




**Hewlett Packard**  
Enterprise

Data Sheet

# HPE Shadowbase Suite for HPE NonStop and Other Servers

Real-time data replication and integration  
for a nonstop world

Mission-critical enterprise users rely on the award-winning HPE Shadowbase product suite to address all aspects of real-time data replication. HPE Shadowbase is HPE's strategic *go-forward* data replication solution for HPE NonStop and other servers.



## HPE Shadowbase Business Continuity software

### Avoid business losses by eliminating planned and unplanned downtime

Today's enterprises need to maintain continuous operations, which requires minimizing downtime – both planned and unplanned. HPE Shadowbase Business Continuity software provides low-latency, uni-directional, and bi-directional data replication between homogeneous and heterogeneous systems and databases with scalability, selectivity, and sophisticated in-flight data transformation and mapping solutions.

#### Key Business Continuity Highlights:

- Eliminate planned downtime
- Eliminate unplanned downtime
- Achieve continuous application availability
- Perform zero downtime migrations, upgrades, and platform refreshes



HPE Shadowbase Business Continuity software provides a solution to meet these key business availability requirements:

1. Uni-directional active/passive disaster recovery architecture for high availability
2. Bi-directional active/almost-active Sizzling-Hot-Takeover (SZT) architecture for higher availability
3. Bi-directional active/active hot-hot disaster tolerant architecture for continuous availability

In active/active architectures, Shadowbase Business Continuity software supports data partitioning to avoid data collisions, as well as data collision detection, notification, and resolution via customized solutions when running in a route-anywhere, load-balancing model.



Figure 1 depicts the source and target platforms and databases supported by Shadowbase Business Continuity software for uni-directional and bi-directional data replication and data streaming. Any source database can replicate to any combination of target databases. Shadowbase is highly customizable to users' needs.

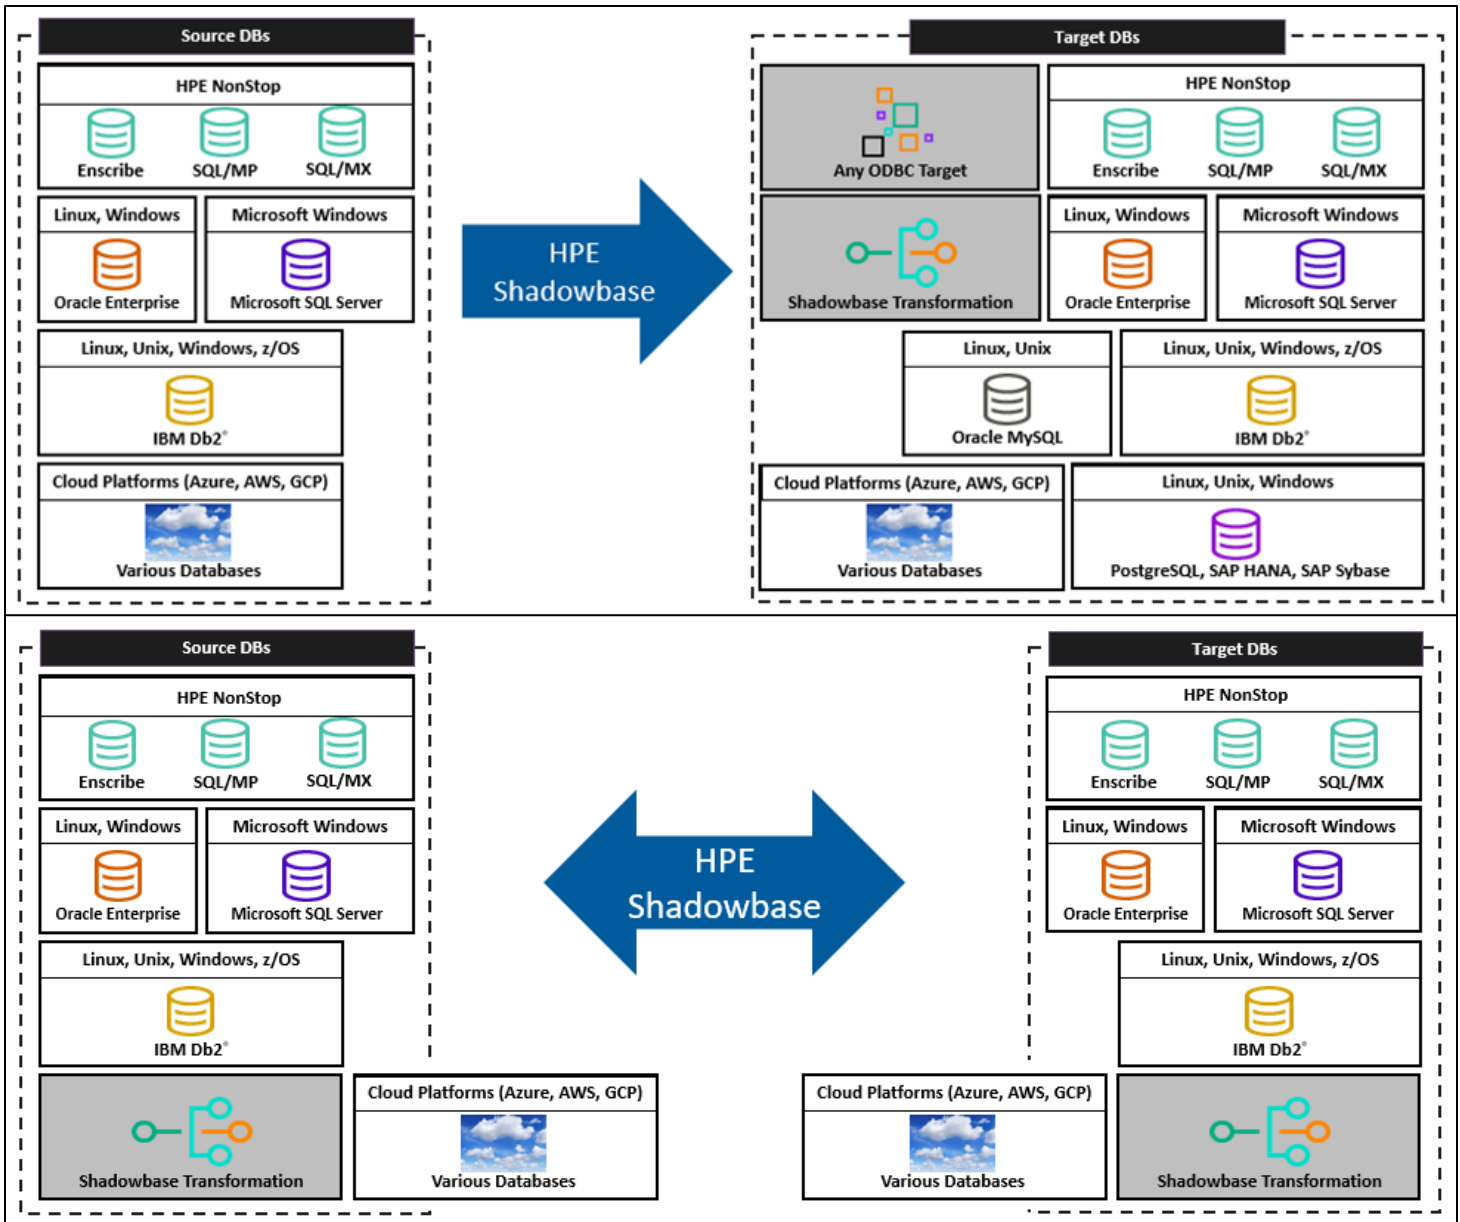


Figure 1: HPE Shadowbase Supported Platforms and Databases

**Note**  
All database and platform combinations are tested in a standard QA cycle. Please check with HPE Product Management (SPML) for the most up-to-date list of supported source and target platforms, databases, and environments.

**Product packaging**

The HPE Shadowbase Business Continuity product is available in four software packages:

- **HPE NonStop Shadowbase Business Continuity Basic**
  - This package is for NonStop to NonStop active/passive (uni-directional), business continuity data replication between homogeneous source and target environments. The package also supports heterogeneous source and target environments for Shadowbase Zero Downtime Migration (ZDM).



- **HPE NonStop Shadowbase Business Continuity Advanced**

- This package is for NonStop to NonStop active/almost-active (Sizzling-Hot-Takeover) and active/active (bi-directional) business continuity data replication. It adds in bi-directional data replication capability. HPE NonStop Shadowbase Business Continuity Basic is a prerequisite product.

- **HPE Shadowbase Business Continuity Basic for Other Servers**

- This package is for Other Server to Other Server in active/passive (uni-directional) business continuity data replication between homogeneous source and target environments. The package also supports heterogeneous source and target environments for ZDM.

- **HPE Shadowbase Business Continuity Advanced for Other Servers**

- This package is for Other Server to Other Server in active/almost-active (Sizzling-Hot-Takeover) and active/active (bi-directional), business continuity data replication. It adds in bi-directional data replication capability. HPE Shadowbase Business Continuity Basic for Other Servers is a prerequisite product.

## HPE Shadowbase Business Continuity software includes additional features

### HPE Shadowbase Online Loading (SOLV)

HPE Shadowbase SOLV utility provides online source database loading into a target database, meaning the source and/or target databases can be open for reading and updating while the load (or copy) occurs. In contrast to other loading products, there is no need to take the source or target application or databases offline while the load occurs.

Shadowbase SOLV can load audited and non-audited HPE NonStop Enscribe source files and HPE NonStop SQL tables into any target environment and database combination supported by the Shadowbase line of data replication products (e.g., HPE NonStop Enscribe or HPE NonStop SQL targets) or Shadowbase for Other Server targets (e.g., HP-UX, Microsoft® Windows®, Red Hat Linux distributions, Oracle Sun Solaris, and more, and support the most popular databases, including IBM DB2, Microsoft SQL Server, MySQL, Oracle, PostgreSQL, SAP® HANA, SAP Sybase, and others). Shadowbase SOLV can perform the load while regular replication is occurring, consuming the source database changes and merging them with the data being loaded to keep the target fully synchronized with the source while the load takes place. This synchronization avoids the need to queue (and ultimately drain) the replication changes while the load takes place.



### HPE Shadowbase File Chaser (also known as SOLVMGR)

HPE Shadowbase Business Continuity software uses the source database transactional TMF audit logs in order to capture and replicate the database change data. However, not all applications are audited. Shadowbase File Chaser reads non-audited, application-created-and-maintained log files or tables (containing transaction events) and injects those events into the replication stream to be applied into the target environment. This technique avoids the need for the source database to be audited, while ensuring such changes are replicated.

### HPE Shadowbase Online Loading and Verification Utility for OSS Regular Files (SOLVUTIL)

SOLVUTIL is available for loading, copying, validating, and file chasing HPE NonStop Open System Services (OSS) Regular Files (or “Flat Files”) that have data inserted at the end of file in append mode. (Note: OSS Regular Files are never TMF-audited.)

### HPE Shadowbase Queue Manager (QMGR)

For NonStop source environments, HPE Shadowbase Data Replication software supports extracting the data from the source environment, transferring it to the target environment, and applying it directly into the target database. The user can optionally configure the Shadowbase QMGR to queue the data to disk while it is being replicated. This step separates the extraction and delivery of the data from the replaying of the data being replicated into the target. This utility allows the target database to be taken offline, e.g., without affecting the delivery of the data to the target environment. Also, when using the QMGR on the target side, it is much faster to write the replicated data into a sequential queue file than to perform random I/O to the actual target database, resulting in less data loss in the event of a complete source system failure when using this utility.



## HPE Shadowbase Zero Downtime Migrations (ZDM)

HPE Shadowbase ZDM software enables continuous application services by eliminating planned downtime across even the most complex migrations, upgrades, platform refreshes, and conversions. It eliminates the risk of classic “big-bang” conversions that typically occur during outage windows.

The Shadowbase SOLV utility also plays an important role in ZDM. In the case of a platform or hardware refresh, the Shadowbase SOLV utility loads the database onto the new system from the database on the existing system while that existing database remains online for application processing. The Shadowbase SOLV utility thereby enables the new database to be built and kept synchronized with no outage of the existing production application or database while the data and application conversion and migration effort occurs. The new system can be thoroughly tested using production data while the existing system remains in full operation. When thorough testing has been completed, Shadowbase ZDM software allows the users to be phased over, or migrated all together, to the known-working, new environment that has an up-to-date database.

### Key Zero Downtime Migration Highlights:

- Eliminate planned downtime
- Avoid the risk of classic big-bang conversions
- Migrate platforms with no impact
- Migrate applications and databases with no impact
- Upgrade to new O/S releases or perform a platform refresh with no downtime

### Notes

1. For HPE NonStop source environments, HPE NonStop Shadowbase software reads and replicates database operations from the HPE NonStop TMF audit trail. If the application does not use HPE NonStop TMF, then HPE NonStop AutoTMF software must be used to enable this capability.

Alternatively, depending on the application, HPE Shadowbase technology provides two options for replicating non-audited Enscribe and SQL data: SOLV “Snap-Shot” Loads (Refreshes) and Shadowbase File Chaser technology. Using the SOLV solution, the user can periodically load (or refresh) the entire non-audited file or table (or a portion thereof) into the target environment. Additionally, Shadowbase File Chaser technology can “chase” the EOF of non-audited application-generated files or tables and inject those events into the stream for replication to a target database. Shadowbase SOLV and Shadowbase File Chaser are bundled with the Shadowbase Business Continuity Basic software package.

2. The Shadowbase Business Continuity software package cannot be used for data and application integration. For this purpose, HPE Shadowbase Data and Application Integration software must be licensed for the appropriate platform(s).

## HPE Shadowbase Zero Data Loss software

### Reduce outage costs with zero data loss (ZDL)

One of the most significant costs of downtime is data loss. With most business continuity replication technologies, there is the possibility of data loss when an unplanned outage occurs (this loss is caused by application database changes made on a source system which have not yet been replicated to a target system before the source system fails). Such data loss can cost a company millions of dollars, not to mention other significant impacts such as loss of customer loyalty, negative publicity, regulatory violations, brand reputation and even threats to human health and safety.

### Key Zero Data Loss Highlights:

- Zero data loss on unplanned system failure or software outage
- Patented low-latency synchronous replication technology
- No specific disk technology requirements
- No source-target system separation distance limits

Shadowbase ZDL software saves a company from all of the costly impacts of data loss. Using unique and patented synchronous replication technology, Shadowbase ZDL ensures that all database changes made on a source system are successfully replicated to a target (backup) system before the source application is allowed to commit (make permanent)



those changes. Thus, if the source system fails, no committed changes made to the source database will be lost; they will be present on the backup system and applied to the backup database. Unlike some synchronous replication products, Shadowbase ZDL does not require any specific disk technology (e.g., an HPE XP array is not required), and there are no hardware-imposed distance limitations between source and target systems.

## Product packaging

The HPE Shadowbase Zero Data Loss product is available in two software packages:

- **HPE NonStop Shadowbase Business Continuity Zero Data Loss**
  - This package is for NonStop to NonStop active/passive (uni-directional) business continuity data replication between homogeneous source and target environments. HPE NonStop Shadowbase Business Continuity Basic is a prerequisite product.
- **HPE NonStop Shadowbase Data and Application Integration Zero Data Loss**
  - This package is for NonStop to NonStop active/passive (uni-directional) data and application integration between heterogeneous source and target environments. HPE NonStop Shadowbase Data and Application Integration Basic is a prerequisite product.

---

## HPE Shadowbase Data and Application Integration software

### Eliminate data and application silos by streaming data for competitive advantage

Data replication is by far the most flexible, least disruptive to existing applications, and easiest to implement of all the various technologies that deliver real-time data integration and data streaming. The HPE Shadowbase Data and Application Integration product suite provides the replication and online copying architecture needed for implementation of effective real-time data integration and data streaming solutions.

When the need arises for replicating data between heterogeneous systems and databases (e.g., when building a data warehouse or populating a data lake), Shadowbase Data and Application Integration software provides support for in-flight user-defined data transformation, filtering, and cleansing.

The Shadowbase Data and Application Integration product suite supports all of the following applications, whether referred to as data integration, data synchronization, application integration, change data capture (CDC), data or event streams/streaming, trickle feeds, micro-batch updating, online data warehousing, etc.

#### Key Data and Application Integration Highlights

- Quickly and easily integrate your NonStop applications and data into the rest of the enterprise
- Master data management
- Rapidly integrate existing applications to create new value-add services
- Online data warehousing
- Real-time fraud detection
- Real-time business intelligence
- Real-time integration of homogenous or heterogeneous data silos
- Big data distribution



#### HPE Shadowbase Data Integration

HPE Shadowbase Data Integration software provides real-time, efficient data integration and synchronization by streaming changes made in one database to another; feeding data warehouses and real-time business intelligence solutions; and driving extract, transform, and load (ETL) utilities. Replication between heterogeneous databases is supported as shown in Figure 1. Any necessary data transformation between formats of the source and target databases is either performed automatically or can be easily customized.




## HPE Shadowbase Application Integration

HPE Shadowbase Application Integration software avoids the need to modify application code by sending or receiving database change events as they occur in real-time via a direct interface or publish/subscribe mechanism to other applications in other systems. As the application data changes, Shadowbase replication picks up the change data events in real-time. It filters, cleanses, and transforms these events as necessary, as well as delivers them to the target environment's application(s) using an efficient, event-driven architecture, regardless of the event input method used by the target application. Supported input methods include appending events into flat files, sending interprocess messages, establishing and utilizing TCP/IP sessions, and feeding middleware such as IBM MQ series. Therefore, Shadowbase application integration allows existing applications to be integrated in real-time to create powerful, new, value-add services.

## HPE Shadowbase Streams

HPE Shadowbase Streams is a method of using the Shadowbase data and application integration software to replicate, or stream, the individual changes made in one environment to another. More specifically, the data generated by one application is streamed to other applications and IT environments. Shadowbase Streams provides the utilities for integrating existing applications at the data or event-driven level in order to create new and powerful functionality and event-driven architectures (EDAs). The events are delivered in real-time as they are created without the need to modify application code, as Shadowbase Streams handles the extraction, transformation, and delivery of the information from the source database into the target database or application.



## HPE Shadowbase SOLV and HPE Shadowbase File Chaser

The HPE Shadowbase Data and Application Integration product also includes Shadowbase SOLV and Shadowbase File Chaser software. (These solutions are discussed above under HPE Shadowbase Business Continuity software.)

## Product packaging

The HPE Shadowbase Data and Application Integration product is available in four software packages:

### 1. HPE NonStop Shadowbase Data and Application Integration Basic

- This package is for NonStop-based active/passive (uni-directional architectures) data and application integration environments.

### 2. HPE NonStop Shadowbase Data and Application Integration Advanced

- This package is for NonStop-based active/active (bi-directional architectures) data and application integration environments. It adds in bi-directional data replication capability. HPE NonStop Shadowbase Data and Application Integration Basic is a prerequisite product.

### 3. HPE Shadowbase Data and Application Integration Basic for Other Servers

- This package is for active/passive (uni-directional architectures) data and application integration environments on Other Servers.

### 4. HPE Shadowbase Data and Application Integration Advanced for Other Servers

- This package is for Other Server to Other Server in active/active (bi-directional architecture) environments. It adds in bi-directional data replication capability. HPE Shadowbase Data and Application Integration Basic for Other Servers is a prerequisite product.

## Notes

1. It is common to interoperate with these software packages between NonStop and Other Server environments, whether using uni-directional or bi-directional replication. Simply order the appropriate software package(s) for each platform (NonStop or Other Server) to match the needs for each user environment. See "Technical Specifications" below for a list of currently supported platforms and database versions (please check with HPE Product Management (SPML) for the most up-to-date list of supported source and target platforms and environments).



2. The HPE Shadowbase Data and Application Integration software package cannot be used for business continuity purposes. For this purpose, HPE Shadowbase Business Continuity software must be licensed for the appropriate platform(s).

---

## HPE Shadowbase Essentials Bundle

### Ensure audit compliance and database consistency

Data replication is at the heart of the Shadowbase product portfolio. This key technology allows data to be copied between systems to keep database copies synchronized. It is also important to be able to monitor and, if necessary, correct that data in order to detect and resolve anomalous behavior, identify improper configurations, ensure continuation of proper business operations, and satisfy audit compliance requirements. The HPE NonStop Shadowbase Essentials software bundle includes additional products to meet these needs.



#### Key Essentials Highlights

- Detect and resolve anomalous behavior
- Satisfy audit compliance requirements
- Restore a corrupted database
- Transform Enscribe records into their SQL table equivalents
- Systematically transform data formats and create customization scripts

#### HPE Shadowbase Audit Log (SAL)

HPE Shadowbase Audit Log creates a searchable archival/reporting database of transactional activity (e.g., inserts, updates, and deletes) for application change data auditing purposes. The database is created on an off-platform target database (e.g., Oracle or SQL Server), which can subsequently be queried to review source database changes, including what and when the data was changed. SAL is available for the HPE NonStop Server platform as a source (Enscribe, NonStop SQL/MP, NonStop SQL/MX).

When configured, SAL will selectively replicate the HPE NonStop TMF audit trail database change activity (the 'after' image of an insert, the 'after' image of an update, and the 'before' image of a delete event) to the target reporting database, storing the change information as a series of rows, one row per insert, update, and delete statement that was originally issued in the source database. SAL prepends certain source environment information to the data as it stores the data, including the original transaction ID, the time the transaction was executed, and the event sequence that occurred at the source.

---

#### Note

1. Prerequisite products for SAL:
  - HPE Shadowbase Audit Log for Other Servers is required in order to create the reporting database.
  - HPE Shadowbase Data and Application Integration for Other Servers Basic is required
  - HPE NonStop Shadowbase Data and Application Integration Basic is required.

---

#### HPE Shadowbase Audit Reader (SAR)

HPE Shadowbase Audit Reader analyzes and displays all audited database activities on HPE NonStop server systems. SAR enables investigating how the application is changing data by showing sequential application transactions to database files and tables. Both current and historical transactional information can be analyzed using a variety of search criteria. SAR mines this information from the HPE NonStop TMF audit trail files. It optionally reads "foreign" audit (audit generated on another HPE NonStop server system that has been copied). SAR supports both Enscribe and NonStop SQL I/O events and displays both before and after I/O images, as well as the "undo" for aborted transactions. SAR software provides an application view of the change data formatted by data type by querying the NonStop SQL catalog or by accessing the Enscribe DDL information (when available). The SAR utility is necessary for reviewing and verifying how the database data is being changed by the applications, and is useful for auditing purposes.





**HPE Shadowbase Database Recovery software – Shadowbase UNDO**

HPE Shadowbase UNDO software allows the user to selectively “undo,” or rollback, the changes made to a database, to a previous date/time. This option is useful to restore a database to an earlier set of values using a selective “as of” approach. Shadowbase UNDO software restores a corrupted database by undoing corrupting database changes without affecting correct changes, thus rolling the database back to a previous known, consistent, and working state. By maintaining an “Undo Queue” of changes that have been made to the database, Shadowbase UNDO software can follow the Undo Queue in reverse chronological order to the initial point of corruption and can reverse any corrupting DML (data manipulation language, e.g., insert, update, and delete) changes that have been made. The rollback of corrupted data can be accomplished by Shadowbase UNDO software while the affected files and tables remain online and the application continues its processing functions. This option is helpful, e.g., if an errant version of an application is deployed, and it corrupts part of the database; Shadowbase UNDO allows for selectively removing that corruption from the database without requiring an application outage. Shadowbase UNDO software supports Enscribe, NonStop SQL/MP, and NonStop SQL/MX databases.

---

**Note**

Either HPE NonStop Shadowbase Business Continuity Basic, or HPE NonStop Shadowbase Data and Application Integration Basic are prerequisite products for Shadowbase UNDO.

---

**HPE Shadowbase Database Recovery software – Shadowbase REDO**

HPE Shadowbase REDO software maintains a “Redo Queue” of all changes that have been made to a database. If a new version of an application is to be run, or if a new database version is to be deployed, or if some other system change is to be made, then a copy of the database is created and saved before executing the change. This step is often accomplished by unmirroring the application database and saving one of the mirrors, or by taking a full database backup. The upgraded system is then run with the remaining mirror (and can be re-mirrored if sufficient additional disk space is available). While the upgraded system runs, the Redo Queue saves a copy of the changes made. If subsequent, serious problems result in database corruption (e.g., a catastrophic latent application bug is uncovered), the saved mirror (or database backup) can be quickly restored. Shadowbase REDO can then be used to selectively “roll forward” through the changes in the Redo Queue to re-apply those that need to be re-applied, even if the original and new data formats are different. Hence, no (new) data will be lost even when a “fail-back” occurs to the original database format. Shadowbase REDO software supports Enscribe, NonStop SQL/MP, and NonStop SQL/MX databases.

**HPE Shadowbase Data Definition Language Utility (SBDDLUTL)**

Many Enscribe users need to replicate their NonStop Enscribe data into external systems and environments, e.g., to feed a data warehouse. This process is usually accomplished by converting the Enscribe information into standard SQL data structures and schemas. The HPE Shadowbase Data Definition Language utility (SBDDLUTL) simplifies this process by providing a powerful interface to convert and map the Enscribe DDL data structures (records, fields) into their SQL equivalents (tables, columns). SBDDLUTL supports all of the target SQL databases that Shadowbase supports, including NonStop SQL/MP, NonStop SQL/MX, Oracle, Microsoft SQL Server, IBM Db2, Oracle MySQL, SAP HANA, and SAP Sybase.

**HPE Shadowbase Data Mapping Utility (SBMAP)**

HPE Shadowbase Data Mapping Utility is a powerful, sophisticated, and extensible scripting “language” that can be used to tell Shadowbase software how to transform source data into target data formats. SBMAP is currently available for NonStop Enscribe and SQL/MP source environments that need to transform, filter, or cleanse the data before replication.

**HPE Shadowbase DDL Command Replication (DCR)**

The HPE Shadowbase DCR Utility automatically replicates and applies Enscribe FUP or SQL/MP SQLCI (DDL) source commands to the target database. It automatically adjusts the source command to match the target environment’s details (e.g., file/table name mapping). These DDL commands integrate with HPE Shadowbase (DML) real-time replication. DCR is particularly helpful to customers migrating from RDF/SDR (a mature product) to HPE Shadowbase replication.

**HPE Shadowbase Extract, Transform, and Load (ETL) Toolkit**

The HPE Shadowbase ETL Toolkit is used to extract database changes or initial load data into flat files for subsequent ETL loading into a data warehouse using comma-separated value (CSV), fixed-position, and tab-delimited formats.

