

HPE Shadowbase Mapping Facility (SBMAP)



An SQL-like Scripting Facility for Transforming Large Data Sets

Transforming data is critical in today's IT landscape. To remain competitive, enterprises must collect, extract, and analyze real-time data in various parts of the organization simultaneously. Transforming large data sets is daunting, whether for migration, real-time analysis, database administration, or another reason.

HPE Shadowbase User Exits have long been a feature to embed custom logic into the replication engine. Prior to SBMAP, all data mapping logic had to be manually programmed as a User Exit. Now, SBMAP automatically transforms in-flight data as it is replicated and defines data mapping logic without the need to write custom C, C++, or COBOL Shadowbase User Exits, by using a familiar and intuitive SQL-like syntax. Designed to work and feel like SQL/MP, SBMAP makes transforming large data sets easier than ever before. SBMAP is sold as part of the HPE Shadowbase Essentials Software Bundle.



Architecture

SBMAP is an extra facility bound into the standard HPE NonStop Shadowbase Consumer that rests downstream in a typical replication environment. Figure 1 depicts a standard uni-directional HPE NonStop Shadowbase replication environment, and shows how SBMAP works with the overall architecture.

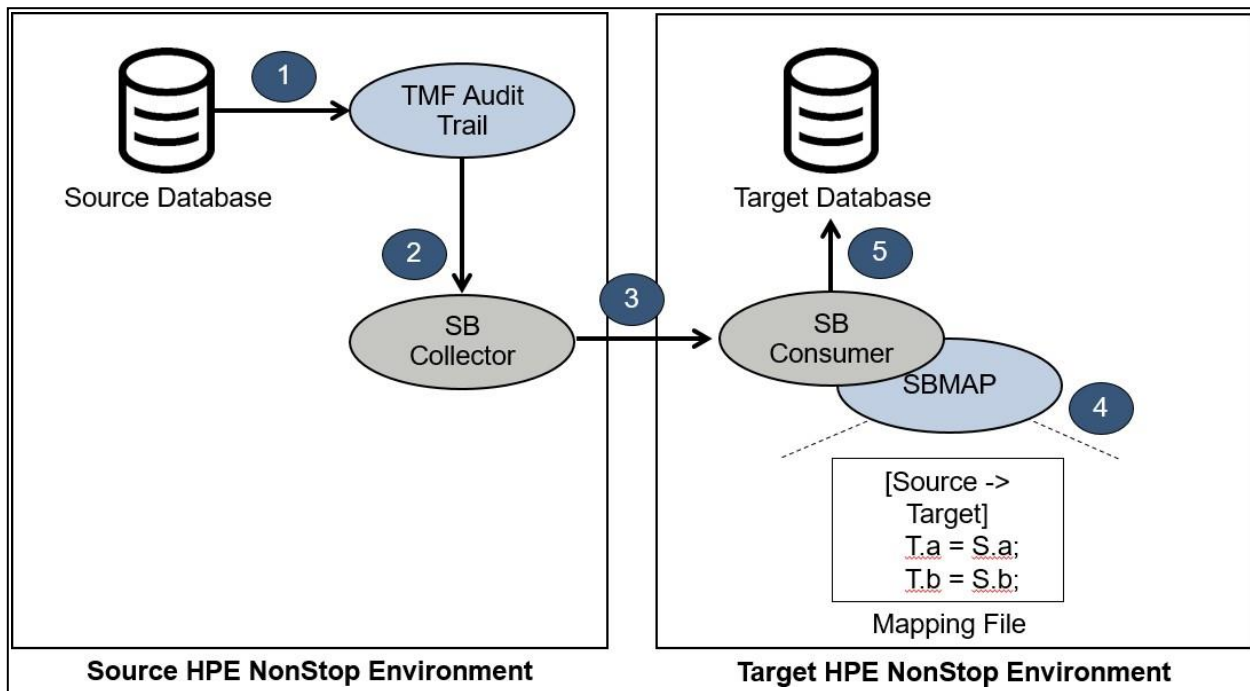


Figure 1 – SBMAP is Bound into the Shadowbase Consumer on the Target System

HPE NonStop Shadowbase Replication Event Steps:

1. A DML event on the source database creates a written record in the TMF audit trail.
2. The upstream Shadowbase Collector picks up and processes this record.
3. DBS settings and other parameters are checked in case the event needs to be sent to the downstream Consumer, which can reside on the same or different node and reached via EXPAND or TCP/IP.
4. At startup, the Consumer processes the SBMAP input file. The SBMAP facility checks the event's source and target tables for any map blocks (contained within the mapping input file) that match the current event.
 - a. If so, SBMAP attempts to apply the mappings and transformations defined by the user.
 - b. Otherwise, the event is replicated without changes.
5. The event is applied to the target database (or sent to a downstream Shadowbase for Other Servers DOC or Direct Writer, if configured to replicate off-platform).

