

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	Shadowbase
Release:	Gravic Version 6.410 HP NonStop Shadowbase: T1122-AAJ (SB Repl/Guardian)
Release Date:	March 31, 2018
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2018 (www.gravic.com)
File Name:	IPM6410.pdf

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
Mature and Deprecated Features	3
Supported Database and Operating System Versions.....	4
Special Notes for Version 6.410	5
Changes in Release 6.410 (since Version 6.320).....	9
New Features	9
Problems Fixed	11
New Features Description.....	17
Reverse Replication Cutoff.....	17
Upgrading Shadowbase Releases with the QMGR	18
SBCOLL Bi-directional Replication of Enscribe and SQL/MP	18
New and Modified Commands	19
STATS CONS Command (Modified)	19
New and Modified SET AUD Parameters.....	20
New and Modified Set COLL Parameters	20
New and Modified SET CONS Parameters.....	21
New and Modified SET DBS Parameters.....	22
New and Modified SET QMGR Parameters	24
New and Modified TACL Parameters	24
MXSTATEIGNOREERR TACL Parameter	24
MXSTATESUPPRESSMSG TACL Parameter	24
SBCOLLMYSTART TACL Parameter for use with SBCOLL.....	24
New or Modified EMS Messages	27
New and Modified User Information Messages	32
New or Modified SOLVCOM User Messages	34
Installation Instructions (HPE NonStop Shadowbase)	36
Known Problems Remaining	38
[***** End of Document *****].....	43

Disclaimer

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

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

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

Mature 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
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

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 to load files and tables should immediately begin planning to move to the Shadowbase Online Loader (SOLV).	NSB 6.300 May, 2017	May, 2019
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 6.410

- 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 6.410 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 6.410.
- 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 6.4nn, the version has thus been updated according to the VNN format to V64. 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 V63 to V64 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, USRXCOB, and USRXCOBA, 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) The license validation has been updated to recognize several new processor models without the need for a shadtier file. The following HPE Integrity NonStop X processors have been added as tier 6:

- NSX-G (NS3 X1)
- NSX-H (NS3 X2)
- NSX-I (NS7 X2)

The following HPE Virtualized NonStop processors have been added as tier 7:

- NSV-B (vNS-EE)
- NSV-C (vNS-EE)
- NSV-D (vNS-EE)
- NSV-F (vNS-EC)
- NSV-G (vNS-EC)
- NSV-H (vNS-EC)

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

20) A new program, SBDDLUTL is included in this release, but is intended for a future release of Shadowbase. Refer to the *HPE NonStop Shadowbase SBDDLUTL Manual* for more information.

21) The default value of the DOIFNOTRGATRUN DBS parameter changed from WARN1ST to STOP. See the New and Modified SET DBS Parameters section for details.

Changes in Release 6.410 (since Version 6.320)

New Features

- 1) When running the INSTALL script, if the specified catalog does not exist, it will now prompt for if you would like it to be automatically created.
- 2) REJECTFILE security is now configurable through the REJECTFILESECURITY parameter.
- 3) When running the SBBIDIR STRTSB script, the prompt for a COLD or WARM start is no longer present, nor is the prompt to clean up the configuration (CU). Shadowbase defaults to starting COLD.
- 4) The SOLV VnV Marker Restart File params VVMKRRETENTION and VVMKRHSTRETENTION now support/can be set to a max of 18,262 days, which is 50 years including leap years. In addition, the SOLVMKRRETENTION and SOLVMKRHSTRETENTION consumer parameters can now be set to a max of 18,000 days.
- 5) When running the CU (cleanup) script (i.e. SRCU, <prefix>CU), the user will be prompted to type 'Y' or 'N' for each file before purging. The file will be purged if 'Y' is selected. The default option is 'Y' if no option is selected by the user. Also, the user can type '!' after the CU script name to run the cleanup without prompts (i.e. SRCU !).
- 6) The CU (cleanup) script (i.e. SRCU, <prefix>CU) will now purge QMGR files.
- 7) SOLVMX has two new parameters MXSTATEIGNOREERR and MXSTATESUPPRESSMSG which give SOLVMX a list of SQLSTATE values to not abend on/not log to EMS, respectively.
- 8) The DBS source file and target file parameters for SBUNIDIR SB_ADD DBS now can only take a fully qualified path name (i.e. \system.vol.subvol.file) if the file isn't omitted.
- 9) STATS AUD output for DIFF (cur) updated to reflect special diff cases. DIFF values less than zero will display as "00:00:00.000*" and values greater than 365 will display as "> 365" as opposed to the output in previous releases ("?:?:?:.???").
- 10) The Shadowbase Consumer now supports monitoring of commit times to aid in diagnostics of performance issues. It will now track the number of commits completed on the main program loop (non-waited,) the number of commits the Consumer had to wait for (that have not completed on the main program loop, waited,) the total amount of time waited for non-waited commits, the total amount of time waited for waited commits, the average time waited for waited commits, and the

overall average time waited. These statistics can be viewed through the STATS CONS command under the COMMIT WAIT TIME INFO section.

- 11) The DOC can now be rolled from a NonStop source node via the SBCOLL Collector object. DOC rolls are either executed immediately if a direct Consumer is configured, or queued if a QMGR is configured. This feature is not yet available for the AUDCOLL Collector object.
- 12) SBCOLL and QMGR processes now recognize AUDMON failure within a much smaller window of time.
- 13) Reverse replication cutoff is now configurable at the Collector level through the REPLICATION parameter.
- 14) SOLV-MX can now handle quoted column names containing reserved words.
- 15) Creating a TAACL define called =SBLICENSE can be used to validate the license for all Shadowbase products. The define needs to be set to type CLASS MAP, FILE <license-file>, where <license-file> is the Shadowbase “shadpass” file. If define =SBLICENSE does not exist, Shadowbase products will use the “shadpass” file they are designated to use in their configuration.
- 16) The obey file “runccomp” has been edited to run the C compiler CCOMP to compile SQL/MX code for user exits. The “runccomp” file has to be manually edited by the user to switch between compiling for SQL/MP and SQL/MX code.
- 17) SBBIDIR and SBBIDROS scripts now accept a new argument for the SB_ADD NODE tacl routine. This new argument coincides with the new bi-directional cut-off feature in the Collector and its settings, and the argument can be set to a value of REPLICATION, SUSPEND, or SKIP. It can also be left blank, in which case it will default to SUSPEND. See details for the bi-directional cut-off feature for more information on how the new SB_ADD NODE argument configures the Collector.
- 18) SBHADW script supports wild carding on the DBS_PARAM, AUD_PARAM, COLL_PARAM, and CONS_PARAM parameters. The ‘*’ character is the wild card.
- 19) SBHADW defaults to using the native Collector instead of the non-native Collector.
- 20) SBHADW does not default to using ERADS EMS.
- 21) SBHADW deploys the Shadowbase TAACL scripts (i.e. *STRTSB, *STOPSB, *CONF, etc.) to the target node, or sub-volume, after prompting. This is similar to the behavior of SBBIDIR.
- 22) Removed COLD/WARM start prompts from SBHADW.

