadowbase Replication for Guardian component. Other softdocs document the releases of the other components in T1122H06, including:

- HPE Shadowbase Audit Reader,
- HPE Shadowbase Enterprise Manager, and
- HPE Shadowbase Compare for SQL.

You are advised to reference those other softdocs for the changes related to those specific components of T1122H06.

NOTE: **Supported Release Version Updates (RVUS):** This softdoc applies to all currently supported database and operating system versions as described by the [Shadowbase Supported Versions Cross-Reference](#). In addition, visit the <http://www.ShadowbaseSoftware.com> web site for the list of supported platforms and databases, and the Shadowbase Software Support Policy for Software Versions by clicking here <https://shadowbasesoftware.com/support/shadowbase-software-product-release-and-support-policies/>.

Table of Contents

SOFTWARE RELEASE DOCUMENT (SOFTDOC).....	1
Table of Contents	1
Disclaimer	2
Note for TCDs.....	3
Mature and Deprecated Features	3
Supported Database and Operating System Versions.....	5
Special Notes for Version 6.600	6
Changes in Release 6.600 (since Version 6.500).....	10
New Features	10
Problems Fixed	13
New and Modified Commands	16
RESUMEDDL Command (New)	16
New and Modified SET AUD Parameters.....	17
New and Modified SET COLL Parameters	18
New and Modified SET CONS Parameters.....	18
New and Modified SET DBS Parameters.....	21
New or Modified EMS Messages	22
New and Modified User Information Messages	33
Installation Instructions (HPE NonStop Shadowbase)	35
Known Problems Remaining	37
[***** End of Document *****]	47

Disclaimer

We are distributing this communication in an effort to bring important information to the attention of users of the affected products. We recommend that all users determine the applicability of this information to their individual situations and take appropriate action. We do not represent or warrant that this information is necessarily accurate or complete for all user situations and, consequently, we will not be responsible for any damages resulting from the user's use or disregard of the information provided. To the extent permitted by law, we disclaim all representations and warranties, whether express, implied, statutory, or otherwise, including the warranties of the merchantability, fitness for a particular purpose, title, and non-infringement.

We expect customers of the Shadowbase product suite to “stay current” on Shadowbase releases. This means that you, the customer, should periodically upgrade your Shadowbase software to a newer release that is under support before support ends on your current release. For most customers, this means that you will want to upgrade while your release is in ‘ACTIVE’ support. Otherwise, you run the risk of not being able to get full (or even any if the release has gone ‘OBSOLETE’ end-of-service-life) support for the version you are running.

The Shadowbase Software Policy for Software Versions is described here: <https://shadowbasesoftware.com/support/shadowbase-software-product-release-and-support-policies/>. We encourage all customers to periodically review this material and plan for periodic upgrades to their Shadowbase software. Contact Support if you need additional information.

Note for TCDs

TCD (Temporary Code Delivery) – A software update delivered via an SPR downloadable from an FTP dropbox. A TCD is an early version, intended for customer testing only (not production usage). A TCD by definition is restricted to certain customers. Note that a “Gravic TCD” is delivered directly from Gravic, not via HPE, but otherwise has the same attributes.

A TCD is provided only to the specified customer for the purposes agreed between the customer and Gravic as to how it will be used. A TCD is provided subject to the following terms and conditions in addition to the existing written license governing the use of Shadowbase:

- A TCD is provided for evaluation and test purposes only for no more than ninety (90) days use, and is not to be used in production systems
- A TCD may not have been fully tested by Gravic, no warranties are implied as to its behavior
- A TCD is delivered directly from Gravic to the customer, it is not available from HPE/SCOUT
- As testing proceeds, iterative TCD deliveries may be necessary as issues are identified/resolved
- A TCD is temporary, after evaluation it is to be withdrawn from use by the customer
- After testing completes, a TCD may or may not subsequently be released as a Shadowbase TCF or otherwise be included in the Shadowbase product line

Please see <https://www.shadowbasesoftware.com/support/shadowbase-software-product-release-and-support-policies/shadowbase-software-release-glossary/> for additional information.

Mature and Deprecated Features

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

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

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

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

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

Table 1 - Mature Features in HPE NonStop Shadowbase

Feature	Description	Mature Release
AUDCOLLN	The original native mode collector, AUDCOLLN is now mature and is being replaced by the new SBCOLL collector. AUDCOLLN will continued to be maintained for critical issues and tested, but will no longer be enhanced with new features and capabilities.	NSB 6.410
CONSTXBQ	A user exit-bound consumer that will process queue file events by use of a mapping file and batches transactions. Consider replicating queue files using the ALLOWQUEUEFILE REPLICATEKEY option along with the CONSTXB user exit-bound consumer instead if transaction batching is needed.	NSB 6.600
Format 1 TMF Audit Trails	While still supported, Gravic recommends that any customers still using Format 1 Audit Trails migrate to Format 2 TMF Audit Trails.	NSB 6.410
ENS2SQL	Tool used to convert Enscribe DDL records to SQL/MP tables, effectively replaced by SBDDLUTL	NSB 6.600

Table 2 - Deprecated Features in HPE NonStop Shadowbase

Feature	Description	Deprecated Release	Last Release Date
Non-native Consumer	The non-native version of the consumer (AUDCONS) is deprecated. Customers still using the non-native version of the consumer should immediately begin planning to move to the native version (AUDCONSN).	NSB 6.410 April, 2018	April, 2020
Loader	The original Shadowbase Loader (LOADER) is deprecated. Customers still using LOADER	NSB 6.300 May, 2017	May, 2019

	to load files and tables should immediately begin planning to move to the Shadowbase Online Loader (SOLV).		
SBLINK	The original program used to connect SEM to a NonStop environment, replaced by listener AUDCOM process.	NSB 6.220 May, 2016	NSB 6.320 May, 2017
SBQFILE	Program used in conjunction with replicating Enscribe queue files to target non-queue key-sequenced files, replaced with exact queue file replication.	NSB 6.220 May, 2016	NSB 6.320 May, 2017
SQL92	Protocol used by Shadowbase components between the NonStop and Other Servers, replaced by SQLSBV1.	NSB 6.320 May, 2017	NSB 6.410 April, 2018

Supported Database and Operating System Versions

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

Special Notes for Version 6.600

- 1) As of version 6.320, the behavior of the Queue Manager ABNORMALSHUTDOWN parameter and the default value has been changed. As a result, when upgrading from a pre-6.320 version of Shadowbase in an environment with Queue Manager objects configured, the associated consumer processes will need to have a TIDFILE configured (if not already), or the ABNORMALSHUTDOWN parameters will have to be set to STOP for the Queue Manager objects.
- 2) If you utilize the SOLV loader related to a DBS object that has ALLPARTITIONS set ON, you should set the SOLV TARGETFNAME parameter to the DBS SOURCEFILE value. TARGETFNAME consists of the node name, volume and subvolume (no filename) of the SOURCEFILE.
- 3) As of version 4.091, the CONTROLFILE record size, associated with the CONS CONTROLFILE parameter, has been increased. You must purge any existing CONTROLFILE's before restarting Shadowbase if you are upgrading from a pre-4.091 version of Shadowbase.
- 4) Enscribe and SQL/MP CREATE, PURGE, PURGEDATA, and ALTER events are now processed directly by the Consumer. If you have the DBS setting for these events enabled (ON) and have user exit code that processes these events, you must change the code to call the SBSETIGNORE API function before returning. Otherwise, the Consumer will attempt to process the events again and this could result in errors.
- 5) For those that have user exits, you must use the 6.600 USRXLIBO (non-native version), or USRXLIBN or USRXLIBX (native versions) API library included in the distribution subvolume. If you have copies of the USRXLIBO, USRXLIBN, or USRXLIBX file in other subvolumes, make sure you replace them with the current version from the distribution subvolume. You should not use earlier versions of USRXLIBO, USRXLIBN, or USRXLIBX with Shadowbase version 6.600.
- 6) Control file changes made in Version 5.001 require that you perform an initial START AUD COLD if you are upgrading from a release prior to 5.001.
- 7) For bi-directional replication environments, due to enhancements made to the internal algorithms within the COLL and CONS, it is recommended that both (or all) nodes be upgraded at the same time. Shut down both, do your upgrade, and then bring each node's Shadowbase environment back up. Contact Shadowbase Support if you have bi-directional environments and you cannot upgrade them both at the same time.
- 8) For SQL/MX IEEE 32-bit floats, 8 digits of precision in scientific notation is the SQL/MX Consumer (new AUDCONXN and AUDCONXR programs delivered in 5.000) assigned value in the generated SQL statement data. The Tandem format 32-bit float will continue to use 6 digits of precision in the non-IEEE Consumer (AUDCONS, AUDCONSN, and AUDCONSR programs).

Additionally, for DDL related to Enscribe files, the Consumer assumes that all floats are in Tandem format and not IEEE format. A user exit is required if they are actually in IEEE format. Otherwise, a generated SQL statement could have an incorrect number.

9) If you are developing user exit code to use with the SQL/MX Consumer, make certain that you compile your code for IEEE floating point format. Only native code can be used with the SQL/MX Consumer. Non-native user exits are not supported.

10) You must install Shadowbase on a physical volume if SQL/MX tables will be replicated. You should not install Shadowbase on an SMF virtual drive, otherwise errors will result during the install process.

11) If SQL/MX tables will be replicated with Shadowbase, NonStop Shadowbase must be installed on the source system(s) and Shadowbase Open Server must be installed on the target system(s) as Shadowbase's SQL/MX target support is implemented in the OSS environment. Shadowbase Open Server for OSS, version 5.000 or higher, is required.

12) Support for SQL/MX has the following limitations:

- Tables with Relational Integrity (RI) constraints. Ordering of I/Os within a transaction for multiple tables cannot be guaranteed in the audit trail due to parallel nature of transaction processing on NonStop systems. For example, consider the case where there are parent and child tables, with a constraint to force a corresponding parent row to exist when a child row is added. If you add both the parent and child row under the same transaction when a child is added with a new parent, and the parent and child tables are on separate data disks, TMF may store them in reverse order in the audit trail and Shadowbase may not be able to add those rows on the target. Depending on which disk flushed to the audit trail first, the audit trail may return the rows with the child row insert first and the parent row insert second. In this case, Shadowbase will fail to insert either row due to the RI constraint being violated on replay.
- Tables with large rows and Cascading Deletes configured. We have seen that TMF may save updates of large rows as a delete event of the old row followed by an insert of the revised row. Shadowbase will replay this event as a delete followed by an insert as well. If cascading deletes are configured on the target, the delete may cause all child rows to be deleted on the target. Note that this behavior may also run afoul of RI constraints noted above, if a child row is inserted between the delete and insert of the parent row.

13) If you are upgrading from a release prior to Version 6.100, note that the Shadowbase EMS subsystem owner and version (the EMS SSID token) changed in Version 6.100. The subsystem owner changed from ITI to GRAVIC, and the subsystem version is now

composed of VNN, where NN are the two most significant digits of the Shadowbase release version id. As a result, the subsystem version has changed from D42. For any upgrade to Version 6.6nn, the version has thus been updated according to the VNN format to V66. Any EMS filters using these attributes for filtering must likewise be updated. Alternatively, the SBEMSOLDSSID TACL parameter can be enabled to have Shadowbase use the old values for the subsystem owner and version.

14) The EMS subsystem version has been updated with this version from V65 to V66 to match the two most significant digits of the current Shadowbase release version id.

15) Errors 8401, 8405, 8406, 8410, 8411, 8412, 8417, 8418, 8423, 8424, 8425, 8426, 8427, and 8428 will no longer be logged and ignored by the Consumer by default. The user exit stubs included with Shadowbase, USRXC, USRXC0B, and USRXC0BA, have been modified to remove the call to SBSETEXCEPTIONCODES. The code to setup the structure and make the call remains as a comment in the stub files. To use the structure and make the call to SBSETEXCEPTIONCODES, uncomment the code from the stub file being used.

16) Some customers experienced a rare issue, which caused the ARLIB2 library to hang when there was no activity on one of the auxiliary audit trails. The Collector now links with the ALV version of the ARLIB2 library, which incorporates a fix for this issue.

17) When using a multi-ported DOC Writer, it is recommended as a best practice not to use Queue Managers for the related threads in the Shadowbase NonStop configuration. Furthermore, it is recommended not to use the SUSPENDUPD/RESUMEUPD commands when replicating to a multi-ported DOC Writer. Contact Support if this issue applies to you and you need to discuss options.

18) As of Version 6.400, the SBLINK and SBQFILE programs have been deprecated will no longer be shipped alongside the rest of Shadowbase.

19) SBHADW configuration scripts now sets the TACL parameter SBCOLLAUTOADVDMODE in the AMON script to a default value of 0. This parameter used to default to 1.

20) If a Shadowbase process on another node does not respond to a STOP request from the AUDMON, the AUDMON will attempt to kill the process using the PROCESS_STOP_ system call. This attempt will often fail if the remote process is PROGID'd, as all of the Shadowbase processes are. Once the AUDMON has passed its STOPPROCESSMAXRETRIES threshold, it will assume the process has been closed and continue on as usual.

When the environment is started up again, the remote process will fail to start because it still exists from before. This problem can be prevented by creating a SAFEGUARD Access Control List (ACL) for the process to allow access to the user that AUDMON is PROGID'd under. Once the user has been whitelisted with the ACL, AUDMON is able to

properly kill the remote process.

Changes in Release 6.600 (since Version 6.500)

New Features

- 1) SOLV can now run in a new ‘bypass’ mode. This mode can be used to eliminate the fact that reading from a source file or table is the cause of a data-loading bottleneck. With this mode enabled, SOLV will read or select data from the source without sending the data events to the HPE Shadowbase consumer. Normally, running SOLV in this mode would be at the request of HPE support.
To enable SOLV bypass mode, remove the CONSUMERFNAME parameter from your configuration before starting SOLV. In addition, the SOLVMKRPREFIX parameter must be supplied. This prefix is necessary because the peer consumer is the supplier of this value to SOLV.
Three SOLVMKR files are created when a consumer is started. These files maintain SOLV’s restart positioning and a history of past SOLV loads. The three files end with ‘A’ (audited), ‘U’ (unaudited) and ‘H’ (history). The SOLVMKRPREFIX is the full name of these files without the trailing A, U or H character.
Use the following procedure to create the marker files if SOLV needs to be started in bypass mode before actually starting the HPE Shadowbase consumer.

```
3> fup
File Utility Program - T6553H02 - (01MAY2018) System \GRAVIC1
(C)Copyright 1984-2016 Hewlett Packard Enterprise Development LP
-set type k
-set format 1
-set code 7685
-set ext(16,64)
-set rec 4062
-set block 4096
-set keylen 16
-set keyoff 8
-set audit
-set maxextents 160
-set buffered
-create markera
CREATED - $DATA.TFS1800.SOLVMKRA
```

```
4> fup
File Utility Program - T6553H02 - (01MAY2018) System \GRAVIC1
(C)Copyright 1984-2016 Hewlett Packard Enterprise Development LP
-set type k
-set format 1
-set code 7686
-set ext(16,64)
-set rec 4062
-set block 4096
```

```
-set keylen 16
-set keyoff 8
-set maxextents 160
-set buffered
-create markeru
CREATED - $DATA.TFS1800.SOLVMKRU

5> fup
File Utility Program - T6553H02 - (01MAY2018) System \GRAVIC1
(C)Copyright 1984-2016 Hewlett Packard Enterprise Development LP
-set type k
-set format 1
-set code 7687
-set ext (64,128)
-set rec 4062
-set block 4096
-set keylen 24
-set keyoff 0
-set maxextents 160
-set buffered
-create markerh
CREATED - $DATA.TFS1800.SOLVMKRH
```

- 2) SOLVMX can now generate explain plans for the queries it executes to load SQL/MX tables.

Add the following parameter to your environment before executing SOLVMX to enable explain plan generation.

```
PARAM EXPLAINPLAN ON
```

Instead of loading the SQL/MX table, SOLVMX will instead generate two files. One is the input file to MXCI and the other is the resulting explain plan output from MXCI.

By default, both files can be found in the same sub-volume where the SOLVMX program resides. Both files are prefixed with the process name used to start SOLVMX. The input file name will have a SXI (Solv mX Input) suffix. The output file will have a SXO (Solv mX Output) suffix.

The file location of these two files can be changed using the EXPLAINPLAININ and EXPLAINPLANOUT parameters.

For example.

```
PARAM EXPLAINPLAN ON
PARAM EXPLAINPLANIN $data.subvol.planin
```

PARAM EXPLAINPLANOUT \$data.subvol.planout

EXPLAINPLANOUT can also be set to a specific terminal name or [#MYTERM].

For example...

PARAM EXPLAINPLANOUT [#myterm]

- 3) The Consumer will now check the DBS ROOTPARTITION parameter for secondary partition events for Enscribe relative and entry-sequenced files when it is loading the source schema for the first time. If ROOTPARTITION is not defined for the DBS, behavior will depend on the new CONS MISSINGROOTPART parameter. The CHECK DBS feature in AUDCOM will also detect this situation before replication begins.
- 4) SBCOLL now supports DDL Audit Replication for uni-directional environments. Previously, this was only supported by AUDCOLLN. The same limitations on the types of events support apply for SBCOLL as with AUDCOLLN, including that only Enscribe and SQL/MP DDL events are supported. DDL Audit Replication is configured using the existing DBS parameters (e.g. ALTERS, ALTEROPTIONS, CREATES, CREATEOPTIONS, PURGEDATAS, PURGEDATAOPTIONS, PURGES, PURGEOPTIONS). SBCOLL does not support DDL Audit Replication for bi-directional replication.
- 5) SBDDLUTL now exits with completion codes; allowing it to be integrated into TACL scripts, etc. It supports the following completion codes:
 - 0: success (no issues during program execution)
 - 2: some or all output was not generated (i.e. command syntax errors, bad options, bad command values)
 - 3: exit fatal (SBDDLUTL encountered an error it could not recover from and exited prematurely)
 - 4: exit no start (SBDDLUTL before fully starting; e.g. SHADPASS license issues, error during initialization)
 - 8: exit with warnings (all output produced, conflict found like duplicate field names that were unresolved, incompatible data types, etc.)
- 6) SBDDLUTL now gives users the ability to decide what to do when the flattening process flattens away the primary key group of a key-sequenced DDL record. A new program setting was created called ONKEYGROUPREMOVAL. Its possible values are:
 - KEEPGROUP: flatten the DDL record, but keep the primary key group in-tact and un-flattened, allow the customer to decide what to do with it manually.
 - NOFILE: comment out the FILE IS clause at the top of the record. This effectively turns the DDL record into an “unstructured” record so the DDL compiler does not complain about a missing key.

- 7) SBDDLUTL now displays command syntax for HELP SET <program option> help text.
- 8) SBDDLUTL now supports the CONCAT command for fields with OCCURS, in addition to groups. A field with an OCCURS clause > 1 will be converted into a field of type PIC X(n) where n is the size of the original field's data type multiplied by the OCCURS clause number.
- 9) SBDDLUTL now allows the user to completely wildcard the REDEFINE command target, i.e. REDEFINE <field or group name> AS *. This produces output where all redefines, including the original field or group definition are preserved in the DDL or SQL output.
- 10) Shadowbase 6.6 includes a new DDL replication feature called Command Stream Replication for SQL/MP (MP CSR). This feature includes updates and enhancements made to the Collector (AUDCOLLN & SBCOLL), Consumer, Queue Manager, AUDMON, and AUDCOM. The existing CONTROLFILE facility was enhanced to include two possible sizes (LEGACY / ENHANCED). A new front-end utility was introduced called SBSQLCI which processes user commands and initiates MP CSR processing. New Shadowbase parameters and User/EMS messages were added. See the HPE Shadowbase Command Stream Replication Manual for detailed information about this new and exciting feature.
- 11) AUDCOM's STATUS AUD, STATUS CONS, STATS CONS, and STATUS COLL output has been updated to reflect new information related to MP CSR. See the HPE Shadowbase Command Stream Replication Manual for detailed information.
- 12) For DDL replication, coordination is now supported when using a QMGR object. Be aware, however, that when coordinating with a QMGR, replication will be held up until the Consumer has processed all of the queued data, as the Consumer needs to catch up to the current audit trail position in replication for coordination.
- 13) The new HPE Shadowbase™ for NonStop SOLVUTIL product has been added to this release. It is used to load target OSS regular files in order to match their corresponding source entities. The *IPM 6.600 SOLVUTIL Addendum* describes the basic functionality of HPE NonStop Shadowbase SOLVUTIL, how to configure for LOADING, and also how to monitor and manage these environments.

Problems Fixed

- 1) The COLL REPLICATION parameter is now properly applied to AUDCOLL when the environment is started with an existing restart file.
- 2) The Consumer will no longer log warnings about skipping events for the configured CONTROLFILE.

- 3) Fixed formatting for Consumer message resulting from trying to build a SQL statement that exceeds the maximum size.
- 4) The AUDCOM HELP SET COLL command would overflow the text output buffer, resulting in a trap when running subsequent commands.
- 5) AUDCOLLN was not properly releasing data buffers when running in REPLICATION SKIP mode. This would eventually result in a situation where the environment would hang once the maximum buffers were in use.
- 6) ALTERing certain DBS parameters after capturing that DBS could result in a trap in AUDCOM.
- 7) Consumer was not properly resetting the AUDMON check timer if a QMGR drain was active. This could lead to situations where the Consumer would hang indefinitely.
- 8) The CONS will now properly roll ILF files when MAPFROMSOURCE is set for a wild-carded DBS SOURCEFILE.
- 9) SBDDLUTL no longer outputs a deprecated “OUT” command to the session SAVE file.
- 10) SBDDLUTL would trap doing a NETFORMAT FLAT command due to missing required fields in the output data.
- 11) SBDDLUTL’s parser couldn’t handle certain situations where there was no whitespace between tokens (for example, “CONVERT TARGET SQLSERV/OUT myfile/;” would cause an error.
- 12) SBDDLUTL had some issues with formatting SET VALUE clauses for DDL output. Regular double-quoted string values would produce extra double quotes in the output, and left-justified or right-justified string values would not be processed by the parse correctly.
- 13) The non-native consumer would trap with an illegal address reference when processing a call to SBGETREALSOURCE called within a USRXSTART function.
- 14) SBDDLUTL displayed out-of-date HELP SET syntax in its help text.
- 15) SBDDLUTL would not properly clear its command cache upon encountering erroneous CONVERT and SHOW commands. This caused it to try and re-execute prior erroneous commands again with subsequent calls to CONVERT and SHOW.
- 16) SBDDLUTL had issues numbering duplicate field names properly.
- 17) Conversion of fields created from the CONCAT command would always have a zero-length in target SQL outputs.

- 18) SBDDLUTL did not properly expand input file names with user defaults.
- 19) The RENAME RECORD command would cause SBDDLUTL to be unable to apply previously set program settings to the output DDL/Create statements.
- 20) SBDDLUTL did not process field-level OCCURS clauses properly when the OCCURS was set to > 1.
- 21) Very long commands would be truncated in SBDDLUTL's FC feature.
- 22) SBDDLUTL would not load primary key information properly for Enscribe DDL records defined with a DEFINITION IS clause.
- 23) SBDDLUTL did not generate correct data type output for SQL/MP CREATE TABLE statements for some Enscribe DDL data type conversions.
- 24) SBDDLUTL did not generate correct data type output for SQL/MP CREATE TABLE statements for some Enscribe DDL binary data type conversions.
- 25) SBDDLUTL dropped a byte of output for TRAILING SIGN EMBEDDED DDL data types when converting to SQL/MP.
- 26) SBDDLUTL converted Enscribe DDL type BIT to an incorrect data type for SQL/MP.
- 27) SBDDLUTL generated incorrect data types for TRAILING SIGN SEPARATE and LEADING SIGN SEPARATE Enscribe DDL data types.
- 28) SBDDLUTL was not properly scanning for the primary key of a record if the primary key was defined after the alternate key(s).
- 29) Misc. typos and factual mistakes in SBDDLUTL documentation and help text were fixed / updated.
- 30) Specifying the target DDL dictionary or record name in a SET command caused SBDDLUTL to be unable to find the proper field or group the program settings are to apply to on a CONVERT command.
- 31) Both Collector objects will now write the configured REPLICATION parameter setting to the RESTART file if it is determined that a valid setting is not already present (possible when upgrading from an older version of HPE Shadowbase). This will avoid repetitive messages when restarting HPE Shadowbase without changing the REPLICATION parameter from its configured value. This also removes the need to run RESTEDIT when upgrading from pre-6.500 environments with existing restart files.

- 32) When using SBCOLL in a bi-directional environment and starting replication from EOF, the restart file would not be written until the pre-reading was completed and the initial audit trail EOF was reached.
- 33) Replicating with SBCOLL starting from a specific point in the audit trail using the ADTSTARTSEQLO and ADTSTARTRBA parameters did not function properly. Events would not be replicated, and the restart file would not be generated.
- 34) AUDCOM would trap if a HISTORY command was issued following a HELP SET COLL command.
- 35) Fixed an issue with bi-directional replication with SBCOLL, where the bi-directional restart position could be set to a position after the replication restart position, when the first TRANSLOG event encountered occurred after the restart position. This was benign, but resulted in bi-directional replication points being shown beyond the restart position, while it was expected to be the same as the restart position or earlier.

New and Modified Commands

RESUMEDDL Command (New)

Support has been added to AUDCOM / AUDMON for the RESUMEDDL command. This command is part of a brand new feature in NSB 6.6 called Command Stream Replication for SQL/MP (MP CSR).

When an MPSERVER Consumer processes an MP CSR command that results in a SQL error its default behavior (configurable via the MPERRORACTION parameter) is to go into a PAUSED state. During this state, Shadowbase replication halts while the MPSERVER consumer waits for a RESUMEDDL command to be issued via AUDCOM. The PAUSED states gives the user the opportunity to manually resolve issues before allowing Shadowbase to continue replication.

If the SQL error caused SQL to abort the current transaction, the MPSERVER consumer will instead go into the PAUSEABT state and wait for a RESUMEDDL STOP or RESUMEDDL SKIPABORTED command.

RESUMEDDL tells the MPSERVER consumer how to handle the encountered error. The syntax follows:

```
RESUMEDDL { [CONS] { [<audmon-name>.<consname> { RETRY }
                                                    { SKIP }
                                                    { SKIPABORTED }
                                                    { STOP } }
            { [<audmon-name>.* { RETRY }
                                { SKIP }
                                { SKIPABORTED }
                                { STOP } } }
```

Where:

- **STOP:** instructs the consumer to abend.
- **SKIP:** instructs the consumer to simply skip the erroneous MP CSR event and continue running
- **SKIPABORTED:** instructs the consumer to skip the erroneous MP CSR event, continue running, and start a new transaction for subsequent MP CSR events.
- **RETRY:** instructs the consumer to try processing the MP CSR event again.

Note: RESUMEDDL SKIPABORTED is only available if the MPSEVER consumer is in the PAUSEABT state.

For more information on RESUMEDDL and MP CSR error handling, see the HPE NonStop Shadowbase Command Stream Replication Manual.

New and Modified SET AUD Parameters

CONTROLFILERECSIZE { ENHANCED }
 { LEGACY }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	LEGACY	n/a	n/a	No

Configures the size of the control file AUDMON creates on start-up (the location of which is determined by the CONTROLFILE CONS param). This parameter is optional, and if omitted, is set to LEGACY, where:

- **LEGACY:** a regular key-sequenced control file is created of record size 4,040 bytes
- **ENHANCED:** an enhanced key-sequenced file with increased limits is created of record size 27,648 bytes

Note that if a control file already exists, Shadowbase will use the existing control file. Look at the HPE Shadowbase Command Stream Replication Manual for more information on upgrade and fallback procedures. Also note that support for Enhanced Key-Sequenced files with increased limits has been added to HPE NonStop H06.28 and J06.17. For more information, see the File Utility Program (FUP) Reference Manual.

New and Modified SET COLL Parameters

COORDINATIONMSG { NONE }
{ VERBOSE }
{ WAITWARN }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	WAITWARN	n/a	n/a	No

Determines the level of verbosity of the collector's coordination messages. This parameter is optional, and if omitted, defaults to WAITWARN. The values are:

- **NONE:** do not issue any EMS messages when coordinating
- **VERBOSE:** issue an EMS message when coordination begins, ends, and at an interval while coordination is taking place (configurable by the COORDWAITINGWARNRATE parameter)
- **WAITWARN:** issue an EMS message on the configured COORDWAITINGWARNRATE interval, only if coordination is still occurring.

Note that these coordination EMS messages apply to coordination triggered by MP CSR as well as DDL Audit Replication.

COORDWAITINGWARNRATE { TIME in SECONDS }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	300	1	32,767	No

Configures an interval at which the collector displays coordination EMS messages. Note that this message will only display if the COORDINATIONMSG COLL param is set to VERBOSE or WAITWARN. This parameter is optional and defaults to 300 seconds, or 5 minutes.