Order of Process Execution

SBMAP and other User Exit-like facilities within the Consumer share a similar point-in-time regarding processing replication events. Utilizing these processes is optional, each process option is independent from the others, but are simultaneously enabled via DBS and CONS parameters. Note, the order of processing may matter.

Figure 2 demonstrates the SBMAP Facility's relative processing order among other processes in a typical HPE NonStop environment. It is possible to write and bind a custom User Exit into the Consumer while also enabling SBMAP processing via DBS and CONS parameters. However, note that SBMAP processing will always trigger prior to User Exit processing, if both are enabled, which may affect how the User Exit is written.

Data Transformation Features

SBMAP flexibly allows the user to choose from multiple commands and operations to ensure that data is transformed and delivered in the proper format.

1. **Add, Drop, Rename Columns, and Apply Event Filtering**
 - a. Set the values of new columns
 - b. Use the SBWHERE statement (a conditional trigger that specifies when SBMAP will execute and what data it will replicate), similar to "WHERE" in standard SQL
2. **Use Built-in Functions (for SQL/MP)**
 - a. Arithmetic operations
 - b. Build "IF...THEN...ELSE" and "SWITCH...CASE" statements
 - c. CAST function (for casting between common SQL/MP datatypes and includes date-time support)
 - d. Numeric manipulation
3. **Use String Processing**
 - a. Concatenate column data
 - b. Create custom strings
 - c. String-based functions
 - i. POSITION
 - ii. SUBSTRING
 - iii. CHAR_LENGTH
 - iv. BYTE_LENGTH
 - v. TRIM
 - vi. UPSHIFT
 - vii. And more
4. **Reference and Utilize External Tables**
 - a. An external table is not a part of the Database Specification (DBS) where the current mapping event belongs, and is like a "third-party" table (neither the source nor the target in the replication architecture)
 - b. Enable event- and streaming-types of operations to occur based on values in tables on other platforms, environments, and systems
 - i. Conditional formatting (based on external table data values)
 - ii. Extract and read data
 - iii. Perform DML operations
5. **Replicate Before and After Image Values**
 - a. Simplify analyzing data that was changed, particularly for active/active architectures
 - b. Enhance data verification and debugging
6. **Support Parallelism**
 - a. Allow for even more fine-grained replication threading by hashing and replicating individual partitions of a table across separate Shadowbase Consumer threads

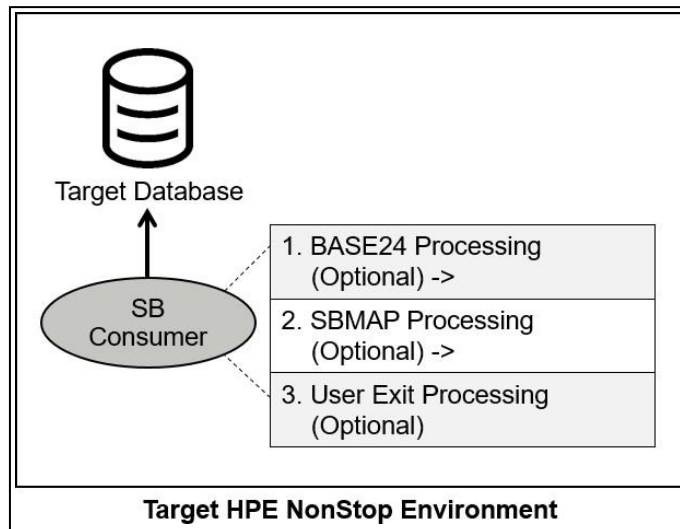


Figure 2 – SBMAP Processing in Relation to Other Processes

SQL/MP Data Types Supported by SBMAP

HPE NonStop Shadowbase replication via SBMAP supports the following SQL/MP data types:

- CHAR and CHAR UPSHIFTed
- VARCHAR and VARCHAR UPSHIFTed
- SMALLINT and SMALLINT UNSIGNED
- INT and INT UNSIGNED

- LARGEINT
- REAL and DOUBLE PRECISION
- DECIMAL UNSIGNED
- DECIMAL LEADING SIGN EMBEDDED
- DATE, TIME, and DATETIME (any range)

The following SQL/MP data types are *not supported* at this time:

- INTERVAL
- DECIMAL LEADING SIGN SEPARATE
- DECIMAL TRAILING SIGN SEPARATE
- DECIMAL TRAILING SIGN EMBEDDED

Note: SBMAP supports replication and mapping of strings containing binary data, including NULL characters.