- 23) SBHADW defaults to setting the Collector parameter ADTAUXMASK to 11111111111111, which specifies that the Collector use all 15 auxiliary audit trails.
- 24) SBHADW produces a *CU TACL script, which is a clean-up script. This *CU script behaves like the SBBIDIR *CU script, except it does not prompt before purging.
- 25) SOLVCOM checks for unprintable characters and displays a more helpful message to the user if any are encountered.
- 26) HPE Shadowbase licensing has been modified to support the NonStop Dynamic Capacity core licensing (also known as Total Capacity on Demand). Shadowbase licensing monitors the number of cores active on each CPU in the system. If the number of cores active exceed the number of cores specified by the Shadowbase license, Shadowbase logs messages to EMS indicating the Shadowbase license is invalid and needs to be updated.

Shadowbase will now monitor whether additional cores have been activated by the NonStop Dynamic Capacity core licensing feature. If they have, and the base number of cores is less or equal to the number specified by the Shadowbase license, the license is valid and Shadowbase will not complain.
- 27) HPE Shadowbase now supports the larger entry-sequence file record sizes – up to 27,576 bytes using exact entry sequence replication. When exact entry-sequence replication is configured, HPE Shadowbase uses unstructured access and builds the entry sequence data blocks internally to update the target file. Prior to this release, the maximum record size supported for exact entry-sequence was 4072 and the maximum block size was 4096.
- 28) SBCOLL now supports TRANSLOG based bi-directional replication for Enscribe and SQL/MP files. See the New Features Description section for details.

Problems Fixed

- 1) There was an issue when replicating from HPE NonStop Enscribe files out to SQL that may cause Shadowbase to erroneously drop some columns from replication in certain configurations. This may have occurred if all of the following were true:
 - You are replicating from Enscribe files to any SQL database except SQL/MP.
 - You have periodic cleanup of inactive files and schemas enabled in the Consumer. Periodic cleanup is enabled and runs every 24 hours by default.
 - You are replicating from files that may become candidates for being removed from memory due to lack of activity over the cleanup period.
 - The Consumer that is replicating the files has very few different schemas. This problem was uncovered in a Consumer that had two DBS's attached. The

- likelihood of the issue occurring drops as the number of different DBS/files increase.
- 2) In a replication environment with a QMGR on the target system, the QMGR process trapped when the source system went down instead of performing processing based on its abnormalshutdown configuration.
 - 3) When replicating DDL create events involving SQL/MP partitioned indexes, the Consumer process abended on a “catalog does not exist” error when it attempted to create an index due to using an incorrect length for the catalog name.
 - 4) The Consumer process would abend during user exit processing when attempting to do a SBPUTCOLUMN() under the following conditions:
 - You are replicating from an Enscribe source file to an HPE Shadowbase for Other Servers target.
 - No SOURCEREC is specified for the DBS.
 - A dummy table is defined via the DBS TARGETFILE.
 - 5) The Consumer process would fail on a false out of memory error following file/schema cleanup processing where an entry for a DBS with a VIEWNAME specified was cleaned up. The workaround for this was to set the tacl params SBSRCCLEANUPAGE 0 and SBSRCCLEANUPRECUR 0.
 - 6) The scripts generated by SBUNDROS and SBBIDROS did not qualify ./sbmon to run sbmon, and so would fail if the current working directory (.) was not in the PATH environment variable. The generated scripts now qualify ./sbmon, and will run without any requirements on the PATH environment variable.
 - 7) The shadparm.ini files generated by SBUNDROS and SBBIDROS included the [GENERAL] section first. Since the first instance of a given parameter is used rather than last, this prevented parameters specified in the [GENERAL] section from being overridden by values in specific sections. SBUNDROS and SBBIDROS will now generate the shadparm.ini file with the [GENERAL] section at the bottom of the file instead of the top.
 - 8) The SBCOLLRESPONSIVENESS TACL parameter did not take on a reasonable default value in response to the Shadowbase configuration type (sync or async.) SBCOLL will now set a reasonable default depending on the status of the ZDLENABLED parameter.
 - 9) The QMGR would hang in an unresponsive state if it were unsuccessfully stopped by the AUDMON. Now Shadowbase objects are automatically killed if the STOP command times out or otherwise results in an error. This behavior now affects all Shadowbase objects.

- 10) SOLV was not checking the SOLVDISPLAYBANNER TACL parameter before deciding whether to print the SOLV banner, resulting in undefined behavior (and therefore unpredictable banner printing).
- 11) When running SOLV, an erroneous message was sent regarding a non-native Collector, despite this message not applying to a SOLV load.
- 12) The Consumer would abend when processing a DBS with a missing target file.
- 13) The Consumer was assuming replication was caught up if the Collector was backed up or not running, leading to confusing “caught up” messages in high-load situations.
- 14) Audited temporary files were not being ignored by Shadowbase.
- 15) Setting SOLVMAXSESSIONS to 0 with a QMGR would result in an AUDMON error 1235.
- 16) Queue files are no longer checked for Guardian file system error 10’s.
- 17) The SBCONSORDERCOLS TACL param checking was incorrect, resulting in the inability to disable column ordering in the Consumer.
- 18) When SOLV loads a SQL/MP table from a restart point, the restart key-column value no longer overflows in the message to print out over 2000 ‘00’s.
- 19) *HPE Case #10-170309-7559: Shadowbase collector won’t start.* Under certain circumstances, the QMGR incorrectly cleared its start position. This occurred under the following conditions:
 - a. Shadowbase replication was shutdown with all data in the queue processed.
 - b. Shadowbase replication was successfully restarted without any data being replicated.
 - c. The consumer was stopped and auto-restarted.

In this case, the QMGR stopped with an EMS Logic Error 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 issue has been corrected.

- 20) Under certain circumstances after a restart, the QMGR displayed negative counts for the events queued to the Consumer when caught up. This problem has been corrected.

