SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product: Shadowbase

Release: Version 6.100 (**full NonStop release**)

Release Date: August 8, 2014

Copyright Notice: Copyright Gravic, Inc. 1995 – 2014 (www.gravic.com)

File Name: IPM6100.doc

NOTE: **Release 6.100 as documented in this softdoc is a full NonStop release of Shadowbase and contains a complete set of the product files.**

NOTE: This softdoc covers new features and corrected problems for Shadowbase for G06, H06, and J06 NonStop operating system versions. This softdoc is available in an Adobe PDF file (.PDF). Softdoc files are named IPMnnnnx (where nnnnx is the Shadowbase version number).

Table of Contents

Special Notes for Version 6.100

1) Due to a known problem outstanding, you should not configure a QMGR process for a Collector if it has one or more DBS's that are processing DDL operations with coordination and one or more Consumers doing TRACKTX processing. TIDFILECONSTANT ON may be used in place of TRACKTX processing.

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. Additionally, if you have a DBS object that uses MAPTOFILE, you should set the SOLV TARGETFNAME parameter to the DBS MAPTOFILE value.

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 6.100 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 6.100. Note that usrxlibn is available for Guardian O/S G0x, and H0x releases of Shadowbase and usrxlibx is only available for Guardian O/S H0x releases of Shadowbase at this time.

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 float. 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 SYSKEY either as the primary key or in a clustering key. Shadowbase cannot set the value of a SYSKEY column when inserting a new row. As a result, the target system will have a different SYSKEY value from the source. Additional programming in the form of user exits is required to reliably replicate data with SYSKEYs. Contact Shadowbase Support if you are contemplating replicating tables with SYSKEYs defined.

- 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, 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, 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.

As noted above, large row support is limited to approximately 28K for replication where the before image is not required, and 14K for rows where before images are required. Contact Shadowbase Support if you are contemplating replicating tables with rows over 28K (no before images) or 14K (before images required).

13) Shadowbase NonStop Version 6.100 is only supported on G06.24 and later releases for NonStop G06 Systems, and H06.03 and later releases for NonStop Integrity Systems.

14) The Shadowbase EMS subsystem owner and version have been updated (the EMS SSID token has changed in this release). 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 to V61. 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. See SBEMSOLDSSID TACL PARAMETER for more details.

15) The Shadowbase NonStop release is now included in a single self-extracting file, SBPAK. The files that were previously included separately in the self-extracting files SBSLPAK and SBLDPAK are now included in SBPAK. See INSTALLATION INSTRUCTIONS (SHADOWBASE NONSTOP) for updated installation instructions.

Changes in Release 6.100 (since Version 5.010E)

New Features

1) The edit files readme and license have been added to the Shadowbase release files. The readme file contains a number of copyright statements. The license file contains the software license agreement text.

2) The Shadowbase password will now be verified for bi-directional replication (BIDIR) if it is configured.

3) Audit-compressed update events with no change data will now be passed through to the CONS for processing if the DBS FETCHSOURCE parameter is set to ON.

4) The start banners for interactive programs, including AUDCOM, AUDMON, SOLV, SOLVCOM, and RESTEDIT has changed to read:

```
<name> - T1122 - V6100J06 - (08AUG14), COPYRIGHT GRAVIC, INC.
1995-2014. PORTIONS COPYRIGHTED BY AND LICENSED FROM THIRD PARTIES. SEE
README FILE. USAGE SUBJECT TO THE TERMS OF A WRITTEN LICENSE AGREEMENT.
PATENTS: SEE WWW.GRAVIC.COM/GRAVICLABS/PATENTS/PRODUCTS.
```

where <name> is the name of the program.

Problems Fixed

1) The datetime and interval qualifier strings were not being generated for the WHERE clause by the Consumer related to KREP/LOADER processing. This caused the Consumer to abend.

2) Related to user exit processing and replication out to a non-NonStop system, the Consumer had a memory leak in logic previously added to generate consistently-ordered column names in SQL statements.

3) Related to a NonStop-to-NonStop replication environment where the TRACKTXFILE was configured, a number of I/O events were missed following a restart.

4) The Consumer trapped on an illegal address reference during RESOLVETRANS processing when IOTRACE was set to 6.

5) A number of compatibility traps were eliminated from the Collector.