Data Literals Supported by SBMAP

Data literal	examples
INTEGERS	-1234, 1234567, 0, -15000000
FLOATING POINT NUMBERS	123.456, -0.5, -3.1415926
STRING LITERALS	"Hello, world", 'This is a string'
DATE LITERALS WITH	DATE "1999-12-31", DATE '2002-08-19'
TIME LITERALS WITH	TIME "12:24:05", TIME '10:02:12'
TIMESTAMP LITERALS	TIMESTAMP "1990-01-22:13:40:05.55", TIMESTAMP '1995-04-03:15:43:12.635'
DATETIME LITERALS	DATETIME "2000-01-01:12:00:00.123456" YEAR TO FRACTION(6)
CURRENT	CURRENT
HEX LITERAL	X'00FF', x'12AB'

Note: The HEX LITERAL is an SBMAP-specific concept and is not supported by SQL/MP.

Operators Supported by SBMAP

Operator	Examples
+	T.VAL = S.VAL + 123;
-	T.VAL = 100 - S.VAL;
*	T.VAL = S.VAL * 2;
**	T.VAL = S.VAL ** 3;
/	T.VAL = 15 / 3;
 	T.VAL = "Foo:" "Bar:" "Baz";
AND	S.VAL > 5 AND S.VAL < 10;
OR	S.VAL = 5 OR S.VAL = 15;
NOT	SBWHERE = NOT (S.VAL = 5);
=	SBWHERE = S.VAL = 5;
<	SBWHERE = S.VAL < 5;
>	SBWHERE = S.VAL > 5;
<=	SBWHERE = S.VAL <= 5;
>=	SBWHERE = S.VAL >= 5;
<>	SBWHERE = S.VAL <> 5;
()	SBWHERE = (S.VAL > 5) AND (S.VAL < 10);

CASE Expressions Supported by SBMAP

1. **IF-CASE** allows the user to assign values based on the results of conditional expressions and behaves like SQL/MP if-statements.
2. **SWITCH-CASE** works more like a C-style switch/case statement, where a single value or the result of an expression is checked against a series of possible options and behaves like C-style switch-case-statements.

IS Expressions Supported by SBMAP

IS Expressions are conditional statements that determine a statement about its argument.

1. **IS [NOT] NULL**
2. **IS [NOT] EMPTY**
3. **IS [NOT] LIST** evaluates whether a given SBEXEC SELECT statement returned multiple rows when processed.
4. **IS [NOT] MISSING** allows the user to determine whether a source column exists within the current replication event.
5. **IS [NOT] VALID** determines if the value of a source column in the audit trail record is "valid" for the given target column data type and size for where it is being replicated.

Summary

Transforming data is critical in today's IT landscape. Shadowbase User Exits have long been used to embed custom logic into the replication engine. Prior to SBMAP, all data mapping logic had to be manually programmed as a User Exit. The HPE Shadowbase Mapping Facility (SBMAP) flexibly allows the user to choose from multiple commands and operations to ensure that data is transformed and delivered in the proper format. Use SBMAP to automatically transform in-flight data as it is replicated, and define datamapping logic without the need to write custom C, C++, or COBOL Shadowbase User Exits, by using a familiar and intuitive SQL-like syntax.

Hewlett Packard Enterprise globally sells and supports Shadowbase solutions under the name HPE Shadowbase. For more information, please contact your local HPE Shadowbase representative or visit our website. For additional information, please view our Shadowbase solution videos: <https://vimeo.com/shadowbasesoftware>.

Learn more:

shadowbasesoftware.com
hpe.com

Contact us:

Gravic, Inc.
17 General Warren Blvd Malvern, PA 19355-1245 USA
Tel: +1.610.647.6250
Fax: +1.610.647.7958
Email Sales: shadowbase@gravic.com
Email Support: sbsupport@gravic.com

Please follow:

