

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

| | |
|-------------------|---|
| Product: | Shadowbase |
| Release: | Gravic Version 6.500 HP NonStop Shadowbase: T1122-AAM (SB Repl/Guardian) |
| Release Date: | September 1, 2018 |
| Copyright Notice: | Copyright Gravic, Inc. 1995 – 2018 (www.gravic.com) |
| File Name: | IPM6500.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 6.500 | 5 |
| Changes in Release 6.500 (since Version 6.410)..... | 9 |
| New Features | 9 |
| Problems Fixed | 11 |
| New and Modified Commands | 12 |
| New and Modified SET AUD Parameters..... | 13 |
| New and Modified Set COLL Parameters | 13 |
| New and Modified SET CONS Parameters..... | 13 |
| New and Modified SET DBS Parameters..... | 14 |
| New and Modified SET QMGR Parameters | 14 |
| New and Modified TACL Parameters | 14 |
| New and Modified SOLV Parameters | 14 |
| New and Modified SOLVMX Parameters..... | 15 |
| New and Modified SOLVMX or SOLV Parameters..... | 15 |
| New or Modified EMS Messages | 16 |
| New and Modified User Information Messages | 20 |
| New or Modified SOLVCOM User Messages | 20 |
| Installation Instructions (HPE NonStop Shadowbase) | 23 |
| Known Problems Remaining | 25 |
| [***** End of Document *****]..... | 31 |

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 6.500

- 1) The Restart file format has been changed for 6.500 in a way that is not directly compatible with older Restart files. If upgrading from an older version of Shadowbase with the intent to use an existing Restart file, it is necessary to upgrade the configured Restart file using the RESTEDIT tool. More information is available in Step 4 of the Installation Instructions section.
- 2) 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.
- 3) 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.
- 4) 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.
- 5) 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.
- 6) For those that have user exits, you must use the 6.500 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 6.500.
- 7) 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.
- 8) 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.

9) 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.

10) 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.

11) 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.

12) 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.

13) 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.

14) 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 6.5nn, the version has thus been updated according to the VNN format to V65. Any EMS filters using these attributes for filtering must likewise be updated. Alternatively, the SBEMSOLDSSID TAACL parameter can be enabled to have Shadowbase use the old values for the subsystem owner and version.

15) The EMS subsystem version has been updated with this version from V64 to V65 to match the two most significant digits of the current Shadowbase release version id.

16) 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.

17) 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.

18) 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.

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

20) SBHADW configuration scripts now sets the TAACL parameter SBCOLLAUTOADVMODE in the AMON script to a default value of 0. This parameter used to default to 1.

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

and continue on as usual.

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

Changes in Release 6.500 (since Version 6.410)

New Features

- 1) The HPE Shadowbase SBCOLL program now reports additional information for errors encountered while making use of TMF ARLIB calls. In the event of an internal library error, the TMF subsystem and ARLIB error code are now displayed in EMS output. See the New or Modified EMS Messages section for more information.
- 2) The Consumer will now emit warnings to EMS if it receives events that are not used. This behavior is configurable via the CONS SKIPWARNRATE parameter.
- 3) SBHADW configuration scripts will now produce a SOLVLD script, which is similar to the SOLVLD script produced by SBBIDIR configuration scripts.
- 4) SBHADW configuration scripts will now support SB_ADD TACL_COMMAND parameters. See the *HPE NonStop Shadowbase High Availability Designated Winner (HADW) Configuration Tool User Manual* for details.
- 5) SBHADW configuration scripts will now support SB_ADD TACL_DEFINE parameters. See the *HPE NonStop Shadowbase High Availability Designated Winner (HADW) Configuration Tool User Manual* for details.
- 6) SBUNDROS SB_ADD CONSOS now can accept the ‘-’ character as a default for the backup TCP/IP address, backup DOCW/Direct Writer process name, and backup DOCW Cleaner process name. Note that if the backup TCP/IP address is set to a ‘-’, then the other two parameters need to be set to a ‘-’, and if the backup TCP/IP address is not set to a ‘-’, then the other two parameters cannot be set to a ‘-’.
- 7) SBBIDROS QMGRs can now have their primary and backup CPUs specified in the SB_ADD QMGR parameter. See the *HPE NonStop Shadowbase Tools and Utilities Manual* for details.
- 8) SBUNIDIR and SBBIDIR configuration scripts no longer switch volume in the AMON script when starting up the AUDCOM or AUDMON program, which means that if an error occurs while starting up either of those programs, then the user will remain in the “configuration” volume and not be moved to the Shadowbase installation volume.
- 9) The USRXEXCEPTIONINIT API now supports tracing via the Consumer parameter TRACE.
- 10) The AUDMON will now retry when it is unable to stop a given process using the PROCESS_STOP_ system call. This happens when a process does not respond to the stop message that the AUDMON sends it. Retry behavior is configurable using the new STOPPROCESSMAXRETRIES and STOPPROCESSRETRYDELAY

parameters.

- 11) The AUDMON will now warn the user if it tries to start a process with the same name of an existing process. If the environment had previously failed to shut down a process, it is now clearer on startup that the old process needs to be stopped before Shadowbase can start.
- 12) Now when replicating ‘create’ events for Enscribe files with alternate-key files, Shadowbase will force the target node name into the alternate-key specifier for the target base file.
- 13) When a Consumer that connects to an Other Servers environment encounters a “SQL statement ‘too big’ error”, a new error message is now being logged to EMS to indicate this issue.
- 14) New SOLV and SOLVMX enhancements.
 - a. Timeout messages will now be log to the console and/or EMS while the program is waiting for an acknowledgment from the consumer. The default timeout is every 15 minutes. Use the existing CONSUMERIPCTIMEOUT parameter to override this value.
 - b. SOLV output will now display any configured parameters not used by the program. This will help assist a user in determining any mistyped parameter names supplied to the program.
 - c. Use PARTONLYFNAME to configure SOLV to load a specific partition of a SQL/MP table. PARTONLYFNAME is not a new parameter. It only applied to partitions of Enscribe files in the past.
 - d. Use PARTONLYMXNAME to configure SOLVMX to process a specific partition of a SQL/MX table. This is a new parameter.
 - e. Removing the CONSUMERFNAME from the configuration will run SOLV in a bypass mode. Use this feature to help determine if IO events into a target file or table is the cause of a bottleneck. In bypass mode, the SOLVMKRPREFIX parameter must be set because normally SOLV retrieves this information from the consumer during the handshake at startup.
- 15) A new Enscribe DDL “Flattener” and SQL table create utility, SBDDLUTL is available. The utility converts Enscribe data definition language (DDL) into a separate/new “flattened” file schema format, and supports translation into a corresponding SQL CREATE TABLE script. The SQL CREATE TABLE script can be created for any of the HPE Shadowbase-supported SQL targets, including SQL/MP, SQL/MX, Oracle, SQL Server, Sybase, MySQL, DB2, SAP Hana, and others. See the HPE NonStop Shadowbase SBDDLUTL Manual for further information.
- 16) RESTEDIT can now be used to display the values of SBCOLL-format RESTART files. To display the values of an SBCOLL-format RESTART file, the file must have

the correct file code of 180 and must not be currently opened by another process.

- 17) RESTEDIT can now be used to set the value of the RRC field (Reverse Replication Cutoff) in AUDCOLLN and SBCOLL Restart files. More information is available in the Shadowbase Tools and Utilities manual.

Problems Fixed

- 1) The AUDCOLLN object was not able to initiate a DOC roll. This problem did not apply to the SBCOLL object.
- 2) SBUNIDIR was not validating that the DBS ID specified in a SB_ADD DBS_PARAM existed or not. SBUNIDIR now checks to see if a SB_ADD DBS for the same DBS ID was made.
- 3) SBHADW was not setting the value on a SB_ADD CONS_PARAM and was generating an erroneous CONF file. SB_ADD CONS_PARAM now works as expected, and the CONF file contains the correct syntax for Consumer overrides.
- 4) The AUDCOM command INFO AUD was not displaying the license features if the license was expired and the PROD_SYSTEM feature was set, or the EXTEND_TWO_WEEKS feature is set and the user is still running within those two weeks. Now if the license is valid to run Shadowbase, it will print out the license features on a INFO AUD command.
- 5) USRXEXCEPTIONINIT was being optimized out of the Consumer by the compiler so that any custom code written in a user-exit under the USRXEXCEPTIONINIT section would not execute. USRXEXCEPTIONINIT now works.
- 6) UNDO functionality was not respecting the configured SBCOLLTURBOSIZE resulting in a Consumer abend.
- 7) SOLV Consumer was mistakenly stopping with duplicate key errors on Enscribe queue files.
- 8) Reverse replication cutoff system was being mistakenly deallocated upon release of a DBS. This would result in a segfault if the DBS was re-added and events were sent through. This deallocation now takes place when its associated Collector process is stopped.

