

SOFTWARE RELEASE DOCUMENT (SOFTDOC)

Product:	Shadowbase
Release:	Gravic Version 6.200 HP NonStop Shadowbase: T1122-AAB (SB Repl/Guardian)
Release Date:	March 16, 2015
Copyright Notice:	Copyright Gravic, Inc. 1995 – 2015 (www.gravic.com)
File Name:	IPM6200.pdf

NOTE: Release 6.200 is a release of HP NonStop Shadowbase for Guardian for the new HP NonStop x86 product line running Guardian OS version L15.02 or later. It includes a complete set of product files for all HP NonStop components for the new architecture. This softdoc applies to the HP NonStop Guardian replication components. Other softdocs document the releases for the HP NonStop Shadowbase Audit Reader, HP NonStop Compare, and the HP NonStop OSS Replication components.

NOTE: This softdoc covers new features and corrected problems for HP NonStop Shadowbase for Guardian, Gravic Version 6.200 (HP TID T1122-AAB). This softdoc is available in an Adobe PDF file (.PDF). Softdoc files are named IPMnnnn.pdf (where nnnn is the Gravic version number).

Table of Contents

SOFTWARE RELEASE DOCUMENT (SOFTDOC).....	1
Table of Contents	2
Special Notes for Version 6.200	3
Changes in Release 6.200 (since Version 6.101).....	5
New Features	5
Problems Fixed	5
New and Modified Commands	5
New and Modified SET Parameters	5
New and Modified TACL Parameters	6
New or Modified EMS Messages	6
New and Modified User Information Messages	6
New or Modified SOLVCOM User Messages	6
Installation Instructions (Shadowbase NonStop).....	6
Known Problems Remaining	8

Special Notes for Version 6.200

- 1) **Important:** The non-native Consumer (AUDCONS, AUDCONSR) and related files are not included in this distribution for the L15 version of HP NonStop Guardian OS. They will continue to be supported for the G06 and H06 OS versions in future releases of HP NonStop Shadowbase Replication for a limited time. However, for the Shadowbase x86 releases, non-native Consumers have entered LIMITED Support and they will attain OBSOLETE Support in a future release. Your use of them should be replaced by using native Consumers (AUDCONSN) at your earliest convenience. Contact Shadowbase Support for more information.
- 2) Shadowbase NonStop Version 6.200 is only supported on L15.02 and later releases for HP NonStop x86 Systems.
- 3) 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.
- 4) 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.
- 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 Release 6.200 USRXLIBN or USRXLIBX API library included in the distribution subvolume.
- 7) 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 (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.

8) 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.

9) 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.

10) 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.

11) 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. (Note: A future version of Shadowbase will allow Shadowbase to assign the target SYSKEY value to match the source SYSKEY value. Contact Shadowbase Support for more information.)
- 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 TMF's 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 the child row due to the RI constraint being violated on replay.
- Tables with large (long) 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 event of the revised row. Shadowbase will replay this sequence 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).

12) The Shadowbase NonStop release is now included in a single self-extracting file, SBPAK. See [INSTALLATION INSTRUCTIONS \(SHADOWBASE NONSTOP\)](#) for installation instructions.

Changes in Release 6.200 (since Version 6.101)

New Features

- 1) Support for the new HP NonStop x86 architecture has been added. All natively compiled objects and executables have been rebuilt targeting the new x86 architecture. The compile, link, and installation scripts have been updated to use the new x86-based development tools. Basic Shadowbase operations remain the same across the H06 and L15 operating systems.

Problems Fixed

Note: This release only supports the L15.02 and later operating systems. As a result, these problems are corrected only for HP NonStop x86 systems only with this release. An SPR will be issued shortly following this release that incorporate the fixes in the H06 version of Shadowbase.

- 1) The collector (COLL) abends if a DBS is dynamically added and captured on a thread that uses the queue manager (QMGR). If restarts are enabled, the COLL will automatically restart with the DBS captured. This has now been corrected: a DBS can be captured and added on a thread using a QMGR with no issues.
- 2) The COLL was not correctly processing to a multi-ported DOC in configurations that replicated data to Other Server targets that also used one or more QMGRs. This could cause the data not to be applied correctly on the target system. This has been corrected.
- 3) One of the COLL's measure counters, ADTEFLDUPDBUSY (a BUSY counter for measuring time spend processing compressed updates), was incorrectly named in the code.

New and Modified Commands

No commands were added or modified since version 6.101.

New and Modified SET Parameters

No SET parameters were modified or added since version 6.101.

New and Modified TACL Parameters

No TACL parameters were modified or added since version 6.101.

New or Modified EMS Messages

No EMS Messages were modified or added since version 6.101.

New and Modified User Information Messages

No User Information Messages were modified or added since version 6.101.

New or Modified SOLVCOM User Messages

There are no new or modified SOLVCOM User Messages in this release.

Installation Instructions (Shadowbase NonStop)

Shadowbase is now built for HP NonStop installations using the HP PAK and UNPAK utilities. 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 L-based systems:

FUP ALTER SBPAK, CODE 500

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 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 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.
- 3) 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.
- 4) With a NonStop AUDMON configured for auto-restart and backup CPUs are configured for the Collectors and Consumers, you may have to issue an AUDCOM RUN command for replication to continue after a CPU failure if you have previously entered a SUSPEND COLL command followed by a RESUME COLL command (prior to the failover of AUDMON). A workaround for this is to enter an AUDCOM run command immediately after the RESUME COLL command.
- 5) During SOLV testing at Gravic, the consumer received an SQL error -1031 for an insert to one of its complex test tables. HP's SQL error description for the error is:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

9) SOLV has the potential of missing records at the end of a file in a sequence when used with SOLVMGR (solv manager) for 'file chasing'. This problem is more likely to

occur with low values of SOLVMGRDELAY, SOLVEOFDELAY, INACTIVITYDELAY and ACTIVITYDELAY are set. By default these values are set to 60 seconds. The minimum values for these configurations are 1 second, however, Gravic does not recommend using values lower than the default to avoid the potential for data loss.

10) SOLV has the potential of generating lots of locks and hitting an error 35 on a source SQL/MP table with mixed ascending and descending primary key columns.

SOLV is hitting a condition where SQL/MX is generating more locks than the volume can handle. The lock limit is set by the value for MAXLOCKSPERTCB of the disk volume. This problem was encountered when SOLV was performing an online load with SOLVCONSISTENCYLEVEL set to TRANASCTION. This also will only occur when there are more than MAXLOCKSPERTCB rows in the source table that have the same value in the first (most significant) column of the primary key.

Work arounds include...

- Perform an 'offline' load by setting SOLVCONSISTENCYLEVEL to ROWORRECORD (will prevent dirty reads) or NONE.
- If the table has an index, specify one of the indexes in the SQLINDEXFNAME configuration parameter. Note that this option will significantly slow down the load.

***** End of Document *****