More specifically, for certain applications, an existing vendor’s ETL loading utility can be used to load either initial data or change data into a data warehouse. The Shadowbase ETL Toolkit enables a Shadowbase user to perform these tasks. For example, when coupled with the Shadowbase SOLV loader, the Toolkit can be used to extract select data from a source database, transform/cleanse it into the proper format, and save it into flat files that the ETL loading utility can understand and load into a target database environment not directly supported by Shadowbase replication. Using the



Shadowbase ETL Toolkit, change data from the source database's audit trail (change log) can be extracted and processed into a flat file format that can be incrementally loaded into the target environment, using a micro-batch incremental update approach to keep the target environment synchronized with the source.

### HPE Shadowbase SQL/MP Schema Conversion (SBCREATP)

The HPE Shadowbase SQL/MP Schema Conversion utility (SBCREATP) converts and maps SQL/MP table schema data structure definitions (rows, columns) into target SQL equivalents (tables, columns), and produces an editable CREATE TABLE statement for the target SQL environment. SBCREATP supports the target SQL databases that Shadowbase supports, including HPE NonStop SQL/MX, IBM Db2®, Microsoft SQL Server, Oracle, Oracle MySQL, PostgreSQL, SAP HANA, and SAP Sybase. It is a key utility for customers performing Shadowbase data integration replication from NonStop source SQL/MP tables out to SQL target tables. Without SBCREATP, the conversion work will have to be done by hand, which is a time-consuming and error-prone process.

## HPE Shadowbase Data Validation software

### Verify source and target database accuracy

#### HPE Shadowbase Compare

HPE Shadowbase Compare software allows a target Enscribe file or NonStop SQL table to be compared to its source and report on all discrepancies found between the two. In most cases, HPE Shadowbase Compare can resolve and repair these discrepancies, should the DBA desire to do so.

Discrepancies may occur, for example, for a variety of reasons including if the user updates the target database (but not the source), misconfigures the replication environment, or accidentally purges necessary audit trails.

Note that Shadowbase Compare is a stand-alone solution that works with any data replication engine. It can also work for environments that do not have a data replication engine.



#### Why is Compare Needed?

Essentially, it's a form of *replication insurance*. The replication engine may be running, but what about operator and procedural errors, configuration errors, malware, and bugs? A data comparison solution will highlight discrepancies, and alert to any potential problems. Data validation is a best practice, can help satisfy regulatory requirements, assure key stakeholders, and avoids unnecessary risks.

Shadowbase Compare allows the application(s) to remain online while running, and is essential for validating that a target database matches its source. It follows auditing best-practices. It is designed for comparing HPE NonStop to HPE NonStop, like-to-like source/target environments where Expand or TCP/IP is available between the nodes (e.g., for validating business continuity environments, both initially as well as periodically).

HPE Shadowbase Compare is particularly useful to ensure replication is setup properly (e.g., in active/passive business continuity environments it identifies anomalies if an application is accidentally run against the standby database instead of the production database). Therefore, it is highly recommended that Shadowbase Compare be used to validate every HPE NonStop replication environment, regardless of the replication engine in use. It provides confidence (or certification) that the source and target data matches, and that the target database is ready to take over on a moment's notice. Shadowbase Compare also helps satisfy the PCI requirement that replicated or copied files/tables match.

#### HPE Shadowbase Repair

HPE Shadowbase Compare also contains a repair function to assist with correcting or repairing highlighted SQL/MP or SQL/MX table differences. HPE Shadowbase Repair works on any two tables for like-to-like databases (MP or MX), and is particularly useful in business continuity environments. The user can determine if, and when, to apply the repair operation.