- 9) The Shadowbase configuration script SBHADW was not setting values when a SB_ADD_CONS_PARAM parameter was added. Now Consumer parameters can be set in SBHADW just like in the other scripts, i.e. SBUNIDIR, SBBIDIR, etc.
- 10) The Consumer was mistakenly abending in certain situations where IPCs are passed through a Queue Manager to a target with a different version of Shadowbase installed.
- 11) The DBS parameter CREATEOPTIONS KEEPALTFILES was being ignored by the Consumer so that whether the parameter was set or not, the Consumer would always create target alternate-key files on replicated creates, purging as necessary. Now with this DBS parameter set, the Consumer will not create the target alternate-key files if they already exist on the target.
- 12) The SOLVMKREXTENT and SOLVMKRHSTEXTENT configuration parameters would be ignored and not get set on Shadowbase startup. This has been fixed.
- 13) SBCOLL would ping-pong data in bidirectional environments configured with AUD TRANSLOG and AUD PEERTRANSLOG. This has been fixed.
- 14) SOLVMGR would improperly send SOLV the SOLVAUDITTYPE parameter, causing SOLV and itself to abend.
- 15) SBCOLL would not notify users if no bidir restart information existed in a RESTART file if in a bidir environment.
- 16) SBCOLL would not properly wait for NOWAITed requests to the audittrail to return information before attempting to respond to AUDMON's START COLL request.
- 17) When the SBCOLL Collector encounters an ARE_DATA_FILE_VERSION error on reading an after image and TA CL PARAM SBCOLLMXSCHEMAERROK is set, then the Collector will ignore the error and not abend. Note that in NSB release v6400 a fix was made for when ARE_DATA_FILE_VERSION error was encountered on a *before image* event and SBCOLLMXSCHEMAERROK was set. Also now for both an ARE_DATA_FILE_VERSION error on a before and after image with SBCOLLMXSCHEMAERROK set, a warning message is logged to EMS.
- 18) AUDCOLL could fail when reading audit data containing audited temporary file events. This problem would occur when lots of temporary file audited events were applied on an AUX audited volume without prescreen volume checking. When a temporary file event was found adjacent to a replicated event in audit, AUDCOLL would try to pre-read the event data resulting in an ARLIB ERROR -14.

New and Modified Commands

New and Modified SET AUD Parameters

STOPPROCESSMAXRETRIES { number of retries }

| <u>Initial Version</u> | <u>Changed Version</u> | <u>Values Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|------------------------|------------------------|-----------------------|----------------|----------------|--------------|
| 6.500 | | 20 | -1 | SHORT_MAX | Yes |

This parameter controls how many times the AUDMON will retry the PROCESS_STOP_ system call when it results in an error code. The AUDMON is forced to stop processes directly when they do not respond to messages telling them to stop. The AUDMON will log an EMS message each time that the PROCESS_STOP_ system call fails. By default, the AUDMON will wait for 15 seconds before retrying the system call. This is configurable through the STOPPROCESSRETRYDELAY parameter.

An error code of 11 will not be retried, since that indicates that the process had already been stopped.

Setting this parameter to -1 will allow the AUDMON to retry the PROCESS_STOP_ system call indefinitely until the process is successfully stopped.

Setting this parameter to 0 will prevent the AUDMON from retrying the PROCESS_STOP_ system call.

STOPPROCESSRETRYDELAY { seconds }

| <u>Initial Version</u> | <u>Changed Version</u> | <u>Values Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|------------------------|------------------------|-----------------------|----------------|----------------|--------------|
| 6.500 | | 15 | 0 | SHORT_MAX | Yes |

This parameter controls how long the AUDMON will wait in between retrying the PROCESS_STOP_ system call. To disable waiting, set this parameter to 0.

New and Modified Set COLL Parameters

None in this release.

New and Modified SET CONS Parameters

SKIPWARNRATE { number of events between warnings }

| <u>Initial Version</u> | <u>Changed Version</u> | <u>Values Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|------------------------|------------------------|-----------------------|----------------|----------------|--------------|
| 6.500 | | 1 | 0 (off) | LONG_MAX | No |

This parameter controls how often the Consumer will log warnings to EMS when it encounters an event that will be skipped.

When set to 1, every skipped event will be logged.

To disable logging, set this parameter to 0.

New and Modified SET DBS Parameters

None in this release.

New and Modified SET QMGR Parameters

None in this release.

New and Modified TACL Parameters

SBCOLLNOFILEOK

| <u>Initial Version</u> | <u>Changed Version</u> | <u>Values Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|------------------------|------------------------|-----------------------|----------------|----------------|--------------|
| 6.500 | | 0 | 0 | 1 | No |

An ARE_DATA_FILE_NOT_FOUND TMFARLIB error (-14) can occur if certain ARLIB API functions are called to obtain information about an event and the file is no longer present on disk. You can have the Collector disregard these errors by entering the following TACL param before starting AUDMON so that the Collector ignores the error and continues reading the audit trail.

```
PARAM SBCOLLNOFILEOK 1
```

New and Modified SOLV Parameters

PARTONLYFNAME [\system.] [\$volume.] [subvolume.] filename

| <u>Initial Version</u> | <u>Changed Version</u> | <u>Values Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|------------------------|------------------------|-----------------------|----------------|----------------|--------------|
| 4.010 | 6.500 | | | | |

This optional parameter is used to enable a part-only load. If specified it must be one of the partitions in a partitioned file or table. The filename specified in LOADFNAME must be the primary partition of the file specified in

PARTONLYFNAME. As of version 6.500 this parameter is now valid for SQL/MP tables in addition to Enscribe files.

New and Modified SOLVMX Parameters

PARTONLYMXNAME partition-name

| <u>Initial</u> <u>Version</u> | <u>Changed</u> <u>Version</u> | <u>Values</u> | <u>Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|----------------------------------|----------------------------------|---------------|----------------|----------------|----------------|--------------|
| 6.500 | | | | | | |

Use this optional parameter to enable a part-only load. If specified, it must be one of the partition-names in a partitioned SQL/MX table. Partition-name is the SQL identifier for the partition you wish to load. Specifying an invalid partition-name to SOLVMX will cause it to display a list of all the possible names for the partitioned SQL/MX table specified in LOADMXTABLENAME.

This is an example of a system generated partition name.

PARAM PARTONLYMXNAME GRAVIC1_DATA1_ZSDLR2BK_DVCGH200

New and Modified SOLVMX or SOLV Parameters

SOLVMKRPREFIX [\system.][\$volume.][subvolume.] fname-prefix

| <u>Initial</u> <u>Version</u> | <u>Changed</u> <u>Version</u> | <u>Values</u> | <u>Default</u> | <u>Minimum</u> | <u>Maximum</u> | <u>Alter</u> |
|----------------------------------|----------------------------------|---------------|----------------|----------------|----------------|--------------|
| 6.500 | | | | | | |

Removing the CONSUMERFNAME from a SOLV configuration will cause it to run in a by-pass mode. Use this mode to help determine if IO events into the target is the cause of a bottleneck. In by-pass mode, the SOLVMKRPREFIX parameter must be set because normally SOLV retrieves the marker file information from the consumer during the handshake at startup.

For example...

PARAM SOLVMKRPREFIX \$DATA1.SOLVMXD.SOLVMKR

New or Modified EMS Messages

| | |
|------|---|
| 2026 | TACL PARAMETER SBCOLLMXSHEMAERROK SET; SKIPPING ERROR AND CONTINUING |
|------|---|

Cause: An ARE_DATA_FILE_VERSION error was encountered on either a before or after image, and the TACL PARAM SBCOLLMXSCEHMAERROK was set

Effect: The Collector will skip the error event without stopping

Recovery: This is a warning message. There is nothing to recover.

| | |
|------|---|
| 2017 | SQL STATEMENT SIZE (<size-in-bytes>) EXCEEDS MAX SIZE (<size-in-bytes>). EVENT LOCATED AT ADT <audit-trail- file> RBA <rba> |
|------|---|

Cause: The SQL statement generated for the event being replicated to an Other Servers target was too big to send over the network.

Effect: The Consumer abends.

Recovery: Contact Shadowbase Support to recover from this issue.

| | |
|------|---|
| 2119 | SBMP: COULD NOT RETRIEVE A SOURCE ROW/RECORD FOR <table name> (KEY COLUMNS: <key columns>) |
|------|---|

Cause: The source row for the compressed update could not be found because the row may have been deleted before the update occurred, or in the case that a MAPTOFILE is being used, the schema of the SOURCEFILE and MAPTOFILE do not match.

Effect: The result of this error depends on the configuration of the respective DBS. The default is to skip the update event.

Recovery: In the case of the schemas being mismatched between the SOURCEFILE and MAPTOFILE tables, either fix the definitions so that they match, or use user exits to add any additional logic. In the case that the row does not exist, i.e. it was deleted before the compressed update, reconfigure the respective DBS so that it behaves as desired (e.g. turn INSERTNOTFOUND ON for the DBS, and make sure audit-compression is off if the desired behavior is to insert into the source table).

| | |
|------|---|
| 2017 | SQL STATEMENT <INSERT/SET/WHERE/DATA> VALUE HAS A 0 LENGTH -> COLUMN=<column name>, DBS OBJECT=<DBS object>; ERROR MAY BE CAUSED BY INCORRECT MAPTOFILE SCHEMA, OLD AUDIT FOR A SOURCEFILE WITH CHANGED SCHEMA, OR NULLS IN EVENT DATA; IF ERROR IS NULLS IN EVENT DATA IN ENSCRIBE SOURCE, THEN FIX THE ERROR BY SETTING THE DBS PARAMETER 'FIXUPDECIMALS' |
|------|---|

Cause: The consumer outputs this message to indicate that a problem occurred trying to format a column set value in an SQL statement. The column name and related DBS object name are included to assist in troubleshooting the source of the problem.

Effect: The Consumer abends.

Recovery: Check the compatibility of the columns in your source and target. Verify the correctness of your user exit code if applicable. It is possible that the incoming data for the column needs is a binary zero or other unknown value, and this data can be scrubbed in a user exit.
If the incoming data is a binary zero in an Enscribe source file replicating out to HPE Shadowbase Other Servers, use the DBS parameter 'FIXUPDECIMAL'.
Make the appropriate corrections and restart.

| | |
|------|---|
| 2017 | SQL STATEMENT SIZE (<event-statement-size>) EXCEEDS MAX SIZE (<max-statement-size>). EVENT LOCATED AT ADT <audit-prefix><audit-sequence-number> RBA <rba> |
|------|---|

Cause: The length of the SQL statement of the current event exceeds the max length for a SQL statement being sent over TCP/IP.

Effect: The Consumer abandons.

Recovery: Contact Gravic Support.

2026 COLL <collector name>: AUDIT READ ERROR (MSG 1 OF 3)
ON CALL TO <name of ARLIB system call> ERROR CODE
<ARLIB error code>, FILE <filename>, POSITION <audit
position> (TEXT FOLLOWS)

COLL <collector name>: AUDIT READ ERROR (MSG 2 OF 3):
arlib error ARE_INTERNAL_ERROR(-1000) on call to <name of
ARLIB system call> (Error <error name> (<error code>) in
Subsystem <subsystem name> (<subsystem code>))

COLL <collector name>: AUDIT READ ERROR (MSG 3 OF 3):
<additional information may follow>

Cause: TMF ARLIB has encountered an internal error.

Effect: The Collector will stop.

Recovery: See the TMF Application Programmer's Guide for an explanation
of the given error codes.