- 21) *HPE Case #10-170602-3803: EN:COLLECTOR on Target system fails to start.*
When upgrading from a previous version of Shadowbase with events in the QMGR queue from the prior release, the QMGR failed with an EMS message indicating that the consumer replied with an error 2 (unsupported operation). The Consumer uses an error 2 to indicate that the message is from a different version of Shadowbase. The QMGR failure prevents the replication environment from successfully starting.
- The QMGR has been improved to support multiple versions of messages in the queue to aid in upgrading. See [Upgrading Shadowbase Releases with the QMGR](#) later in the Softdoc for more information.
- 22) Internal testing uncovered several situations where the QMGR could send extra events that would be applied by the Consumer in the event of one of the processes (Collector, QMGR, or Consumer) failing and being restarted. This could result in events being re-applied at the target even when either TRACKTX or TIDFILE CONSTANT is configured. These issues have been fixed.
- 23) SOLVMX would abend on SQL/MX error/warning 6008 (table statistics unavailable) instead of logging the error and continuing. This issue has been fixed.
- 24) When running the SBUNIDIR script, if the DBS source file and target file are not given a fully qualified path name, i.e. the system/node name is not specified, the SRCONF file that gets generated does not include the system/node name. This issue has been fixed.
- 25) The QMGR was not honoring the MAXCONSWRITES parameter – instead, it would always use the maximum (15) value when sending messages to the Consumer. This release fixes the issue.
- 26) *HPE Case #10-170624-8951: URGENT – Shadowbase SOLV abended with run time error 4.* When the SOLV parameter IPCMAXBUFFERSIZE was set to a value higher than 32,767, the consumer failed with an arithmetic overflow trap. This issue has been fixed.
- 27) SOLV would overwrite its internal source file name value during a SOLV LOAD of a partitioned source table if the TARGETFNAME parameter was set. This was fixed and it no longer overwrites it – instead the source file name will contain the value of the current partition being loaded, and the data file field will contain the name of the primary partition.
- 28) There were cases that when doing a SOLVMX load, SOLVMX would use the range partition keys instead of the hash partition keys when the source table was hash partitioned. This created an error when SOLVMX tried to load, because SOLVMX would try to load the range partition key into the hash partition key on the target. This problem has been fixed so that when SOLVMX loads a table that is hash partitioned, it will always use the hash partition key.

29) The Consumer's performance could become extremely slow under the following conditions:

- Resolve locks are enabled and TIDFILECONSTANT is disabled.
- The environment is shutdown while resolve locks processing is started.
- The restart point is adjusted prior to restarting.

In this situation, records can become "hung" in the TIDFILE resulting in the slow down. In previous releases, the only workaround was to delete the hung records from the TIDFILE. This release corrects the issue, the Consumer will continue to perform normally and the hung records will eventually be removed.

30) Collectors were not being started in HIGHPIN mode on L06 systems. This release fixes the issue.

31) The new Collector, SBCOLL, does not yet have support for TRACKTX processing when connected directly to a consumer. However, you were able to configure TRACKTX processing for SBCOLL, which could result in previously committed data being re-applied during restart. Configuring TRACKTX processing will now be prohibited for configurations with SBCOLL connected to a consumer. Note that the QMGR does support TRACKTX. Configurations using TRACKTX where SBCOLL connects to a QMGR are valid and are permitted.

32) When replicating a SQL DDL Create Index event, the Consumer could overrun its buffer, leading to a trap.

33) When reaching EOF on the last SQL/MP table in a series, SOLV could produce an error 002017 (INVALID COLUMN KEY COUNT) when using SOLVMGR.

34) When running out of buffers for reads on \$RECEIVE, the Consumer would abend instead of allocating another buffer.

35) Through internal testing, we found that AUDCONSN would incorrectly abend with an O2N-CONVERSION ERROR ON INPUT when replicating from a SQL table to Enscribe. This occurred whenever a column that corresponded to a DDL PIC type with a leading implied decimal point (such as PIC SV9(1)) was replicated. This issue has been fixed.

36) SBHADW configurations produced the *STRTSB and *STOPSB Shadowbase TACL scripts that would fail when executed on systems which do not have the TACL #DEFAULTS included in their #PMSEARCHLIST. The *SRTRTSB and *STOPSB scripts were altered so that they will work regardless of the #PMSEARCHLIST.

37) Setting TRACKTXFILE no longer results in an erroneous message about the use of a non-native Collector object.

- 38) The Consumer could hit a logic error related to running out of IPC buffers and abend.
- 39) *HPE Case #10-170824-4833*: The Consumer could hit an illegal address reference and trap when creating a SQL index.
- 40) All Shadowbase objects are now stopped with the same procedure, which eliminates a problem where the QMGR was hanging on exit.
- 41) The Consumer now ignores confusing Error 10s on Enscribe queue files.
- 42) Allow STOP CONS when the Consumer is in recovery mode, preventing confusion when trying to manually stop Consumer during autorestart.
- 43) Messages from the Consumer to a Direct Writer could be corrupted while attempting to reconnect and resend under the following conditions:
- The TCP/IP connection between the Consumer and Direct Writer is broken.
 - NETMAXERRORS is set to a value > 0, so the Consumer retries the connection without stopping.
 - Large buffers (>32K) are being sent by the Consumer.
- Under these circumstances, the buffers immediately after the reconnect could have several bytes corrupted. This corruption could lead to unpredictable behavior, including the DW repeatedly failing, database update failures, or corrupted data being inserted in the database.

This problem has been corrected.

- 44) Under certain rare circumstances, an aborted transaction can become stuck in the collector after a restart, eventually resulting in the transaction being aged cleared after 12 hours. This problem has been corrected.
- 45) The collector's restart point can hang if it fails and restarts while replicating DDL operations. This problem has been corrected.
- 46) *HPE Case 10-171002-6816, STATS Coll show processing inserts and deletes, but target files aren't changed*: The Shadowbase Consumer (AUDCONSN) uses a call to FILE_GETINFOLISTBYNAME_ to get information about the source file. This information is required to interpret the audit trail information sent from the Collector. In prior releases (Shadowbase Version 6.320 and earlier), the Consumer skipped, without logging an error, all events for which the FILE_GETINFOLISTBYNAME_ failed for any Guardian error other than error 11, *File Does Not Exist*. This can cause a loss of data when Expand is down and the source file is not accessible, for example.

The Consumer only does the check when it gets an event for a file or table that match its replication selection criteria (DBS SOURCEFILE setting) but it has not seen before (since process startup) to be replicated. Once the information has been

obtained, it is cached in memory. If the information cannot be obtained, the file/table events will not be replicated (with no error being logged). As a result, data loss will occur. This issue is most likely to occur at Consumer startup or soon thereafter.

Issuing a `STATS CONS` command can help diagnose if the error is occurring. Events skipped due to this issue will show up in the event counts as received (in the `NUM RCVD` column), but not as used (`NUM USED` column). In general, any discrepancies between the `NUM RCVD` and `NUM USED` columns should be investigated.

With this patch, the Consumer will now log `FILE_GETINFOLISTBYNAME_` errors and will shutdown, preventing any data loss.

- 47) `SBCOLL` was incorrectly stopping on an `ARE_DATA_FILE_VERSION` error when reading the before image for a compressed updated when the collector parameter `SBCOLLMXSHEMAERROK` is set. This issue has been corrected.
- 48) `SBCOLL` was not handling the restart correctly in a bi-directional environment with multiple threads and could incorrectly pong data on startup. It was not going far enough back in the audit trail to insure that it saw the `SHADTRS` event for all peer transactions at startup. This could cause it to incorrectly identify a peer transaction, that shouldn't be replicated, as an application transaction that is replicated. This problem has been corrected.
- 49) The `QMGR` could not read queue files that have been duplicated due to an internal version of the file create date versus the file create date recorded in the block – the `FUP DUP` operation changes the file create date. This check has been removed.
- 50) HPE Case 10-170520-3856: *Could not issue STOP CONS. Continuous autorestart.* `CONS` could not be stopped while autorestarting.
- 51) Shadowbase processes that do not respond to the `AUDMON`'s `STOP` message will be automatically killed after 60 seconds.

New Features Description

Reverse Replication Cutoff

Reverse replication cutoff is now configurable with the `REPLICATION` parameter currently available for Collector objects. This parameter allows audit events to be replicated or skipped depending on its setting. Audit reading can also be suspended when

an event is read from the audit trail. Various EMS logging features are also enabled through this parameter, including periodic summaries of audit event operation counts.

Upgrading Shadowbase Releases with the QMGR

Shadowbase expects all programs running in the same environment (e.g., all replication programs started by the same AUDMON monitor process) to be the same version and enforces this compatibility by including a version number in its messages. You can upgrade one node at a time using TCP/IP configurations (where the source and target programs are started by different AUDMON processes), as Gravic does maintain inter-version compatibility for the TCP/IP messages (unless specifically identified in the SoftDoc).

Prior to this release, Version 6.410 (T1122^AAI), you could also encounter version issues when upgrading configurations with QMGRs that have data from a previous version in the queue. The old data would not pass the version check, and, as noted in [Item 21](#) in the Problems Fixed section, would cause the environment to fail to start. Because the QMGR receives and sends data asynchronously, it was also possible for the QMGR to write events from the new release into the queue before failing. If this occurred, the queues were not useable with either release and you could not upgrade or failback without purging the data in the queue.

The QMGR has been made to be backward compatible with data from prior releases. The new QMGR will be able to process and convert ipms in queues made by any previous versions of the QMGR after version 6.100. Note that you still will not be able to revert back to the previous version while there is data in the queue.

This release will also be forward compatible with future releases of the QMGR, unless otherwise specified in the future release. Future upgrades from this release to a new release will be able to be reverted back to this release without emptying the queue.

SBCOLL Bi-directional Replication of Enscribe and SQL/MP

SBCOLL now supports bi-directional replication of Enscribe and SQL/MP events. Previously, SBCOLL only supported bi-directional replication of SQL/MX events.

Note: If replicating both SQL/MX events, and SQL/MP or Enscribe events, separate collector processes must be used for SQL/MX replication, and non-SQL/MX replication.

For bi-directional replication of Enscribe and SQL/MP events, a similar mechanism is used to that of AUDCOLLN. TRANSLOG files must be configured for the Consumer processes, and special EXCLUDE DBS objects must be added to the Consumer process under the SBCOLL collector process to specify the peer consumer TRANSLOG file. These special DBS objects must have SPECTYPE set to EXCLUDE, SCREENTRANS

ON, and SOURCEFILE set to the TRANSLOG of the peer consumer process for which the DBS is excluding events from replication. One DBS must be added for each TRANSLOG file used by a peer consumer for which events are to be excluded from replication. See the CONS TRANSLOG parameter, and the DBS SCREENTRANS parameter descriptions for further details.

By default, SBCOLL only allows two starting points when replicating in a bi-directional Enscribe and/or SQL/MP environment - at the current RESTARTFILE's position, or if that does not exist then at the current master audit trail's EOF. None of the other Collector audit starting positions (such as ADTSTARTTIME, ADTSTARTRBA, etc) are supported. This approach guarantees that bi-directional environments will always avoid transaction ping-pong situations (the incorrect looping of a transaction from a source to a target, then back to the source).

More specifically, the default Shadowbase processing is to always check for a Collector RESTARTFILE, and if it exists to use the audit reading positioning information in it. If it does not exist, then any "peer Consumer" TRANSLOG files are searched for a "sync" position; if found, the earliest audit position is used. If no peer Consumer sync positions are found, the Collector will then start replicating from the current master audit trail's EOF.

However, in a bi-directional environment, there are certain situations when it is desirable or even necessary to be able to start reading audit at specific event times or RBA positions in the audit trails. In order to support these cases for bi-directional SQL/MP and Enscribe replication with SBCOLL, the SBCOLLMYSTART TACL parameter can be used. Note that this parameter only pertains to bi-directional replication environments as uni-directional replication environments can start at any Collector ADTSTARTxxx location. See the SBCOLLMYSTART TACL parameter description for further details, and fully read the section before use.

New and Modified Commands

STATS CONS Command (Modified)

Support has been added to the Consumer for monitoring commit times. The STATS CONS command has been updated to show commit wait time information:

```
SQL LOCK ERROR INFO:
...

COMMIT WAIT TIME INFO:
-----WAITED COMMITS----- --NO-WAITED-- ---OVERALL---
NUM COMMITS          TWAIT          AVG TWAIT          NUM COMMITS          AVG TWAIT
-----
          #####  00:00:00.000  00:00:00.000          #####  00:00:00.000

ENSCRIBE I/O INFO:
...
```

NUM COMMITS – represents the total number of each type of commit used by the Consumer. Waited commits result in the program delaying until the commit completes. No-Waited commits allow the program to continue other processing in parallel while the commit completes.

TWAIT – represents the total amount of time the program is inactive waiting for commits to complete.

AVG TWAIT – represents the average amount of time the program was inactive for waited commits (TWAIT / NUM WAITED COMMITS) and for waited + non-waited commits respectively (TWAIT / (NUM WAITED COMMITS + NUM NO-WAITED COMMITS)).

New and Modified SET AUD Parameters

PEERTRANSLOG { peer-translog-filename }
 { (peer-translog-filename [, peer-translog-filename] ...)}

Initial Version	Changed Version	Values Default	Value Specifications	Alter
6.400		None	filename or filename-list	No

Specifies between one and twenty global peer translog filenames. All collectors under this audmon will internally generate SPECTYPE EXCLUDE DBS objects with SCREENTRANS ON for each of the peer translog files for all threads. When using a global AUD TRANSLOG for simple bidirectional replication, specify the peer translog filename here.

TRANSLOG { translog-filename }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
6.400		None	filename	No

Specifies a global TRANSLOG file for simple bidirectional replication environments. All target consumer processes under this audmon will use this TRANSLOG filename if they do not have one specified via the CONS TRANSLOG parameter. For bi-directional replication over Expand, this was designed for configurations with one audmon environment for each direction.

See the PEERTRANSLOG parameter for information on how to specify the global peer translog filenames, which can be used in place of SCREENTRANS EXCLUDE SPEC DBSes for simple bidirectional environments.

New and Modified Set COLL Parameters

REPLICATION { replication options }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
6.400		See Below	See Below	Yes

The REPLICATION parameter has three components. These components can be entered individually, or as comma-separated values in parentheses. Any values that should remain default can be omitted.

1. Action: The replication action represents the action the Collector will take when it reads an event from the audit trail.
 - a. ON – Events are sent/processed for replication.
 - b. SKIP – Events are skipped.
 - c. SUSPEND – Replication flow suspends. The event will be processed when flow is resumed by ALTERing the REPLICATION action’s value to ON or SKIP.
2. Warn: The warn option specifies the frequency of detailed EMS messages as events are read from audit.
 - a. NOWARN – No detailed messages are logged.
 - b. WARNALL – A detailed message is logged for each processed event.
 - c. WARN – A detailed message is logged for the first processed event in each file.
3. Interval: The interval option specifies the frequency (in minutes) of the periodic summary message of audit event operation counts. To disable periodic summary messages, set this component to -1. To leave this value as the default, either omit it or set it to 0.

When this parameter is not set, it receives its default value, depending on the configuration. Any unset components also receive their individual default values.

- Default REPLICATION value for uni-directional Collector:
SET COLL <coll>, REPLICATION (ON, NOWARN, INTERVAL -1)
Replication is enabled with no warnings.
- Default REPLICATION value for bi-directional Collector:
SET COLL <coll>, REPLICATION (SUSPEND, WARN, INTERVAL 60)
Collector will suspend at the first event, print a detailed message for the event, and issue a periodic summary message every 60 minutes.

To explicitly use these defaults, use this command:
SET COLL <coll>, REPLICATION DEFAULTALL

New and Modified SET CONS Parameters

QFILEMAXRETRIES { number }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.400		60	-1	2147483647	No

Sets the maximum number of retries of an Enscribe queue file IO operation following an error 10 or 11.

Setting this parameter to -1 will allow the Consumer to retry Enscribe IO operations indefinitely as long as there is an error 10 or 11. Settings this parameter to 0 will prevent the Consumer from retrying these operations.

QFILERETRYDELAY { centiseconds }

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.400		50	0	2147483647	No

Sets the delay between retries of an Enscribe queue file IO operation following an error 10 or 11. The delay is represented as centiseconds.

REJECTFILESECURITY { security string }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
6.400		UUUU	security string	No

Sets the security permissions of the REJECTFILE configured through the other REJECTFILE parameters.

New and Modified SET DBS Parameters

CREATEOPTIONS { create-option }
{ (create-option [, create-option]) }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
4.091	6.400	<See Below>	Valid Create-Options	No

Added the create-option 'AUDITED' which allows for replicated CREATE events to be audited by TMF.

This option conflicts with the DBS parameter ENTRYSEQEXACT ON, and will trigger a configuration error if both are present.

This option also is not allowed for bidirectional configurations, and will trigger a configuration error if it is present in such configurations.

DOIFNOTRGATRUN { CONTINUE }
 { STOP }
 { WARN }
 { WARN1ST }

Initial Version	Changed Version	Values Default	Value Specification	Alter
4.091	6.410	STOP	Valid Action	No

This parameter is used to enable run-time checking in the CONS to ensure that the target file/table, associated with the current event for the DBS, exists. Note that this run-time checking is only performed if the CONS CHECKFORMISSING parameter is set to ON. When DOIFNOTRGATRUN is set to CONTINUE, the CONS does not perform this check for the DBS associated with the current audit trail event. When set to STOP, WARN, or WARN1ST the CONS does perform the check and report the condition to EMS. WARN causes the CONS to report the condition for every event. WARN1ST causes the CONS to report the condition for the first event only associated with the target file/table. The CONS abends if STOP is set and it does not find the target file/table. The default is STOP.

Note: If DOIFNOTRGATRUN is set to anything other than STOP, the Consumer will still fail if it tries to apply an event to a target file that does not exist. If DOIFNOTRGATRUN is set to anything other than STOP, it does not override that default Consumer behavior.

PURGEOPTIONS { purge-option }
 { (purge-option [, purge-option]) }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
4.091	6.400	<See Below>	Valid Purge-Options	No

Added the purge-option ‘AUDITED’ which allows for replicated PURGEDATA commands to be audited by TMF.

PURGEDATAOPTIONS { purgedata-option }
 { (purgedata-option [, purgedata-option]) }

Initial Version	Changed Version	Values Default	Value Specifications	Alter
4.091	6.400	<See Below>	Valid Purgedata-Options	No

Added the purgedata-option 'AUDITED' which allows for replicated PURGEDATA events to be audited by TMF.

This option conflicts with the PURGEDATAOPTIONS PARTONLY option, and will trigger a configuration error if both are present.

New and Modified SET QMGR Parameters

No SET QMGR parameters were added or modified since version 6.320.

New and Modified TACL Parameters

MXSTATEIGNOREERR TACL Parameter

<u>Initial Version</u>	<u>Changed Version</u>	<u>Values Default</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Alter</u>
6.400		n/a			No

Specifies a list of SQL/MX errors to ignore during a SOLVMX LOAD operation. Such errors will be logged, but SOLVMX will not abend. The parameter takes a list of SQLSTATE values in a space-separated string. For more information about SQL/MX SQLSTATES look at the SQL/MX Release 3.x Messages Manual.

MXSTATESUPPRESSMSG TACL Parameter

<u>Initial Version</u>	<u>Changed Version</u>	<u>Values Default</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Alter</u>
6.400		n/a			No

Specifies a list of SQL/MX errors/warnings to suppress during a SOLVMX LOAD operation. Such errors will NOT be logged to the terminal or EMS. This parameter does not affect whether or not SOLVMX abends. MXSTATESUPPRESSMSG takes a list of SQLSTATE values in a space-separated string. For more information about SQL/MX SQLSTATES look at the SQL/MX Release 3.x Messages Manual.

SBCOLLMYSTART TACL Parameter for use with SBCOLL

Note: This description of the SBCOLLMYSTART TACL parameter is intended for use with SBCOLL only. Do not use this parameter for AUDCOLLN, except as instructed by Shadowbase Support.

<u>Initial Version</u>	<u>Changed Version</u>	<u>Values Default</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Alter</u>
6.400 (SBCOLL)		0	0	3	No

You can override the initial TMF audit trail starting position where the Collector starts in a bi-directional environment using the TACL SBCOLLMYSTART parameter.

Initial Audit Reading Locations in a Bi-Directional Enscribe or SQL/MP Environment in Shadowbase

By default, Shadowbase only allows two starting points when replicating in a bi-directional Enscribe and/or SQL/MP environment - at the current RESTARTFILE's position, or if that does not exist then at the current master audit trail's EOF. None of the other Collector audit starting positions (such as ADTSTARTTIME, ADTSTARTRBA, etc) are supported. This approach guarantees that bi-directional environments will always avoid transaction ping-pong situations (the incorrect looping of a transaction from a source to a target, then back to the source).

More specifically, the default Shadowbase processing is to always check for a Collector RESTARTFILE, and if it exists to use the audit reading positioning information in it. If it does not exist, then any "peer Consumer" TRANSLOG files are searched for a "sync" position; if found, the earliest audit position is used. If no peer Consumer sync positions are found, the Collector will then start replicating from the current master audit trail's EOF.

However, in a bi-directional environment, there are certain situations when it is desirable or even necessary to be able to start reading audit at specific event times or RBA positions in the audit trails. In order to support these cases, the SBCOLLMYSTART TACL parameter can be used (note that this parameter only pertains to bi-directional replication environments as uni-directional replication environments can start at any Collector ADTSTARTxxx location):

TACL param SBCOLLMYSTART – when supplied and set non-zero, and no RESTARTFILE exists (a RESTARTFILE position ALWAYS overrides the SBCOLLMYSTART parameter), the Collector will change the default audit reading starting positions as follows:

0 = Default mode of operation, as described above. The collector may start only at the current RESTARTFILE's position, or if that does not exist, then at the current master audit trail's EOF.

1 = When there is no "peer Consumer" TRANSLOG sync point and no RESTARTFILE record, you can use the COLLECTOR ADTSTARTxxx parameters to specify a position where the COLLECTOR should start from in the master audit trail. Note that the user accepts the fact that ping-pong could occur in this situation because they could be starting from an unpredictable spot in the audit trail (i.e. the peer Consumer could have been generating audit and without the sync point, this data may be replicated back to the sender). Note that when there is a peer Consumer sync point or RESTARTFILE record, SBCOLLMYSTART 1 will not be considered. **Never enable this mode unless instructed to by Shadowbase Support.**

SBCOLL does not support SBCOLLMYSTART 1. If configured, SBCOLL will log a message and stop.

2 = When there is no RESTARTFILE record and regardless whether a peer Consumer TRANSLOG sync point record exists or not, you can use this setting to tell the Collector to disregard the peer Consumer's sync point and to use any supplied Collector ADTSTARTxxx parameters for the initial audit reading positioning. In this mode, the user accepts the fact that ping-pong could occur because they could be starting from an unpredictable spot in the audit trail (i.e. the peer Consumer could have been generating audit and without the sync point, this data may be replicated back to the sender). Note that when there is a RESTARTFILE record, SBCOLLMYSTART 2 will not be considered. **Never enable this mode unless instructed to by Shadowbase Support.**

SBCOLL does not support SBCOLLMYSTART 2. If configured, SBCOLL will log a message and stop.

3 = When there is no RESTARTFILE record and regardless whether a peer Consumer TRANSLOG sync point record exists or not, you can use this setting to tell the Collector to disregard the peer Consumer's sync point and to use any supplied Collector ADTSTARTxxx parameters for the initial audit reading positioning (with SBCOLL, either using ADTSTARTTIME, or SBCOLLADTSTARTSEQLO and ADTSTARTRBA are allowed). However, unlike SBCOLLMYSTART 2 where ping-pong (and data corruption) could occur, this mode eliminates this possibility by computing an additional "go back" interval earlier than this initial audit reading position. The "go back" interval is computed from the ADTSTARTTIME by subtracting (SBCOLLAUTOADVFACT * FLUSH_BUSY) from this initial position. The Collector will start reading audit from this new position, and not turn on replicating the data until the ADTSTARTTIME position is reached. By doing this additional "go back" interval, the Collector should have gone back far enough to read sufficient peer Consumer TRANSLOG entries to avoid ping-pong (the peer Consumer TRANSLOG entries are flushed to the TRANSLOG every SBCONSFLUSHBUSY interval). When using SBCOLLMYSTART, this should generally be the mode selected.

This is the only option, other than the default setting of 0, that SBCOLL supports.

Enter the following TACL PARAM before starting AUDMON to set SBCOLLMYSTART:

```
PARAM SBCOLLMYSTART <1 or 2 or 3 as defined above>
```

New or Modified EMS Messages

2017	SQL STATEMENT SIZE (<statement_size>) EXCEEDS MAX SIZE (<max_statement_size>). EVENT LOCATED AT ADT <audit_trail_rba>
------	---

Cause: The SQL statement being cached exceeds the max size.

Effect: Consumer willabend.

Recovery:

2129	COLL RESUMING AUDIT READING
------	-----------------------------

Cause: A change in the REPLICATION parameter has resulted in the Collector resuming its reading of the audit trail.

Effect: Audit trail reading will resume.