## Technical specifications<sup>i</sup>

**Table 1: Technical specifications**

<b>Operating systems<sup>ii</sup></b>	<p>HPE NonStop OS, J, and L- series (H06.25 and later, J06.10 and later, L15.02 and later) HPE Virtualized NonStop HPE-UX Itanium 2 (11.31) IBM AIX (7) IBM (z/OS (and others) for Db2 sources when using a Data Appliance) Red Hat Enterprise Linux (RHEL) (7, 8, and 9) Oracle Solaris (11) Oracle Solaris x86/x64 (11) SuSE Linux Enterprise Server (SLES) (12 and 15) Windows Server 2016 through 2022</p>
<b>Source databases and environments<sup>iii</sup></b>	<p>HPE NonStop Enscribe HPE NonStop SQL/MP HPE NonStop SQL/MX (SQL/MP Tables) HPE NonStop SQL/MX (Native ANSI Tables) IBM Db2 10.5 and 11.x on AIX, RHEL, and Windows Microsoft SQL Server 2014 through 2022 on Windows<sup>iv</sup> Oracle 19c and 21c on HPE-UX, AIX, RHEL, Solaris, Solaris x86, and SLES<sup>v</sup> Application created log files (via customer-provided user exits)</p>
<b>Target databases and environments</b>	<p>HPE NonStop Enscribe HPE NonStop SQL/MP HPE NonStop SQL/MX (SQL/MP Tables) HPE NonStop SQL/MX (Native ANSI Tables) IBM Db2 (10.5, and 11.x) on AIX, RHEL, and Windows IBM Db2 (z/OS (and others) when using a Data Appliance) Microsoft SQL Server 2014 through 2022 on Windows MySQL 8 on RHEL and Windows Oracle 19c and 21c on HPE-UX, AIX, RHEL, Solaris, Solaris x86, and SLES SAP Sybase ASE 16 on RHEL, SLES, and Solaris SAP HANA 2 on RHEL, SLES, and Windows PostgreSQL 11 through 15 on RHEL ETL Extract Files (via customer-provided user exits) Flat Files (via customer-provided user exits) Inter-process message delivery/notification (via customer-provided user exits); message feeds into middleware (via customer-provided user exits)</p>
<b>Operating systems, databases, and environments covered under a statement of work (SOW)<sup>vi</sup></b>	<p>HPE NonStop S series G06.xx and above (source or target) HPE NonStop OS, H, and J series (H06.20 and earlier, J06.09 and earlier) 32-bit versions of HPE Shadowbase for Other Servers Newer versions of operating systems or databases as they release SAP Sybase 12.5, ASE 15 (source; requires corresponding Sybase Replication Server version) Sybase IQ (target) Teradata target and other ODBC target databases Flat files (source or target) Interprocess message delivery/notification (target) Message feeds into middleware (target) systems and versions not listed above</p>

Footnotes on next page



## Summary

Catastrophic events happen. It's not a question of *if*, but *when* one will impact IT services. That failure could happen *later today or even tomorrow*. And if it happens today, it may also happen **again** tomorrow. Are you prepared? Once the event happens, then what?

An HPE Shadowbase Business Continuity solution architected for continuous availability can help minimize the damage, providing the best available recovery time (RTO) and recovery point (RPO) objectives at a total cost of ownership (TCO) that makes sense for the business.

As workload volumes and the complexity of heterogeneous IT environments increase, the pressure is on IT to deploy simple, yet elegant solutions to these complex problems. To further complicate matters, customers are increasingly expecting continuous services with no tolerance for downtime and instant access to current information from any access point.

HPE Shadowbase real-time data replication and integration solutions meet the most demanding business continuity as well as data and application integration requirements, delivering the performance and reliability that HPE customers need to manage their complex IT landscape.

Since HPE Shadowbase software is optimized for performance and reliability, it achieves very low latencies under extreme workload volumes. With its modular architecture, it provides solutions to a wide range of business challenges, from eliminating downtime and data loss to enabling rapid development of new solutions.

HPE Shadowbase is HPE's strategic data replication and data streaming solution for homogeneous and heterogeneous platforms and databases. It has proven to be successful in the most demanding mission-critical industries, such as finance, banking, communications, manufacturing, and healthcare.

Find out how HPE Shadowbase software can help you prepare for the inevitable.

## Resources

For more information, refer to the HPE Shadowbase material available at HPESC [hpe.com/info/nonstop-ldocs](https://hpe.com/info/nonstop-ldocs), or visit [shadowbasesoftware.com](https://shadowbasesoftware.com). Learn more at [hpe.com/info/nonstop](https://hpe.com/info/nonstop)

© Copyright 2024 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

4AA4-4123ENW, April 2024 DRAFT




---

### Footnotes:

<sup>i</sup> For the most recent supported databases by version, please reference [Supported Databases and Platforms](#).

<sup>ii</sup> HPE Shadowbase for Other Servers software (non-NonStop platforms) ships as 64-bit executables unless otherwise noted as 32-bit.

<sup>iii</sup> Source databases support uni-directional and bi-directional replication. IBM Db2 is supported as a source and a target, uni-directionally and reciprocally (meaning the same table cannot be both a source and a target), at this time

<sup>iv</sup> SQL Server source replication is *trigger*-based.

<sup>v</sup> Oracle source replication is *log*-based.

<sup>vi</sup> The listed operating systems, databases, and environments may be supported via a customized services Statement of Work (SOW) agreement. Contact HPE Product Management for more details.

---