2108 WARNING: CONSUMER SKIPPING <event type> EVENT
FROM SOURCE FILE=<source filename> (DBS=<db name>).
ADT POS=<audit position>

Cause: The Consumer encountered an event that will be skipped. Events
skipped via user exit processing will not result in warnings.

Effect: The Consumer will skip the event and continue processing.

Recovery: Determine whether or not this event was intended to be skipped,
and investigate whether this event could be skipped at the
Collector or Queue Manager level.

2014 UNABLE TO KILL PROCESS <process name>:
STOP_PROCESS_FAILED WITH ERROR CODE <code>

- Cause:** The AUDMON's attempt to stop the process via the PROCESS_STOP_ system call failed with the given error code.
- Effect:** Depending on your configured STOPPROCESSMAXRETRIES parameter, the AUDMON will retry the system call or continue under the assumption that the process was stopped.
- Recovery:** Ensure that the listed process was actually stopped before starting the listed process again.

| | |
|------|---|
| 2014 | UNABLE TO KILL PROCESS <process name> AFTER <number> RETRIES. PROCESS WAS NOT STOPPED |
|------|---|

- Cause:** The AUDMON's attempt to stop the process via the PROCESS_STOP_ system call failed and the configured STOPPROCESSMAXRETRIES parameter was such that the AUDMON will no longer retry the system call.
- Effect:** The AUDMON will continue under the assumption that the process was stopped.
- Recovery:** Ensure that the listed process was stopped before starting the listed process again.

| | |
|------|--|
| 2013 | PROCESS <process name> ALREADY EXISTS. STOP THE EXISTING PROCESS BEFORE ATTEMPTING CREATION OF <object logical name> PROCESS |
|------|--|

- Cause:** The AUDMON attempted to start a process with the same name as an existing process. This can occur if Shadowbase did not properly stop all objects the last time the environment was stopped.
- Effect:** The object will not start.
- Recovery:** Ensure that the listed process is stopped before trying to start the object again.

New and Modified User Information Messages

The new user information messages are:

| | |
|---|--|
| 4 | ALTERING REPLICATION PARAMETER WITH ACTIVE TRANSACTIONS IS NOT ALLOWED. REPLICATION SETTING NOT CHANGED. |
|---|--|

Cause: The REPLICATION parameter's ACTION component was modified as part of an ALTER COLL command while one or more transactions are being tracked by the Collector. This is disallowed to prevent situations in which changing the REPLICATION ACTION while a transaction is in progress can result in hung transactions.

Effect: The REPLICATION parameter is unchanged.

Recovery: Ensure all active transactions are completed before retrying the ALTER command.

| | |
|------|--|
| 2129 | COLL USING REPLICATION SETTING <restart file setting> FROM RESTART FILE INSTEAD OF CONFIGURED SETTING <configured setting> |
|------|--|

Cause: The Collector's REPLICATION parameter's ACTION component was modified using the ALTER command to something other than its configured value. This change is persisted using the configured RESTART file to avoid problems associated with changing the REPLICATION ACTION while a transaction is in progress.

Effect: The Collector will continue using <restart file setting>.

Recovery: If <restart file setting> is not desired, the REPLICATION parameter can be modified using the ALTER command at a time when there are no outstanding transactions that the Collector is tracking.

New or Modified SOLVCOM User Messages

New or Modified SOLV User/EMS Messages

The new user messages are:

| | |
|--------|---|
| 002084 | SOLV: CONSUMER HAS <n> UNACKED MSGS (OLDEST SENT AT <time>) CHECK AUDIT LATENCY OR SUSPENDED PROCESSES. |
|--------|---|

Cause: The consumer has not acknowledged an IPC message in a 15-minute timeout period.

Effect: Informational message only that could indefinitely repeat every 15 minutes.

Recovery: Check audit latency or look for a suspended consumer process. Use CONSUMERIPCTIMEOUT parameter to adjust the timeout period and frequency of the reported message.

| | |
|--------|--|
| 002084 | SOLV: UNUSED PARAM(S): <parameter-name-list> |
|--------|--|

Cause: SOLV identified an unused parameter name(s).

Effect: None unless a required configuration parameter for SOLV or SOLVMX was mistyped.