Recovery: n/a

2129	COLL WILL SUSPEND AUDIT READING UPON NEXT EVENT
------	---

Cause: The REPLICATION parameter's action option was set to SUSPEND.

Effect: Audit trail reading will be suspended when the next event is encountered.

Recovery: n/a

2126	AUDIT EVENT PROCESSED (<explanation of replication action>): ADT POS: <audit position>, SOURCE TRANSID: <source transid>, EVENT TYPE: <event type>, FILE NAME: <source filename>[, DBS NAME: <DBS name if available>]
------	---

Cause: An audit event was processed by the Collector and it is the first event for its file, or the REPLICATION warn option is set to WARNALL.

Effect: The event will be processed, skipped, or saved for later, depending on the REPLICATION action option setting.

Recovery: n/a

2127	REPLICATION SUSPENDED DUE TO REPLICATION PARAMETER SETTINGS AT <time of settings change> (ELAPSED TIME <elapsed time since settings change>). REPLICATION CAN BE RESUMED BY ISSUING AN 'ALTER COLL <collector name>, REPLICATION ON' COMMAND OR EVENTS CAN BE SKIPPED BY ISSUING AN 'ALTER COLL <collector name>, REPLICATION SKIP' COMMAND.
------	--

Cause: The REPLICATION parameter's action option was set to SUSPEND and an event was read from the audit trail.

Effect: Audit trail reading will be suspended until the REPLICATION parameter is altered to set its action option to SKIP or ON.

Recovery: n/a

2128	COLL <collector name>: <number of events since last summary> EVENTS <PROCESSED SKIPPED> SINCE <time of last summary> (ELAPSED <time since last summary>): INSERTS: <number of inserts>, UPDATES: <number of updates>, DELETES: <number of deletes>, COMMITS: <number of commits>, ABORTS: <number of aborts>
------	--

Cause: The REPLICATION parameter's interval option has triggered a summary message.

Effect: n/a

Recovery: n/a

2017	PROCESS <process name> KILLED BY AUDMON FOLLOWING FAILURE OF STOP MESSAGE
------	--

Cause: The AUDMON tried to send a STOP message to the process but the message was sent unsuccessfully.

Effect: The process is automatically killed.

Recovery: Informational only.

2013	SYNC SUBSYSTEM NOT LICENSED FOR <mode> MODE SOME OBJECTS IN THE SYNC SUBSYSTEM WERE CONFIGURED TO RUN IN MODE [sync] – WHICH SHADOWBASE IS NOT LICENSED FOR
------	--

Cause: The user runs a START SYNC command when the coordinators were configured to start in SYNC mode, and Shadowbase is licensed for ASYNC or MONITOR mode.

Effect: The sync subsystem is not started.

Recovery: Modify the sync configuration file by setting the Coordinator MODE property value to ASYNC or MONITOR; or contact Shadowbase Support to resolve licensing issues.

2173	AUDITED CREATES NOT ALLOWED WHEN ENTRYSEQEXACT IS ON
------	--

Cause: A DBS is added with both the ENTRYSEQEXACT parameter set to ON, and the AUDITED option set in the CREATEOPTIONS parameter.

Effect: The DBS is not added.

Recovery: Change the ENTRYSEQEXACT parameter setting to OFF or remove the AUDITED option from your CREATEOPTIONS parameter setting.

2174	AUDITED CREATES NOT ALLOWED IN BIDIRECTIONAL ENVIRONMENT
------	--

Cause: A DBS is added to a bidirectional environment with the CREATEOPTIONS AUDITED option set.

Effect: The DBS is not added.

Recovery: If this is a unidirectional configuration, verify that the TRANSLOG and PEERTRANSLOG parameters are not set, as these parameters are not needed in a unidirectional configuration. Otherwise, it is necessary to remove the AUDITED option from the CREATEOPTIONS parameter.

2175	AUDITED PURGEDATAS NOT ALLOWED WITH PARTONLY SETTING IN PURGEDATAOPTIONS
------	--

Cause: A DBS is added with a PURGEDATAOPTIONS setting that includes both PARTONLY and AUDITED options.

Effect: The DBS is not added.

Recovery: Remove either the PARTONLY or AUDITED options from the DBS settings.

2125	ERROR 2 ON WRITEREAD TO CONS <name>(<process name>); IF TARGET CONSUMER IS RUNNING A VERSION OF SHADOWBASE PRIOR TO 6201, THE COLL TURBOMODE PARAM MUST BE ON AND THE SBCOLLTRBOSIZE TACL PARAM MUST BE SET TO 29964 TO INTEROPERATE
------	--

Cause: QMGR received an Error 2 from the Consumer, which may be caused by an out-of-date version of Shadowbase.

Effect: The QMGR is unable to communicate with the target Consumer.

Recovery: Set the COLL TURBOMODE param to ON and set the SBCOLLTRBOSIZE TACL param to 29964.

2017	PLEASE CHECK TARGET <EMS SYSTEM> FOR MESSAGES TO HELP DIAGNOSE THE PROBLEM.
------	---

Cause: The source node encountered an error caused by an inability to communicate with the target node in a TCP/IP environment.

Effect: The source node is unable to communicate with the target node.

Recovery: The target EMS/system logs may provide more insight into the problem than is available on the source node.

2019	CAN'T DETERMINE SOURCE FILE STRUCTURE; FILE_GETINFOLISTBYNAME_ERROR <error> ON FILE <file_name>
------	---

Cause: The AUDCONS process was not able to obtain the source file structure due to an error on the FILE_GETINFOLISTBYNAME_ call. The error that occurred is specified by <error>, the file it was checking is specified by <file_name>

Effect: The Consumer will stop.

Recovery: There are a several methods of correcting this issue, depending on the root cause of the error. If possible, correct the Guardian error. For example, if the Consumer is running on the target node and the

error indicates that Expand is down, restore Expand connectivity. If the error cannot be corrected, add a MAPTOFILE parameter to the DBS configuration to a file having the same structure on the target system that is accessible.

2002	NO BIDIR RESTART POINT FOUND IN RESTART FILE FOR BIDIR, WILL USE GENERATED BIDIR RESTART EVENT TIME OF <generated time> (CALCULATED BY ADJUSTING <derived-from-text> <value>)
------	---

Cause: SBCOLL is configured for bi-direction replication, and a restart file exists on startup, but the restart file does not have a bi-directional restart position. Typically, this occurs if converting when converting an existing uni-directional replication configuration to a bi-directional one.

Effect: The Collector starts reading audit prior to the restart point for bi-directional cutoff records only. Replication of events begins at the replication restart point present in the restart file. The bi-directional restart time used, as well as the basis for the calculation, are output in the message.