6) Related to DDL create processing involving partitioned sources, events got skipped and not applied to the target if the first create event received was not for the primary partition.

7) SQL DDL creates and alters were ping-ponging due to flushing order of the events.

8) When the ROWORRECORD option was specified as the consistency level while running SOLVMX, the load would abnormally abend before completing the load. The problem has been corrected.

9) A long running SOLVMX load would sometime display a PROCESS_DELETE system message when the SQL/MX compiler (mxcmp) program timed out. It is normal for mxcmp to eventually time out as it is no longer needed by the Shadowbase loader. The message was enhanced to include the process handle of the mxcmp program.

10) Corrected a problem in SOLV loading of SQL/MP tables that had multiple columns that were defined as varchar with a length greater than 255 bytes.

11) Corrected a problem in regards to replicating creates of SQL/MP tables that caused a memory trap if the table had only one partition.

New Features Description

Software Licensing Changes

As a result of the Gravic/Shadowbase HP global reseller deal, the Shadowbase license processing has been reworked. All previously released Shadowbase license files (SHADPASS and shadpass.ini) are now obsolete.

Please contact Gravic support if you are upgrading from a release prior to this version 6.100 as you will require a new Shadowbase license file.

New and Modified Commands
No commands were added or modified since version 5.010E.

New and Modified SET Parameters
No SET parameters have been modified since version 5.010E.

New and Modified TACL Parameters

SBEMSOLDSSID TACL Parameter

Initial Version	Changed Version	Values Default	Minimum	Maximum	Alter
6.100		0	0	1	No

By default, for the EMS subsystem identifier (SSID), Shadowbase will use GRAVIC as the subsystem owner, and VNN for the subsystem version, where NN are the two most significant digits of the Shadowbase release version. For example, for Shadowbase release version 6.100, the subsystem version is V61. This TACL parameter can be set to 1 to enable use of the old Shadowbase subsystem owner and version. In this case, Shadowbase will use ITI as the subsystem owner, and D42 as the subsystem version. Note that the subsystem owner and version are used when logging messages to EMS, and the values used will impact EMS filters if the subsystem owner and/or version are being used for filtering.

New or Modified EMS Messages

No EMS messages have been modified since version 5.010E.

New and Modified User Information Messages

No user information messages have been modified since version 5.010E.

New or Modified SOLVCOM User Messages

No SOLVCOM messages have been modified since version 5.010E.

Installation Instructions (Shadowbase NonStop)

Shadowbase is now built for HP NonStop installations using the HP 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 **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 non-native consumer, AUDCONS, may hang on a restart if the NETSWAPCONNECT parameter is enabled. The native consumer should be used in this case.

2) The SQL/MX collector, SBCOLL, will fail with an ARLIB ARE_DATA_FILE_VERSION error message if a SQL/MX source table is restored with audit events still to be processed in the audit trail from prior to the restore point. In this case, the table must be dropped from replication until there are no more events to be processed in the audit trail from the version of the table prior to the restore.

3) SQL/MX will not replicate rows from tables with reserved words as column names without a user exit.

4) The consumer may trap during cleanup of its in-memory tables if FETCHSOURCE is enabled for partitioned files. In this case, the consumer will restart successful and continue processing. The cleanup processing can be disabled by setting the CONS parameter FETCHSRCCLNUPTM to 0.

5) 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.

6) In rare cases, a File System error 1061 (THE CURSOR IS NO LONGER DEFINED IN THE DISK PROCESS) on an SQL/MP table that results in an -8204 SQL error (AN EXISTING OPEN FAILED ON THE TABLE <table>) may result in the operation being retried twice. If the operation was an insert against an entry-sequenced table, a duplicate row may get added. This error is under investigation.

7) For SQL/MX, the documented maximum value for the 32-bit float is 3.40282347e+38. However, the Shadowbase Open ODBC Server gets a range error starting with 3.40282337e+38 (difference highlighted and underlined). The cause of this issue remains under investigation.

8) For SQL/MX, Shadowbase can handle a maximum row size of about 28K bytes, half that if BEFOREVALUES are being sent to the Consumer for user exit purposes. Note that SQL/MX 3.0 and above can handle a maximum row size of 32K bytes.

9) 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.

10) 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.

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

12) 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.

13) 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.

End of Document