New and Modified SET CONS Parameters

MISSINGROOTPART { NOWARN }
{ WARN }
{ SKIP }
{ STOP }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	STOP	n/a	n/a	No

This parameter controls the behavior of the Consumer when it encounters an event for the secondary partition of an Enscribe relative or entry-sequenced file without having the DBS ROOTPARTITION defined.

NOWARN: The event is replicated without logging a message. This setting is not recommended due to the possibility of data corruption.

WARN: The event is replicated with a message. This setting is not recommended due to the possibility of data corruption.

SKIP: Events for this file are skipped.

STOP: The Consumer will abend. This is the default and recommended setting.

MPSERVER { ON }
 { OFF }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	OFF	n/a	n/a	No

Designates the given Consumer as an MPSEVER consumer. MPSEVER consumers are capable of processing MP CSR events. At least one consumer within a Shadowbase configuration must be configured with this parameter for MP CSR to be enabled and work properly. It is recommended to configure a dedicated MPSEVER consumer, and let other consumer(s) in the environment perform regular audit data replication.

MPLOG { filename }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	None	n/a	n/a	No

Configures the location of the MPSEVER log file. This parameter has no default value and if not set, the MPSEVER consumer will not log CSR-related informational messages.

MPPRI { number }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	None	1	199	No

Configures the priority of the underlying SQLCI process spawned by the MPSEVER consumer during certain MP CSR events. This parameter is optional, and if omitted, the priority of the MPSEVER consumer is used.

MPPROCESS { processname }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	None	n/a	n/a	No

Configures the process name of the underlying SQLCI process spawned by the MPSEVER consumer during certain MP CSR events. This parameter is optional, and if omitted, a system generated process name is used. Note that if supplied, the node name portion of the process name is required.

MPCOORDINATION { ON }
{ OFF }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	ON	n/a	n/a	No

Determines whether MP CSR events will trigger Shadowbase coordination. This parameter is optional, and defaults to ON. Coordination occurs when an MP CSR event is read from the audit trail, causing the collector to stop reading audit and wait for downstream consumers to drain (close target files, free SQL resources, etc.) Once drained, the MPSEVER consumer processes the MP CSR event that triggered the coordination, and upon completion, the collector resumes reading audit and sending new audit data to its downstream consumers. In most cases it is recommended to leave MPCOORDINATION ON if processing MP CSR events.

MPERRORACTION { PAUSE }
{ SKIP }
{ SKIPNOLOG }
{ STOP }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	PAUSE	n/a	n/a	No

Determines how the MPSEVER consumer responds to errors encountered while processing MP CSR events. This parameter is optional and defaults to PAUSE if omitted. The parameter values are:

- **PAUSE:** wait until the user issues a RESUMEDDL command via AUDCOM
- **SKIP:** log a message to the log file (denoted by the MPLOG parameter,) skip the erroneous MP CSR event, and continue running
- **SKIPNOLOG:** skip the erroneous MP CSR event and continue running
- **STOP:** log a message and abend

Note that while the MPSEVER consumer is in the PAUSED state, the upstream collector is not reading audit, as it is waiting for the consumer to finish processing the CSR event. Periodic EMS messages will be logged as the MPSEVER consumer continues to wait for a RESUMEDDL command.

MPPROGRAM { filename }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.600	n/a	\$SYSTEM.SYSTEM.SQLCI	n/a	n/a	No

Configures the location of the SQLCI program file that MPSEVER consumer will use to spawn a SQLCI session when processing some MP CSE events. This parameter is optional, and if omitted, the MPSEVER uses \$SYSTEM.SYSTEM.SQLCI.

New and Modified SET DBS Parameters

DTSOURCEFORMAT { ON }
 { OFF }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
5.002	6.600	ON	n/a	n/a	No

The default of this parameter was changed due to the deprecation of SQL92 in favor of SQLSBV1.

New or Modified EMS Messages

2116	ERR 2180: COULD NOT GET ALL REQUESTED INFO FROM CALL TO FILE_GETINFOLISTBYNAME_ON CONTROLFILE; CODE: <#>
------	--

Cause: The user configured Shadowbase and set the CONTROLFILE CONS_PARAM. Shadowbase attempted to validate the given control file but ran into an issue.

Effect: Shadowbase displays this EMS message and abends.

Recovery: Look at the EMS message; an associated Guardian error code is printed with it. If that doesn't help, contact Shadowbase support.

2125	ERR 2181: EXISTING LEGACY CONTROL FILE DETECTED BUT CONTROLFILERECSIZE SET TO ENHANCED. SEE THE HPE SHADOWBASE COMMAND STREAM REPLICATION MANUAL FOR UPGRADE PROCEDURES.
------	--

Cause: The user configured Shadowbase and set the CONTROLFILE CONS_PARAM. The user also set the CONTROLFILERECSIZE parameter to ENHANCED. The control file itself already existed at the time of Shadowbase start-up, but it was a LEGACY control file.

Effect: Shadowbase displays this EMS message and abends.

Recovery: Either upgrade the existing control file by following the procedure outlined in the CSR manual under the title "Controlfile Upgrade and Fallback," or set the CONTROLFILERECSIZE parameter to LEGACY. Then, restart Shadowbase.

2116	MORE THAN ONE MPSEVER CONSUMER CANNOT BE CONFIGURED FOR THE SAME UPSTREAM COLLECTOR; THIS ISSUE DETECTED FOR AUDMON <audmon>, COLLECTOR <coll>.
------	---

Cause: The user configured Shadowbase such that more than one consumer downstream to the same collector is designated as an MPSEVER.

Effect: Shadowbase displays this EMS message and abends.

Recovery: Fix the Shadowbase configuration by making sure no more than one MPSEVER consumer is configured per collector.

2116	MP CSR CANNOT BE CONFIGURED FOR OTHER SERVERS TARGET. TARGETTYPE MUST BE 'NSDIRECT'
------	---

Cause: The user configured Shadowbase to replicate out to OPEN/OTHER and set the MPSERVER CONS_PARAM to ON for one or more of the configured consumers.

Effect: Shadowbase displays this EMS message and abends.

Recovery: Remove the MPSERVER parameter and re-start Shadowbase.

2116	MPPROCESS IS REQUIRED WHEN MPSERVER IS SET TO 'ON'
------	--

Cause: The user configured Shadowbase and set the MPSERVER parameter to ON for one or more consumers, but omitted the MPPROCESS parameter.

Effect: Shadowbase displays this EMS message and abends.

Recovery: Either disable MP CSR by commenting out, removing, or setting the MPSERVER parameter to OFF; or configure the MPPROCESS parameter for the given MPSERVER consumer(s); then re-start Shadowbase.

2008	DBS <db name> INCLUDES SECONDARY PARTITION FILE <file> WITHOUT ROOTPARTITION DEFINED. <consumer action>. CONSIDER SETTING ROOTPARTITION FOR THIS DBS OR CHANGING CONSUMER MISSINGROOTPART SETTING.
------	--

Cause: An event was processed for a secondary partition of an Enscribe relative or entry-sequenced file <file>, but ROOTPARTITION is not specified for the associated DBS <db name>.

Effect: The result will depend on the setting of the CONS MISSINGROOTPART parameter, and will be displayed in the above message's <consumer action>.

Recovery: Set the ROOTPARTITION parameter for the given DBS to the primary partition volume of the file in question.

2125	LEGACY CONTROLFILE CREATED. SBSQLCI COMMANDS WILL BE LIMITED TO <#> BYTES. CONSIDER SETTING CONTROLFILERECSIZE TO ENHANCED.
------	---

Cause: The user configured Shadowbase and set the CONTROLFILE CONS_PARAM. The user also either omitted the CONTROLFILERECSIZE parameter, or explicitly set it to LEGACY. The control file itself did not exist at the time of Shadowbase start-up.

Effect: Shadowbase created the specified control file as a LEGACY control file. This means the file is a standard key-sequenced file with a record size of 4040 bytes.

Recovery: N/A

2125	EXISTING LEGACY CONTROLFILE DETECTED. SBSQLCI COMMANDS WILL BE LIMITED TO <#> BYTES. CONSIDER SETTING CONTROLFILERECSIZE TO ENHANCED.
------	---

Cause: The user configured Shadowbase and set the CONTROLFILE CONS_PARAM. The user also either omitted the CONTROLFILERECSIZE parameter, or explicitly set it to LEGACY. The control file itself already existed at the time of Shadowbase start-up.

Effect: Shadowbase uses the found control file.

Recovery: N/A

2108	BEGINNING CONS COORDINATION FOR <truncated text> COMMAND. EVENTS WILL NOT BE SENT FROM COLL UNTIL CONFIRMATION OF COORDINATION COMPLETION IS RECEIVED.
------	--

Cause: The user has entered an SBSQLCI command, initiating Command Stream Replication for SQL/MP while the MPCOORDINATION parameter is turned ON.

Effect: The collector displays this EMS message containing the truncated DDL command text and waits for coordination completion confirmation from the consumer.

Recovery: No recovery required. To disable the message, set the COORDINATIONMSG parameter to NONE.

2108	BEGINNING CONS COORDINATION FOR COMMAND: '<full text>'. EVENTS WILL NOT BE SENT FROM COLL UNTIL CONFIRMATION OF COORDINATION COMPLETION IS RECEIVED.
------	--

Cause: The user has entered an SBSQLCI command, initiating Command Stream Replication for SQL/MP while the MPCOORDINATION parameter is turned ON.

Effect: The collector displays this EMS message containing the full DDL command text and waits for coordination completion confirmation from the consumer.

Recovery: No recovery required. To disable the message, set the COORDINATIONMSG parameter to NONE.

2108	COLL WAITING ON COORDINATION FOR THE FOLLOWING CONSUMER(S): [<consumer list>]. EVENTS WILL NOT BE SENT FROM COLL UNTIL COORDINATION IS COMPLETED BY EACH LISTED CONSUMER.
------	---

Cause: HPE Shadowbase MP CSR has been in coordination for as long as the value of the COORDWAITWARNRATE parameter. The COORDINATIONMSG parameter is set to WAITWARN or VERBOSE.

Effect: The collector displays the given EMS message.

Recovery: No recovery required. To disable the message, set the COORDINATIONMSG parameter to NONE.

2108	COORDINATION COMPLETED FOR ALL CONFIGURED CONSUMERS. COLL WILL NOW RESUME READING AUDIT AND SENDING EVENTS.
------	---

Cause: HPE Shadowbase MP CSR has completed coordination; i.e. the collector has received coordination completion confirmation for any/all coordinating consumers.

Effect: The collector displays the given EMS message.

Recovery: No recovery required. To disable the message, set the COORDINATIONMSG parameter to NONE.

2124	MPSERV UNABLE TO OPEN MLOG FILE '<file name>', ERROR <Guardian error code>
------	--

Cause: The given MPSEVER consumer was unable to open its log file, denoted by the MPLOG cons_param.

Effect: The MPSEVER consumer displays the given EMS message and continues running.

Recovery: Look at the given error code to determine the cause of the problem, and proceed accordingly.

2124	MPSERV UNABLE TO WRITE TO MPLOG FILE '<file name>', ERROR <Guardian error code>
------	---

- Cause:** The given MPSEVER consumer was unable to write a message to its log file, denoted by the MPLOG cons_param.
- Effect:** The MPSEVER consumer displays the given EMS message, closes its log file, and stops trying to log further message to the file. The MPSEVER consumer continues running.
- Recovery:** Look at the given error code to determine the cause of the problem, and proceed accordingly.

2124	MPSERV MPLOG FILE <file name> TRUNCATED DUE TO FILE FILLING, PREVIOUS MESSAGES HAVE BEEN CLEARED
------	--

- Cause:** The given MPSEVER consumer attempted to write to its log file, which was full.
- Effect:** The MPSEVER consumer displays the given EMS message, clears the contents of the log file, and writes the new messages to the log file.
- Recovery:** N/A

2124	MPSERV SKIPPING SQLMP DDL EVENT DUE TO RESUMEDDL SKIP REQUEST
------	---

- Cause:** The MPSEVER consumer was in a PAUSED state due to an MP CSR error (configurable via the MPERRORACTION parameter.) The user issued a RESUMEDDL SKIP command in AUDCOM.
- Effect:** The MPSEVER consumer displays the given EMS message, skips the given SQL/MP DDL statement, and goes back into the IDLE state.
- Recovery:** N/A

2124	MPSERV RETRYING SQLMP DDL EVENT DUE TO RESUMEDDL RETRY REQUEST
------	--

Cause: The MPSEVER consumer was in a PAUSED state due to an MP CSR error (configurable via the MPERRORACTION parameter.) The user issued a RESUMEDDL RETRY command in AUDCOM.

Effect: The MPSEVER consumer displays the given EMS message and retries the given SQL/MP DDL statement. If the retried command succeeds, the consumer goes into the IDLE state. If the retried command fails, the consumer goes back into the PAUSED state to await another RESUMEDDL command.

Recovery: N/A

2124	MPSERV DDL EVENT ERROR IN SESSION <transaction id>. ACTION SKIP - <sql error message>
------	---

Cause: The MPSEVER consumer encountered an error processing an MP CSR event. The MPERRORACTION parameter was set to SKIP.

Effect: The MPSEVER consumer displays the given EMS message and skips the MP CSR event that caused the error. The consumer continues running.

Recovery: Look at the give SQL error message and resolve accordingly.

2124	MPSERV DDL EVENT ERROR IN SESSION <transaction id>. ACTION STOP - <sql error message>
------	---

Cause: The MPSEVER consumer encountered an error processing an MP CSR event. The MPERRORACTION parameter was set to STOP.

Effect: The MPSEVER consumer displays the given EMS message and abends.

Recovery: Look at the give SQL error message and resolve accordingly.

2124	MPSERV DDL EVENT ERROR IN SESSION <transaction id>. ACTION PAUSE AND WAIT FOR RESUMEDDL COMMAND FROM AUDCOM - <sql error message>
------	---

- Cause:** The MPSERVER consumer encountered an error processing an MP CSR event. The MPERRORACTION parameter was set to PAUSE.
- Effect:** The MPSERVER consumer displays the given EMS message and enters the PAUSED state.
- Recovery:** Look at the give SQL error message and resolve accordingly. In most cases, the user may open AUDCOM and enter the appropriate RESUMEDDL command to resolve the issue.

2124	MPSERV TRANS SWAP FATAL ERROR: UNABLE TO FIND ORIGINAL MP CSR TRANSACTION FOR ID: <transaction id>
------	--

- Cause:** The MPSERVER consumer was in a PAUSED state when the user attempted to run a RESUMEDDL RETRY command. An internal error occurred causing the consumer to be unable to complete a transaction swap.
- Effect:** The MPSERVER consumer displays the given EMS message and abends.
- Recovery:** This is an internal error that shouldn't happen. Contact Shadowbase support and send over the ZZSA.

2124	MPSERV DDL EVENT ERROR IN SESSION <transaction id>. ACTION UNDEFINED. STOPPING - <sql error message>
------	--

- Cause:** The MPSERVER consumer encountered an error processing an MP CSR event. The MPSERVER consumer was unable to determine what state to fall into upon this error.
- Effect:** The MPSERVER consumer displays the given EMS message and abends.
- Recovery:** This is an internal error that shouldn't happen. Contact Shadowbase support and send over the ZZSA.

2124	MPSERVER INTERNAL RESUMEDDL RESPONSE TOKEN IS INVALID; FATAL ERROR
------	--

Cause: The MPSERVER consumer encountered an error processing an MP CSR event. The MPERRORACTION parameter was set to PAUSE. The RESUMEDDL command initiated from AUDCOM contained an invalid action. The only valid RESUMEDDL actions are: RETRY, SKIP, and STOP.

Effect: The MPSERVER consumer displays the given EMS message and abends.

Recovery: Re-enter the RESUMEDDL command with the correct action.

2124	MPSERVER RESUMEDDL RESPONSE WAS TO STOP
------	---

Cause: The MPSERVER consumer encountered an error processing an MP CSR event. The MPERRORACTION parameter was set to PAUSE. The user entered a RESUMEDDL STOP command in AUDCOM.

Effect: The MPSERVER consumer displays the given EMS message and abends.

Recovery: No recovery required. Optionally, look at the SQL error displayed in a prior EMS message that caused the consumer to go into the PAUSED state, and correct the error.

2124	MPSERVER SQL STATEMENT BUFFER OVERFLOW, INCREASE SIZE FROM <#> TO AT LEAST <#>
------	---

Cause: The user entered a command into SBSQLCI that was too large for the MPSERVER consumer's command buffer.

Effect: The MPSERVER consumer displays the given EMS message and abends.

Recovery: Shorten the SQL command entered in some way. The maximum number of characters in a single command supported by MP CSR is 27502, assuming the user is using an ENHANCED control file. If the user's command isn't longer than the maximum buffer size, contact Shadowbase support with the ZZSA.

2124	MPSERV SHADOWBASE RESUMEDDL COMMAND HAS NOT BEEN ISSUES FOR FAILED TARGET DDL OPERATION <transaction id> <entered sql command>
------	---

Cause: The MPSEVER consumer was in the PAUSED waiting for a RESUMEDDL command when it received a timeout (error 40).

Effect: The MPSEVER consumer displays the given EMS message and continues waiting for the RESUMEDDL command.

Recovery: No recovery necessary.

2124	MPSERV NO SQLCI OUTPUT GENERATED FOR COMMAND '<entered sql command>'
------	--

Cause: The MPSEVER consumer attempted to executed a SQLCI command (PURGE, PURGEDATA, or SECURE) and encountered an error.

Effect: The MPSEVER consumer displays the given EMS message and abends.

Recovery: Contact Shadowbase support with the ZZSA file.

2125	XSBCTRLF UNDEFINED RECORD TYPE <#>
------	------------------------------------

Cause: The MPSEVER consumer was sent a control file event that is not supported or is undefined.

Effect: The MPSEVER consumer displays the given EMS message and abends.

Recovery: Contact Shadowbase support with the ZZSA file.

2132	COORDINATION, AND SUBSEQUENT CSR EVENT(S) CANNOT BE PROCESSED UNTIL A RESUME OR RESUMEUPD IS ISSUED FOR THE DOWNSTREAM CONSUMER(S).
------	---

Cause: The user issued a SUSPEND or a SUSPENDUPD command at one point, then entered a SQL/MP DDL command in SBSQLCI, initiating a CSR event while at least one of the downstream consumers are still suspended.

Effect: AUDCOLLN issues this EMS message and continues waiting. Note that AUDCOLLN is NOT reading audit during this conflicted state.

Recovery: Issue a RESUMEUPD or RESUME command to resume any consumers that are downstream to the given collector.

2133	MPSERV SKIPPING SQLMP DDL EVENT DUE TO RESUMEDDL SKIPABORTED REQUEST – A NEW TRANSACTION WILL BE STARTED FOR SUBSEQUENT EVENTS.
------	---

Cause: The MPSEVER Consumer was in a PAUSEABT state when the user issued a RESUMEDDL SKIPABORTED command via AUDCOM.

Effect: The MPSEVER Consumer issues this EMS message, starts a new transaction, and skips the erroneous MP CSR event.

Recovery: No recovery necessary.

2133	TARGET-SIDE TRANSACTION ABORTED – (ALL MP CSR EVENTS SINCE THE START OF SESSION <collector trans id> ROLLED BACK). OVERRIDING MPERRORACTION AND PAUSING; WAITING FOR RESUMEDDDL - <sql error text>
------	--

Cause: Shadowbase replicated an MP CSR event that failed on the target and caused SQL to automatically abort the MPSERVER Consumer's transaction.

Effect: The MPSERVER consumer issues this EMS message and goes into the PAUSEABT state where it waits on a RESUMEDDDL command from AUDCOM. RESUMEDDDL STOP and RESUMEDDDL SKIPABORTED are the only commands allowed to be issued while in this state. Note that a "session" refers to the all commands executed within a unit of work. Usually this will be one command unless a BEGIN WORK is executed; then all commands between the BEGIN WORK and COMMIT WORK are considered a "session." The session in this case is identified by the transaction id sent down by the collector to the consumer.

Recovery: The consumer is now in the PAUSEABT state waiting for a RESUMEDDDL command. While it's in PAUSEABT, manual intervention will be needed to recover any lost/aborted events that need to be applied to the target. If the MPLOG parameter was configured prior to running, open the corresponding MPSERVER log file, and look for the events corresponding to the session, copy them, and run them manually on the target. The user will need to manually resolve any issues that may result from the SQL DDL commands entered this way. If the MPLOG parameter was NOT configured, and the SBSQLCI session is still running, issue a HISTORY, VERBOSE command in SBSQLCI. Scroll up and collect all of the commands whose "Transaction ID" field match the session id displayed in the EMS message. These will need to be mapped, copied, and manually executed on the target. This process is easier if you already have these commands in a script on the source system – if so, run SBSQLCI, turn on CHECKMODE, and run all of the commands through the program, collect the output, then run the resulting commands on the target.

New and Modified User Information Messages

The new user information messages are:

1643	BAD REQUEST – RESUMEDDL SKIPABORTED CAN ONLY BE ISSUED IF THE MPSEVER CONSUMER IS IN THE PAUSEABT STATE
------	---

Cause: The MPSEVER consumer is in the PAUSED state due to an MP CSR error and the user issued a RESUMEDDL SKIPABORTED command via AUDCOM.

Effect: The message is displayed on the terminal and no other effect is produced.

Recovery: No recovery needed, but perhaps the user intended to issue a RESUMEDDL SKIP command instead.

1644	BAD REQUEST – RESUMEDDL SKIP CAN ONLY BE ISSUED IF THE MPSEVER CONSUMER IS IN THE PAUSED STATE
------	--

Cause: The MPSEVER consumer is in the PAUSEABT state due to an MP CSR error causing the consumer to lose its transaction. The user issued a RESUMEDDL SKIP command via AUDCOM.

Effect: The message is displayed on the terminal and no other effect is produced.

Recovery: No recovery needed, but perhaps the user intended to issue a RESUMEDDL SKIPABORTED command instead.

1645	BAD REQUEST – RESUMEDDL RETRY CAN ONLY BE ISSUED IF THE MPSEVER CONSUMER IS IN THE PAUSED STATE
------	---

Cause: The MPSEVER consumer is in the PAUSEABT state due to an MP CSR error causing the consumer to lose its transaction. The user issued a RESUMEDDL RETRY command via AUDCOM.

Effect: The message is displayed on the terminal and no other effect is produced.

Recovery: No recovery needed, but perhaps the user intended to issue a RESUMEDDL SKIPABORTED or RESUMEDDL STOP command instead.

Installation Instructions (HPE NonStop Shadowbase)

Shadowbase is now built for HPE NonStop installations using the HPE PAK and UNPAK utilities – refer to the G06.16 and above operating system documentation (or contact Shadowbase Support) for more information. The SBPAK file described below is delivered as a self-extracting file.

1) Binary FTP the SBPAK from your PC (if you obtained the files from a Gravic FTP site) or DVD. If FTP was used, enter the following commands to alter the file code.

For G-based systems:

FUP ALTER SBPAK, CODE 700

For H-based and J-based systems:

FUP ALTER SBPAK, CODE 800

For L-based systems:

FUP ALTER SBPAK, CODE 500

2) To “unpak” the files from the self-extracting SBPAK, enter the following command:

RUN SBPAK, *.*.*, VOL <restore volume and subvolume>, LISTALL, MYID

Where <restore volume and subvolume> is the location you want the Shadowbase files installed into, e.g. \$DATA.SHADBASE

Note: Using UNPAK to extract the files instead can result in errors and fail to unpack the files with older versions of UNPAK. Using the RUN command as listed above does not have this issue.

3) Follow instructions in the Installing Shadowbase on a NonStop System section of the ***HPE Shadowbase NonStop Installation and Planning Manual***. You must now run either the INSTALL or ENSINST script.

IMPORTANT NOTES ABOUT USER EXITS:

IF YOU HAVE USER EXITS DEVELOPED WITH COBOL, YOU MUST RECOMPILE THE USER EXIT “PROGRAMS” SO THAT THE LATEST VERSION OF USRXLIBO OR USRXLIBN (the native library) IS ADDED TO THE RESULTING OBJECT FILE THAT WILL SUBSEQUENTLY BE INCLUDED IN THE FINAL BIND OR LINK.

AFTER COMPILING THE USER EXITS (IF NECESSARY AS DISCUSSED ABOVE), RUN BINDUSRX, BINDXYPR OR LINKUSRX TO GENERATE THE

CUSTOM CONSUMER (AUDCONS OR AUDCONSN). NOTE THAT THIS SHOULD BE DONE AFTER THE STANDARD INSTALL IS RUN.

Known Problems Remaining

- 1) The DOC Roll command is not supported in configurations that use a QMGR out to either the OSS SQL/MX target or an Other Servers target.
- 2) For SQL/MX, the documented maximum value for the 32-bit float is 3.40282347e+38. However, the HPE Shadowbase for Other Servers ODBC Server gets a range error starting with 3.40282337e+38 (difference highlighted and underlined). The cause of this issue remains under investigation.
- 3) With a NonStop AUDMON configured for auto-restart and backup CPUs are configured for the Collectors and Consumers, you may have to issue an AUDCOM RUN command for replication to continue after a CPU failure if you have previously entered a SUSPEND COLL command followed by a RESUME COLL command (prior to the failover of AUDMON). A workaround for this is to enter an AUDCOM run command immediately after the RESUME COLL command.
- 4) During SOLV testing at Gravic, the consumer received an SQL error -1031 for an insert to one of its complex test tables. HP's SQL error description for the error is:

A supplied numeric value exceeds the declared precision of the column, some of the supplied values for DECIMAL or VARCHAR columns are invalid, or the supplied row is too long. Also, the disk process might have encountered a bad column in a stored row or a value in an update on a row that would change the length of a VARCHAR column in an entry-sequenced table.

After investigation of the problem, the error was subsequently reproduced with SQLCI and submitted to HP for further investigation. At this time, it appears that this is caused by a problem in SQL/MP. A fix by HP has yet to be released. HP development is working on the problem.

It has not been determined as to the exact cause of the error -1031. It could be related to the order of certain data types in the schema definition for the SQL table, resulting in an alignment issue that is not handled properly in SQL/MP. The error appeared to be related to SQL processing on the column named C63 in the following series of columns in the schema:

,c62	SMALLINT	DEFAULT NULL
,c63	varchar(1)	default NULL
,c64	DATETIME	year to fraction

HP has indicated that a workaround could resolve the problem. Add the following define before starting AUDMON.

```
ADD DEFINE = _SQL_CMP_FORCE_EXE_EXPR,  
FILE <any Shadowbase TARGETFILE file name>
```

Contact Shadowbase support if the SQL error -1031 continues after you have attempted the workaround.

5) The SOLVMGR process will abend if configured with an INACTIVITYDELAY greater than 2147 seconds (approximately 35 minutes).

6) In certain conditions, the Consumer may continuously fail (with the Collector restarting at the same audit position) replaying audit and not be able to advance. This problem may occur under the following conditions:

- The Consumer is replicating to a file (table) with a unique alternate key (index).
- The Consumer receives a Guardian error 40 (timeout) on an insert operation due to a lock collision between two active transactions on the alternate key (index).
- The Consumer receives a Guardian error 59 (disk file is bad) on its subsequent delete of the record (row) to 'reverse' the failed insert.

If this issue occurs, either set the Consumer parameter TRACKTX to enable the TRACKTX processing or set the Consumer parameter TIDFILECONSTANT to ON to enable TIDFILE processing to allow the Consumer to continue after the initial failure.

7) A failed insert into a file (table) with alternate keys (index) may leave the record inserted in the base file (table) with no corresponding alternate key records (index entries) under the following conditions:

- The maximum number of active transactions in a Consumer has been reached so the Consumer is doing the "RESOLVETRANS" processing (pre-committing transactions to allow additional simultaneous transactions).
- The Consumer receives a Guardian error 40 (timeout) on an insert operation due to a timeout due to a lock collision on the alternate key (index).

In this case, the record (row) inserted into the base file (table) may not be properly removed.

This condition can be avoided either by turning RESOLVETRANS off or by ensuring that enough Consumers are configured to avoid RESOLVETRANS processing.

8) SOLV has the potential of missing records at the end of a file in a sequence when used with SOLVMGR (solv manager) for 'file chasing'. This problem is more likely to occur with low values of SOLVMGRDELAY, SOLVEOFDELAY, INACTIVITYDELAY and ACTIVITYDELAY are set. By default these values are set to 60 seconds. The minimum values for these configurations are 1 second, however, Gravic does not recommend using values lower than the default to avoid the potential for data loss.

9) For SQL/MX replication, running the Consumer process on the target system is problematical. In several cases, the Consumer needs to access file / schema information from the catalog that is only available on the source system. To prevent issues, the Consumer process should be run on the source system. This is the normal configuration, and will be configured by default when using the SBUNDROS and SBBIDROS scripts.

10) The maximum key length supported by SOLV, SOLVMX, and SOLVNV remains 255 bytes.

11) SOLVMX cannot handle fully qualified SQL/MX table names more than 48 characters in length.

12) For SQL/MX replication, the maximum column name length that can be successfully replicated is 76 characters.

13) When performing a SOLV load of a partitioned file or table from NonStop to NonStop, no warning message is output if the TARGETFNAME is not specified. A warning is output when replicating to a HPE Shadowbase for Other Servers target when configured with the basefile as the source with the ALLPARTITIONS DBS parameter ON. This warning indicates the number of events for which no associated DBS was found in the Consumer during the SOLV load. This warning message should be output in this case if performing a SOLV load to a NonStop target as well.

14) The Consumer does not properly handle statement text that is generated in excess of 64KB for a single statement. This can occur in certain extreme cases with SQL/MX replication or with excessive user exit “put” function calls. To generate statement text exceeding this limit requires the equivalent of a table with approximately 1000 columns with an average column name length in excess of 60 bytes.

15) For DDL replication, certain DDL events require multi-stage processing by HPE Shadowbase replication. Create events, for example, are processed by creating the base file, then performing an alter of the file, and then creating alternate key files (if any). In the event of a failure in the middle of such multi-stage processing, the remaining stages may not be properly applied. If the base file was created, but the Consumer was stopped before the alternate key files were created, for example, on restart the Consumer will attempt to reprocess the event. Since the base file already exists, however, attempting to recreate it fails with an error 10. After receiving this error, the remaining stages are not attempted, and the alternate key files are not created.

16) When replicating from an Enscribe source file using a DDL definition out to a HPE Shadowbase for Other Servers target, the proper where clause is not generated if no dummy TARGETFILE is specified. This results in an empty where clause, causing a SQL error in the Consumer process. The workaround is to specify a dummy TARGETFILE.

17) SQL/MX tables may not be replicated if the SQL/MX catalog they are located in has multiple definition schema versions defined. This can occur if a new schema is added to a catalog following an upgrade to a newer version of SQL/MX. This can result in a consumer process replicating MX data to log errors similar to the following:

```
\NODE1.$DATA1.ZSDQAG7S.G2CS4Q00 / MXCAT_V.MXSCH_T.TABL_1  
(DBS=DBS-SB-01) -> TARGET IS MISSING. CONSUMER ABENDING.
```

In order to get around this issue, the MX catalog metadata must be upgraded so that there is only one version of the schemas define. The following command can be run to upgrade the catalog metadata:

```
UPGRADE ALL METADATA IN CATALOG catalog  
  [ RESTRICT | CASCADE ]  
  [ output-spec ]  
  
output-spec is:  
  [ log-to-spec ] [ REPORTONLY ]  
  
log-to-spec is:  
  { [ LOG TO ] OUTFILE oss-file [ CLEAR ] | LOG TO HOMETERM }
```

18) After a QMGR restart with the CONS in a suspended (SUSPEND CONS command issued) state, you will have to issue another SUSPEND CONS command prior to the RESUME CONS command. Issuing just a RESUME CONS command will result in an erroneous “ERR 1244: DIDN'T FIND ANY OBJECTS TO RESUME” message to be issued. Issuing another SUSPEND CONS command followed by the RESUME CONS command will resume the replication flow.

19) There is an issue with the QMGR that may cause it to stop with the message:

```
LOGIC ERROR IN C:\Release\nsb\audqmgr\TNS-E\..\src\qmgr.cpp AT LINE 607:  
ASSERTION cons_seqno > last_cons_position.get_queue_sequence_number()  
FAILED. PLEASE CONTACT SHADOWBASE SUPPORT AT SBSUPPORT@GRAVIC.COM
```

This may occur if the attached Consumer goes down with no data outstanding and no data is flowing from the Collector. The QMGR will fail to restart with the message:

```
QMGR QMGR-SRC-CJCQ1(\GRAVIC1.$CJCQ1): UNABLE TO FIND QUEUE FILE SEQUENCE  
NUMBER (1) MATCHING $*.CJCQ1.QF* REQUIRED BY CONSUMER FOR RESTART
```

This can only occur if there is no data queued for the Consumer. If this error does occur, you can purge the QMGR restart file to allow it to successfully restart without data loss.

20) Enabling the DBS ALLPARTITIONS parameter does not work if using the SBCOLL Collector program. The workaround for this is to use either a DBS SOURCEFILE parameter value with a wildcarded volume, or to use multiple DBSes for the

individual partitions. In either case, the DBS TARGETFILE should be set to the primary partition of the target file.

- 21) SBDDLUTL cuts out the display of subfields within groups containing OCCURS in the DIFF command.
- 22) SBDDLUTL flattens redefines in a DIFF command for the non-flattened DDL output.
- 23) SQL/MX bidirectional replication sometimes sends events to the target even if those events are part of a transaction that will ultimately be aborted. These events can trigger the Collector to change from SUSPNEXT to REPLSUSP mode, which can seem misleading since the events will not be applied to the target.
- 24) The DUMPTX COLL command does not display information for any tracked peer consumer transactions when using bidirectional Enscribe or SQL/MP replication with SBCOLL. The peer consumer transactions are counted in the number of active transactions displayed by the command, but they are not output.
- 25) The AUD level parameters for global translog support, PEERTRANSLOG and TRANSLOG, do not work for bidirectional replication with SBCOLL.
- 26) Replicating a created table using the new CREATEOPTIONS AUDITED parameter setting can lead to inconsistency between the Collector stats for the source and target tables.
- 27) Enscribe Compare (SOLVNV) does not work correctly for entry-sequenced or relative files. Incorrect comparison results are output when attempting to compare either entry-sequenced or relative Enscribe files.
- 28) When the DOIFNOTRGATRUN DBS parameter is set to WARN1ST, WARN, or CONTINUE, if the target file is not present when the first event is replicated, the Consumer will log a message that the target file is missing and that it is continuing. The Consumer will still abend, however, when it gets an error attempting to open the file (for Enscribe target) or get the schema information (for SQL/MP targets).
- 29) The checking for the DOIFNOTRGATRUN DBS is only performed when an event is encountered for a target file for which the consumer does not have the schema information in memory. For SQL/MP tables, if the target table is dropped after one or more events have been replicated, and the schema is already loaded, the consumer will ultimately receive an error 11 for the event, ignore it, and continue, even if DOIFNOTRGATRUN is set to STOP.
- 30) When replicating ALTER events between multiple tables in succession (e.g. Table A to Table B, then Table B to Table C), extra ALTER and COMMIT events can be logged in the STATS.

- 31) Consumer processes that are on a different node from their corresponding AUDMON process are not automatically killed when a non-responsive state is detected.
- 32) AUDCOLL processes do not currently handle NEXTDOC commands.
- 33) The HADW script enables the SBCOLLAUTOADVMOODE parameter by default. If configuring an HADW config, be sure to explicitly disable this parameter to ensure that the Shadowbase collector does not automatically advance when the needed audit trail is no longer on the system.
- 34) The Consumer's I/O Tracing output can contain garbage characters when the Consumer encounters an error 73 (table or record is locked).
- 35) Regarding SOLV/SOLVnV shadparm input processing, Shadowbase currently only supports KEYVALUESTART and KEYVALUESTOP values up to 45 bytes long. With other methods SOLV can be pushed to use a KEYVALUESTART value above 239 or 240 bytes, but goes into an infinite loop.
- 36) SBGETREALTARGET will cause the Consumer to trap with an illegal address reference if called for an "end trans" (commit or abort) event, and tracing is enabled. SBGETREALTARGET should never be called for an end trans event.
- 37) The DBS ALLPARTITIONS parameter does not work with the SBCOLL Collector program.
- 38) The Enscribe compare program, SOLVNV, does not compare entry sequenced or relative files correctly. Incorrect counts of matches and mismatches are returned.
- 39) The CONS SOLVMAXSESSIONS parameter cannot be set to 0 if a QMGR is configured. Attempting to do so will result in error 1235: REQUIRED PARAMETER NOT SET -> COLLNAME when attempting to add the Consumer object.
- 40) No validation occurs when the DBS VIEWNAME parameter is set to confirm that the DBS SOURCEFILE is a SQL/MP table. This parameter is only valid for SQL/MP source tables.
- 41) The Enscribe compare program, SOLVNV, traps with an illegal address reference when attempting to compare unstructured files.
- 42) SBDDLUTL duplicate field name numbering logic does not work as expected. Currently it is zero-relative by default, skipping the second index. This is the case regardless of the DUPOFFSET parameter. Unique field names are still generated, but might be less predictable than expected.

43) Internal Gravic testing revealed SOLV traps when running in VnV mode for Enscribe Compare. The issue occurs when SOLV attempts to create its marker history file. There are two workarounds to this issue:

- Temporarily run a Shadowbase environment configured with a SOLVONLY Consumer, and use the marker history file generated by the consumer
- Manually create the SOLV marker history file using FUP, see commands below (extents can be set as desired):

```
- SET TYPE K
- SET FORMAT 1
- SET CODE 7687
- SET EXT(64, 128)
- SET REC 4062
- SET KEYLEN 24
- SET MAXEXTENTS 160
- CREATE SOLVMKRH
- SECURE SOLVMKRH, "NNNN"
```

44) Internal Gravic testing revealed that setting the VVMKRHSTEXTENT overwrites the SOLV VVMKREXTENT value. The workaround is to create the SOLVMKRH file with the desired extents via one of the steps described above.

45) When attempting to SOLV load a SQL/MP or SQL/MX table with more than 3 primary key columns with a mix of ascending and descending columns in the key, the SQL optimizers for both SQL/MP and SQL/MX will choose a poor access path to the data which results in SOLV or SOLVMX eventually encountering an error 35 or slowing to a crawl. A different method must be used to load this table other than SOLV.

46) Attempting to use UNPAK to unpack the shadowbase self-extracting PAK files will fail with older versions of UNPAK up to T1255H01^AAI. The workaround is to run the self-extracting PAK files instead of using UNPAK, using the method outlined in the installation instructions.

47) Modifications to a DDL record definition during replication will not be picked up by the consumer when either performing a RELEASE DBS followed by a CAPTURE DBS, or removing the existing DBS and adding a new DBS. The consumer process must be restarted in order for the change to be picked up.

48) The Queue Manager will abend with the following error if MAXCONSWRITES is set to 1 and the Consumer process is on a remote system:

```
QMGR QMGR01(\GRAVIC1.$GSQM1): ERROR 28 CALLING PROCEDURE
SETMODENOWAIT ON FILE \VIV1.$GSCS1 , OPERATION 71, PARAM1 0,
PARAM2 115
```

```
QMGR QMGR01(\GRAVIC1.$GSQM1): ABENDING: UNABLE TO SET
EXPAND MESSAGE TRANSMISSION PRIORITY
```

The workaround for this issue is the set MAXCONSWRITES to a value of 2 or higher. The default value for the MAXCONSWRITE parameter is 15.

- 49) The Shadowbase Consumer sporadically encounters SQL error -8204 / FS error 1061 when doing inserts into partitioned entry-sequenced SQL/MP tables. Based on internal testing, a work-around for this issue is to do a SET DBS CONTROLTABLEEXT1 (SEQUENTIAL INSERT OFF) for the related DBS object. Note that the CONTROLTABLEEXT parameter has superseded CONTROLTABLEEXT<n>, and if present, CONTROLTABLEEXT<n> will be ignored (CONTROLTABLEEXT1, in this case).
- 50) Shadowbase BASE24 handling has a potential efficiency issue when processing events associated with a DBS where the source or target file definitions could be associated with multiple files. This is especially apparent when using "true dynamic" selection mapping (i.e. no MAPFROMSOURCE and SOURCEFILE/TARGET is used only for mapping) and when using MAPFROMSOURCE ON and the SOURCEFILE has a wild-card (*) in the volume or sub-volume. For now a workaround would be to try and limit the use of these types of DBS specifications for BASE24.
- 51) Enscribe DDL definitions with data types containing negative scale causes SBDDLUTL to generate invalid SQL/MP output. For instance, an Enscribe DDL data type of TYPE BINARY 64,-16 will generate a corresponding SQL/MP type of NUMERIC(18, -16).
- 52) SBDDLUTL traps if supplied an input file it cannot open (e.g. if the input file does not exist, or if SBDDLUTL does not have permission to open it)
- 53) When VARCHAR fields are included in the primary key of a SQL/MP table with AUDITCOMPRESS on, SQL/MP will pad VARCHAR columns in the WHERE clause of UPDATE statements with spaces, up to the configured size of the VARCHAR. This can result in data corruption.
- 54) MP CSR name mapping does not currently have awareness of the relationship between physical and virtual drives, i.e. when performing name mapping, SBSQLCI does not know ahead of time if the mapping performed for the PHYSVOL clause will generate a command with a physical volume that is in the same storage pool as the virtual volume used in the rest of the command. For example:

Given the following name mapping definitions:

```
MAP NAME $VDRV1.*.* TO $VDRV1B.*.*
MAP NAME $QA TO $QAX
```

The following CREATE TABLE statement will cause the target MPSEVER consumer to fail:

```
CREATE TABLE $VDRV1.TEST.NEWTAB (  
    ...  
)  
CATALOG $VOL.CAT  
PHYSVOL $QA;
```

The failure will occur if the physical drive \$QAX is not within the same storage pool as the virtual drive \$VDRV1B.

The work-around here is to just be conscientious about what name mappings are configured for the MP CSR environment if you're going to be running commands with the PHYSVOL clause. Check the storage pools on the target node ahead of time to make sure you properly configure name mapping. To do so:

```
$VOL SUBVOL 10> SCF  
SCF - T9082H01 - (01MAY18) (02MAR18) - 08/02/2019 17:56:25 System \GRAVIC1  
(c) Copyright 1986-2016 Hewlett Packard Enterprise Development LP  
(Invoking \GRAVIC1.$SYSTEM.SYSTEM.SCFCSTM)  
1- INFO POOL $*, DETAIL
```

- 55) When generating an explain plan using SOLVMX, certain conditions can cause the explain output to contain this error: '*** ERROR[2105] This query could not be compiled because of incompatible Control Query Shape (CQS) specifications. Inspect the CQS in effect.' This is the result of the sql compiler selecting an invalid SPLIT_TOP plan from the supplied select statements and CQS generated by SOLVMX. This problem does not seem to affect the SOLVMX load.
- 56) The MPSEVER consumer will not back out target-side transactions following a PAUSE that happens within the transaction/session.

For now, it is recommended to either avoid running multiple commands under a single explicit transaction on the source, or avoid using MPERRORACTION PAUSE until this issue is resolved.

- 57) Running an environment with some form of coordination (CSR coordination / DDL coordination), a Queue Manager on any thread, and TRACKTXFILE(S) has the potential to lock up the environment in some scenarios. It is recommended to avoid this by doing one of the following things:
- Not using a QMGR in the Shadowbase environment if MP CSR or DDL Replication with coordination will occur under the same AUDMON
 - Turning MPCOORDINATION OFF for MP CSR or turning DDL coordination off (SET DBS CREATES NOCOORDINATION) for DDL replication
 - Using TID files with TIDFILECONSTANT ON instead of TRACKTXFILE

- 58) SBSQLCI currently has a problem parsing SQL DDL commands that contain SQL/MP reserved/key words as part of column/file names. For instance, any place in a command that could take any user-specified data value cannot be a reserved/key word. Example:

```
SBSQLCI 1> CREATE TABLE CSRSRC.PARTTB (
mul(1, 2)> ID INT NOT NULL,
mul(1, 3)> NAME CHAR(30),
mul(1, 4)> PRIMARY KEY(ID))
mul(1, 5)> CATALOG $DATA9.VPCAT
mul(1, 6)> PARTITION (
mul(1, 7)> $VDV001.CSRSRC.PARTTB FIRST KEY(1001) CATALOG $DATA9.VPCAT,
mul(1, 8)> $VDV002.CSRSRC.PARTTB FIRST KEY(2001) CATALOG $DATA9.VPCAT,
mul(1, 9)> $VDV003.CSRSRC.PARTTB FIRST KEY(3001) CATALOG $DATA9.VPCAT
mul(1,10)> );
-- SQL operation complete.

SBSQLCI 2> CREATE INDEX CSRSRC.PARTIND1 ON CSRSRC.PARTTB (ID, NAME) PARTITION (
mul(2, 2)> $VDV001.CSRSRC.PARTIND1 FIRST KEY(1001) CATALOG $DATA9.VPCAT,
mul(2, 3)> $VDV002.CSRSRC.PARTIND1 FIRST KEY(2001) CATALOG $DATA9.VPCAT,
mul(2, 4)> $VDV003.CSRSRC.PARTIND1 FIRST KEY(3001) CATALOG $DATA9.VPCAT);
...CSRSRC.PARTTB (ID, NAME) PARTITION ( $VDV...
      ^
SBSQLCI ERROR: syntax error (unexpected "NAME")
```

In the preceding example, the column name in an index column list cannot be NAME because NAME is a SQL/MP reserved/key word (in this case in a WITH SHARED ACCESS clause).

- 59) The QMGR occasionally abends on a Shadowbase restart with a message similar to the following:

```
19-06-26 10:56:03 \H2.$CJCQM *GRAVIC.100.V66 004530 QMGR
QMGR-SRC-CJCQM(\H2.$CJCQM): COLDSTART REQUIRED: UNABLE TO FIND FIRST
EXPECTED EVENT IN THE RESTART IPC; ADT POSITION (AA000001,282391212); MAT
POSITION (1, 282423416); MATCH MAT ONLY (FALSE); FIRST MAT POSITION IN IPC
(1, 282423416); LAST MAT POSITION IN IPC (, )
```

Allowing Shadowbase to restart the QMGR will usually clear up this issue.

- 60) AUDMON will not allow a DIRECT environment to be configured without any DBSs and Shadowbase will fail to start if this is done. This may impact CSR-only environments where the only desired replication thread is the CSR thread, and no DBSs are needed. To work around this issue, add a DBS to the *dir configuration script. The source and target table specifications can be anything as this DBS is not effectively in use.
- 61) When SBSQLCI evaluates a TACL define contained within a SQL DDL statement, it does not remove the node name from the statement. This will cause a problem if replicating CSR events to systems that have a different node name than the source system. A work around is to use table/catalog/physvol names directly instead of TACL defines.

- 62) SBSQLCI's config file settings CTRL_FILENAME and CONFIG_FILENAME override the =SBCONTROLFILE and =SBSQLCICONFIG TACL defines. A work-around to this is to delete the CTRL_FILENAME and CONFIG_FILENAME parameters from the config file prior to defining the =SBCONTROLFILE and =SBSQLCICONFIG TACL defines.
- 63) The MPSEVER consumer does not currently have the ability to process certain SQL DDL commands such as ALTER or DROP when replicating MP CSR commands on un-audited target tables. The user will see a SQL error 1353 when attempting to UPDATE or DROP a SQL table that is not audited and perform the actions dictated by the MPERRORACTION parameter. To work around this issue:
- MPERRORACTION is set to SKIP, either
 - Alter the given table to turn audit ON in a separate SQLCI session, then re-execute the DDL commands via SBSQLCI to have them re-replicate, or
 - Manually apply the DDL commands to the target table in a separate SQLCI session
 - MPERRORACTION is set to PAUSE: ALTER the table to turn AUDIT ON in a separate SQLCI session, then do RESUMEDDL RETRY in AUDCOM to retry the command
 - MPERRORACTION is set to STOP: either manually enter the DDL commands in a separate SQLCI session, or ALTER the table to turn audit ON, then restart Shadowbase

[*** End of Document *****]**