Recovery: Informational only. None needed.

New and Modified User Information Messages

The new user information messages are:

7013	TRACKTX IS NOT SUPPORTED BY SBCOLL WITHOUT A QMGR; CONSIDER ADDING A QMGR TO CONS <cons-name>
------	---

Cause: Configurations enabling TRACKTX for consumers directly connected to the new collector, SBCOLL, are not valid and may lead to committed data being replayed on restart.

Effect: The collector does not successfully start.

Recovery: Correct the configuration. In many cases, replaying committed data on restart does not have any effect on the target database and will not lead to mismatches. If this is the case, consider just disabling TRACKTX. If replaying committed data on restart is an issue, consider adding a QMGR between the specified consumer and the collector. Finally, if adding a QMGR is not an option and you cannot replay committed data on restart, disable TRACKTX and enable TIDFILECONSTANT ON.

7014	OLD OR INCORRECT FILE CODE <file-code> FOR TRANSLOG <translog-fname>
------	--

Cause: SBCOLL detected an existing TRANSLOG file with an incorrect or unsupported file code.

Effect: The collector does not start.

Recovery: Correct the configuration. Change the configuration to use a new file, or remove the existing unsupported or invalid TRANSLOG file. Note that this impacts both sides of the bi-directional replication environment, as the SCREENTRANS ON SPECTYPE EXCLUDE DBS object must match the related peer consumer TRANSLOG. Also, if removing the TRANSLOG file, the corresponding target (peer) consumer must be restarted to regenerate the TRANSLOG file with an appropriate sync point.

7015	SBCOLL ONLY SUPPORTS SBCOLLMYSTART OPTIONS 0 AND 3; OPTIONS 1 AND 2 ARE NOT SUPPORTED
------	---

Cause: The SBCOLLMYSTART TAACL param was specified with a value of 1 or 2, which are not supported by SBCOLL.

Effect: The collector does not start.

Recovery: Correct the configuration. Shutdown the environment (including AUDMON), specify a valid setting for SBCOLLMYSTART, and then restart.

New or Modified SOLVCOM User Messages

<p>IGNORING MALFORMED SOLVCOM PARAMETER(S): <param> UNPRINTABLE CHARACTER <character> FOUND. THIS CAN OCCUR WHEN COPYING AND PASTING A COMMAND FROM A SOURCE THAT SUPPORTS FORMATTED TEXT. IF THE COMMAND LOOKS CORRECT, TRY RE-TYPING IT MANUALLY.</p>

Cause: SOLVCOM detected an unprintable character in user input.

Effect: The input is ignored.

Recovery: Follow instructions given in message. In particular, watch out for the em-dash character as this character can sometimes replace dashes in formatted text documents, resulting in problems when copying and pasting from these sources.

New or Modified SOLV User Messages

?	SOLV: ROW/RECORD LOCKED, <NUMBER OF LOCK RETRIES> OCCURRENCE, FILE SYSTEM ERROR <SYSTEM ERROR>
---	--

Cause: When LOCKWARNINGS ON is set for a SOLV load, this message will be printed for each time SOLV tries to read from a locked record.

Effect: No effect.

Recovery: A warning message only. Unlock the record SOLV is trying to load.

?	SOLV: KEY COLUMNS OF LAST ROW/RECORD LOADED:
---	--

Cause: When LOCKWARNINGS ON is set for a SOLV load, this message will be printed once the max LOCKRETRIES has been reached, and the last row/record loaded by SOLV for that session will be printed. If no row/records are loaded, this does not get printed.

Effect: Shows the last row/record that SOLV loaded for that session.

Recovery: No recovery needed.

?	SOLV: MXSTATEIGNOREERR PARAM MISSING VALUE:
---	---

Cause: Supplying an empty string or a string containing only whitespace as the value for the MXSTATEIGNOREERR param.

Effect: SOLVMX logs message to EMS and abends.

Recovery: Fix the value provided for the MXSTATEIGNOREERR param by providing valid SQL/MX SQLSTATE values in the string.

?	SOLV: MXSTATEIGNOREERR PARAM <PARAMETER LIST SUBVALUE> IS AN INVALID SQL/MX SQLSTATE:
---	---

Cause: The MXSTATEIGNOREERR param contains illegal characters.

Effect: SOLVMX logs the message to EMS and continues running.

Recovery: No recovery needed; illegal parameter value will be ignored.

?	SOLV: MXSTATESUPPRESSMSG PARAM MISSING VALUE:
---	---

- Cause: Supplying an empty string or a string containing only whitespace as the value for the MXSTATESUPPRESSMSG param.
- Effect: SOLVMX logs message to EMS and abends.
- Recovery: Fix the value provided for the MXSTATESUPPRESSMSG param by providing valid SQL/MX SQLSTATE values in the string.

?	SOLV: MXSTATESUPPRESSMSG PARAM <PARAMETER LIST SUBVALUE> IS AN INVALID SQL/MX SQLSTATE:
---	---

- Cause: The MXSTATESUPPRESSMSG param contains illegal characters.
- Effect: SOLVMX logs the message to EMS and continues running.
- Recovery: No recovery needed; illegal parameter value will be ignored.

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

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

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 new 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.

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) Loading Enscribe queue files with UPDATEDUPLICATES ON and ALLOWQUEUEFILE REPLICATEKEY set causes an abend due to an Error 2.

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) Using the ALTER command to change a Collector's REPLICATION setting from ON to SKIP before committing an outstanding transaction will prevent the transaction from being cleared if it has already been seen by the attached Consumer.
- 26) 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.
- 27) The AUD level parameters for global translog support, PEERTRANSLOG and TRANSLOG, do not work for bidirectional replication with SBCOLL.
- 28) 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.
- 29) 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.
- 30) SBCOLL will incorrectly stop on an ARE_DATA_FILE_VERSION error when reading the after image for a compressed updated when the collector parameter SBCOLLMXSHEMAERROK is set, if a MX schema error has already been logged for the file.
- 31) 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).
- 32) 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.
- 33) 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.
- 34) Replicating Enscribe queue files using the CONS UPDATEDUPLICATE ON setting can lead to WRITEREAD Error 2, which can cause failures and the deletion of records.

- 35) Consumer processes that are on a different node from their corresponding AUDMON process are not automatically killed when a non-responsive state is detected.
- 36) AUDCOLL processes do not currently handle NEXTDOC commands.
- 37) 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.

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