Recovery: Check the spelling of all parameter names in the list and correct if necessary.

| | |
|--------|---|
| 002084 | SOLV: ADD MISSING CONSUMERFNAME PARAM OR FOR BY-PASS MODE ADD SOLVMKRPREFIX PARAM |
|--------|---|

Cause: A consumer process name was not specified.

Effect: SOLV does not run.

Recovery: If this was intentional because you wish to run SOLV in BY-PASS MODE, then the SOLVMKRPREFIX param must be specified. If this is not the case, then supply the CONSUMERFNAME parameter with valid consume process name.

| | |
|--------|---|
| 002084 | SOLV: RUNNING IN BY-PASSMODE, USING %s AS SOLVMKRPREFIX |
|--------|---|

Cause: Informational message only.

Effect: SOLV does not send load events to the consumer. Use this mode to help identify bottlenecks in data loading.

Recovery: None.

| | |
|--------|--|
| 002084 | SOLV: CANNOT USE PARTONLYFNAME FOR MX, USE PARTONLYMXNAME. |
|--------|--|

Cause: The PARTONLYFNAME parameter was specified to SOLVMX for loading a SQL/MX table.

Effect: SOLVMX does not run.

Recovery: Use PARTONLYMXNAME instead for PARTONLYFNAME.

| | |
|--------|--|
| 002084 | SOLV: <PARTONLYMXNAME-value> MUST BE ONE OF THE <number-of-partitions> PARTITION NAMES: <partition-name-list>" |
|--------|--|

Cause: The value specified for the PARTONLYMXNAME parameter specified to SOLVMX was invalid.

Effect: SOLVMX does not run.

Recovery: the value for <PARTONLYMXNAME> must be one of the partition names in the <partition-name-list>.

| | |
|--------|---|
| 002084 | SOLV: PARTONLYMXNAME SET TO <partition-name> PART#<part-number> |
|--------|---|

Cause: Informational message only. <Part-number> identifies to which partition-number the partition-name mapped.

Effect: None

Recovery: None.

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

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.

4) If upgrading from an older version with an existing Restart file, it is necessary to run the RESTEDIT program to upgrade this Restart file for use with the 6.500 environment. Before upgrading your Restart file, it is recommended to create a backup copy (e.g. via FUP DUP) of the existing restart file. Upgrading the restart file can be done using the following syntax:

```
RESTEDIT <configured Restart file> -SET-RRC=<desired setting>
```

To use the default REPLICATION parameter setting for the configured Collector, use (ON, NOWARN, INTERVAL -1) as the desired setting in the above command.

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) 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) 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.
- 29) 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).
- 30) 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.
- 31) 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.
- 32) 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.
- 33) Consumer processes that are on a different node from their corresponding AUDMON process are not automatically killed when a non-responsive state is detected.
- 34) AUDCOLL processes do not currently handle NEXTDOC commands.
- 35) The HADW script enables the SBCOLLAUTOADVDMODE 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.

- 36) The Consumer's I/O Tracing output can contain garbage characters when the Consumer encounters an error 73 (table or record is locked).
- 37) 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.
- 38) 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.
- 39) The DBS ALLPARTITIONS parameter does not work with the SBCOLL Collector program.
- 40) The Enscribe compare program, SOLVNV, does not compare entry sequenced or relative files correctly. Incorrect counts of matches and mismatches are returned.
- 41) 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.
- 42) 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.
- 43) The Enscribe compare program, SOLVNV, traps with an illegal address reference when attempting to compare unstructured files.
- 44) 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.
- 45) Internal Gravic testing revealed SOLV traps when running in VnV mode for Enscribe Compare. The issue occurs when SOLV attempts to create its marker history file. There are two workarounds to this issue:
 - Temporarily run a Shadowbase environment configured with a SOLVONLY Consumer, and use the marker history file generated by the consumer
 - Manually create the SOLV marker history file using FUP, see commands below (extents can be set as desired):

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

- 46) 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 in #46 above.
- 47) Using the Collector REPLICATION parameter's SKIP action component can result in the Collector mistakenly holding buffers until it has no more buffers to allocate. At this point, the Collector will hang.
- 48) 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.

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