

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	Shadowbase
Release:	Gravic Version 6.610 HP NonStop Shadowbase: T1122-AAP (SB Repl/Guardian)
Release Date:	August 1, 2020
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2020(www.gravic.com)
File Name:	IPM6610.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 6.610	6
Changes in Release 6.610 (since Version 6.600).....	10
New Features	10
Problems Fixed	10
New and Modified Commands	12
STATS COLL Command	12
New and Modified SET AUD Parameters.....	12
New and Modified SET COLL Parameters	13
SKIPMPINDEXEVENTS	13
SKIPMXINDEXEVENTS.....	13
New and Modified SET CONS Parameters.....	13
New and Modified SET QMGR Parameters	13
New and Modified SBBIDIR SB_SET Parameters.....	13
SB_USE_3DIGIT_DBS_NAMES	13
SB_USE_CONS_ID_IN_NAMES.....	14
New or Modified EMS Messages	15
New and Modified User Information Messages	15
Installation Instructions (HPE NonStop Shadowbase)	15
Known Problems Remaining	17
[***** End of Document *****].....	28

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	NSB 6.300 May, 2017	May, 2019

	LOADER 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 6.610

- 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.

21) The `START AUD COOL` and `START AUD WARM` command functionality is dependent on stored configuration data, whose structure may change each release. When upgrading or downgrading a Shadowbase environment between release versions, `START AUD COLD` must be used.

22) The secondary extent size used for the Shadowbase `AUDMON` control file has been increased from 64 to 1024, to allow for larger environment configuration sizes.

Changes in Release 6.610 (since Version 6.600)

New Features

- 1) Enscribe Compare now supports comparing over TCP/IP, in addition to EXPAND. See the separate document, IPM6610-SBGW for details.
- 2) Several performance enhancements were made for SBBIDIR when generating configurations for large environments. These changes, however, involve changing the names generated for DBS, CONS, QMGR objects, as well as certain generated files. As a result, these changes are disabled by default. Enabling them when generating configurations for existing environments is not recommended, as it may cause problems due to changes in object and file names used. See the New and Modified SBBIDIR SB_SET Parameters section for details.
- 3) SOLVUTIL was enhanced to allow for sub-second file modifications by setting ACTIVITYDELAYSECONDS to 0. The previous minimum was 1 second. Files with the same exact modification timestamp will be processed in pathname/filename order. This can happen in OSS because some operations set the fractions portion of the modification timestamp to zero.
- 4) Added two new collector params to allow for skipping SQL index events at the collector level, that affect both AUDCOLLN and SBCOLL:
 - **SKIPMPINDEXEVENTS**
 - **SKIPMXINDEXEVENTS**

Refer to the New and Modified SET COLL Parameters section below for more details.

Problems Fixed

- 1) SOLVNV (Enscribe Compare) would not properly compare entry sequenced or relative files.
- 2) Setting the VVMKRHSTEXTENT parameter for Enscribe Compare would cause the marker history file creation to fail with an error 21, and for SOLV to abend. The prior workaround was to use the default extents or manually pre-create the file with the extent sizes desired.
- 3) For Enscribe Compare, SOLV could trap when running in VnV mode for Enscribe Compare. The issue occurred when SOLV attempted to create its marker history file. There were 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"
```

- 4) The MISSINGROOTPART parameter processing, and the CHECK command, could fail to identify cases where the ROOTPARTITION parameter was required, when using ALLPARTITIONS ON, or a wildcarded volume in the DBS SOURCEFILE.
- 5) When using DDL Command Replication with a TRACKTX file enabled, the drain events could be replayed on a restart. This would result in the coordination being performed again, when not needed.
- 6) When performing bidirectional replication of Enscribe ALTER events using DDL audit replication, cases were encountered in internal testing where slow auxiliary audit trail flushing could cause the ALTER events to be ponged (i.e. replicated a second time back to the originating node).
- 7) An issue was encountered in internal testing where SBCOLL could abend due to an audit reading error ARE_INTERNAL_ERROR(-1000) on a call to ARGETANSINAME on SQL/MX DDL events not being replicated. In this case, the collector would restart and continue successfully after the initial abend.
- 8) No indication would be given if the CPU AUDMON was run in, and the CPU configured for AUD BACKUPCPU were the same. An EMS message is now logged in this case, similar to the following:

20-06-16 12:46:58 \VIV1.\$CJCMN GRAVIC.100.V66 002013 BACKUP
AND PRIMARY CPU NUMBERS ARE THE SAME - AS A RESULT, NO
BACKUP AUDMON PROCESS WILL BE LAUNCHED
- 9) MISSINGROOTPART was not displayed in the INFO CONS and OBEYFORM CONS command output.

New and Modified Commands

STATS COLL Command

The STATS COLL command was modified to display a message when the Collector is at EOF, and the restart lag may be growing as a result of there being no further audit events occurring.

The updated description of the RESTART INFO (MAT POSITION) section is as follows:

RESTART INFO (MAT POSITION) – This section of the output is to show the RESTARTFILE checkpoint information (i.e., the point in the TMF audit trail where the Collector will start reading from if it has to restart).

SEQ shows the current audit trail Sequence number saved in the RESTARTFILE.

RBA shows the restart position within the audit trail expressed as the relative byte address.

TIME shows the restart point timestamp.

LAG (CUR) shows the difference between current time and the restart point timestamp.

When HPE Shadowbase is caught up, and is not receiving any audit (none being generated) the restart point timestamp reflects the last audit event read. In this case, the following text is displayed on the next line:

Collector is caught up to audit EOF (restart lag may grow)

DIFF LTS shows the difference between the most recent audit trail event read and the restart position timestamp. If the difference is greater than 100 hours, it will be displayed as “>100 HOURS”. (A “*” is displayed next to the “DIFF LTS” field if the restart point is lagging behind the threshold defined for the COLL CHKPTWARNDIFF parameter.)

New and Modified SET AUD Parameters

None in this release.

New and Modified SET COLL Parameters

SKIPMPINDEXEVENTS

SKIPMPINDEXEVENTS { ON }
 { OFF }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.610		ON	OFF	ON	NO

Determines whether to replicate or skip all SQL/MP index events at the collector level. This applies to both SBCOLL and AUDCOLLN. Defaults to ON – meaning both collectors skip SQL/MP index events by default as of 6.610. Note that this was not the case in prior versions.

SKIPMXINDEXEVENTS

SKIPMXINDEXEVENTS { ON }
 { OFF }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.610		ON	OFF	ON	NO

Determines whether to replicate or skip all SQL/MX index events at the collector level. This applies only to SBCOLL. Defaults to ON – meaning SBCOLL skips SQL/MX index events by default as of 6.610. Note that this was not the case in prior versions.

New and Modified SET CONS Parameters

None in this release.

New and Modified SET QMGR Parameters

None in this release.

New and Modified SBBIDIR SB_SET Parameters

SB_USE_3DIGIT_DBS_NAMES

This parameter extends the dbs limit in the SBBIDIR configuration. The original 'legacy' limit was 1,296 (36*36) DBSs. The new limit is 46,656 (36*36*36) DBSs. When this parameter is set to the default, zero, SBBIDIR uses the legacy limit.

Note: Enabling this feature is not recommended when generating configurations for pre-existing environments, as it may result in changes to the DBS object names generated, which may cause problems.

```
SB_SET SB_USE_3DIGIT_DBS_NAMES <value>
```

```
where
```

```
<value> is 0 or 1.
```

SB_USE_CONS_ID_IN_NAMES

This parameter changes how the consumer, qmgrs and associated files are named. This parameter is only applicable when there are only two nodes in the configuration. Set this parameter to '1' to enable this feature. Enabling this feature will tell the script use the two character SBBIDIR consumer object id in all the Shadowbase consumers, qmgrs and associated file names. The default setting for this parameter is '0'. This is the legacy setting. This setting uses an internal sequence number in the naming of consumers, qmgrs and associated files names.

Note: Enabling this feature is not recommended when generating configurations for pre-existing environments, as it may result in changes in object and file names used, which may cause problems.

```
SB_SET SB_USE_CONS_ID_IN_NAMES <value>
```

```
where
```

```
<value> is 0 or 1.
```

New or Modified EMS Messages

2013	BACKUP AND PRIMARY CPU NUMBERS ARE THE SAME - AS A RESULT, NO BACKUP AUDMON PROCESS WILL BE LAUNCHED
------	--

- Cause:** AUDMON was started in the same CPU as configured for the AUD BACKCPU parameter.
- Effect:** No backup AUDMON process will be created.
- Recovery:** If no backup process for AUDMON was intended, none needed. However, to create a backup AUDMON process, shutdown the environment and alter the configuration by either specifying a different CPU to launch AUDMON in, or changing the configured AUD BACKUPCPU parameter value.

New and Modified User Information Messages

None in this release.

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

- 2) 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.
- 3) 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.
- 4) 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.
- 5) 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.

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

7) 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.

8) 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.

9) 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.

10) 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.

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

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

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

14) 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.

15) 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.

16) 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.

17) 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.

18) 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 }
```

19) 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.

20) 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.

21) 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.

- 22) SBDDLUTL cuts out the display of subfields within groups containing OCCURS in the DIFF command.
- 23) SBDDLUTL flattens redefines in a DIFF command for the non-flattened DDL output.
- 24) 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.
- 25) 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.
- 26) The AUD level parameters for global translog support, PEERTRANSLOG and TRANSLOG, do not work for bidirectional replication with SBCOLL.
- 27) 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.
- 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) AUDCOLLN processes do not currently handle NEXTDOC commands when a QMGR is configured.
- 33) The HADW script enables the SBCOLLAUTOADV MODE 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 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.
- 39) 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.
- 40) 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.
- 41) 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.
- 42) 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.

43) 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.

44) 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.

45) 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.

46) 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).

47) 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.

48) 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).

49) 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)

- 50) 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.
- 51) 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.SCFSTM)
1- INFO POOL $*, DETAIL
```

- 52) 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.

- 53) 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.

- 54) 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

- 55) 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).

- 56) 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.

- 57) 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.
- 58) 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.
- 59) 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.
- 60) 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

- 61) The use of the NEXTDOCTIME and NEXTDOCTRIGGER collector parameters causes several issues in the DOC rolling logic:
- The DOC rolls repeat once per second, even if NEXTDOCTRIGGER is set to TIMEOFDAY
 - The NEXTDOCTIME parameter is ignored, and the DOC rolls begin as soon as the Shadowbase environment is started
 - The EMS message pertaining to the initiated DOC roll shows a timestamp that is off by several hours; between 2 and 6 hours seen in internal testing.
- 62) Due to existing volume recovery logic, the Queue Manager holds onto transaction abort records until it sees another event without the volume recovery flag set (typically an IUD event.) In some cases, when no IUD events are sent down the Queue Manager's thread for long enough (12 hours), the Queue Manager will clear these transactions and output an EMS message saying that transactions are being "age cleared." However, the Queue Manager still acknowledges to the collector that the abort event is stored, so the collector's restart point is NOT held back for the 12 hour duration. In this scenario, either a subsequent replicated event without the volume recovery flag, or the eventual age clear successfully cleans up the transaction from memory without causing any issues.
- 63) SBSQLCI, the front-end tool for Shadowbase SQL DCR, gets a SQL error 1353 when creating a constraint with the DEFERRED option. SBSQLCI creates implicit transactions for commands entered while an explicit transaction is not in effect; this transaction is active at the time a CREATE CONSTRAINT with option DEFERRED is run, causing error 1353. There is no work-around for this.
- 64) With older versions of SQL/MX (versions prior to 3.7), SYSKEY values lower than the current maximum value in the target table cannot be replicated via exact SYSKEY replication (i.e. when SHAD_OVERRIDE_SYSKEY is enabled). As a result, the following limitations apply:
- Uni-directional replication only
 - Partitioned files are not supported
 - 'Missing' rows cannot be re-replicated to insert them (i.e. the restart point cannot be reset, expecting missing rows to be inserted)
- With SQL/MX 3.7 and newer, these limitations do not apply.
- 65) The new Shadowbase Collector (SBCOLL) supports uni-directional DDL audit replication for Enscribe and SQL/MP files; but does not currently support bi-directional DDL audit replication.
- 66) SOLV cannot be used to load Enscribe key-sequenced queue files.
- 67) HPE Shadowbase Compare for Enscribe cannot be used to compare key-sequenced queue files or unstructured files.

- 68) When configuring HPE Shadowbase Compare for Enscribe, the VVPROCESSNAME process name portion of the value must be 4 characters or less (not counting the \$ character).
- 69) The SOLV KEYVALUESTARTINC parameter cannot be set to the value OFF. This applies to both SOLV loading and HPE Shadowbase Compare for Enscribe.
- 70) When using KEYSPECIFIER with HPE Shadowbase Compare for Enscribe to access the source and target files by an alternate key, the KEYVALUESTART and KEYVALUESTOP options cannot be used.

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