HPE Shadowbase Total Replication Solutions
Product Datasheet

A Gravic, Inc. White Paper

HPE Shadowbase Total Replication Solutions Product Datasheet

Executive Summary

In today's business world, access to real-time online transactional data is a competitive advantage. To realize the advantage, this data must be available at any time, all the time, from anywhere, and it must be current. The corollary is that the inability to access or update this current data carries a significant business cost, possibly measured in many thousands of dollars per second or large regulatory fines. These requirements necessitate an IT infrastructure that is continuously available, and where transactional data is rapidly distributed wherever it is needed, to other systems and applications. This environment is likely to be heterogeneous, with many different platform types and databases.

Gravic, Inc. is a world leader in providing innovative data collection, transformation, and distribution solutions. Shadowbase is Gravic's real-time data replication and data integration software solution for HPE NonStop, Linux, Unix, and Windows platforms. The award-winning HPE Shadowbase Total Replication Solutions product suite provides the means to deliver on the promise of always available data, via reliable low-latency real-time data replication and distribution across heterogeneous systems, databases, and applications. With these powerful capabilities, HPE Shadowbase solutions provide your business with the tools needed to realize the competitive advantage of continuous access to real-time transactional data across the enterprise, and to avoid the significant costs of system and data unavailability.

HPE Shadowbase software solutions provide your business with the tools needed to realize the competitive advantage of continuous access to real-time transactional data across the enterprise, and to avoid the significant costs of application service and data unavailability.

Table of Contents

Table of Figures

HPE Shadowbase Total Replication Solutions Product Datasheet

HPE Shadowbase Business Continuity (BC)

Avoid business losses by eliminating planned and unplanned downtime.

Stuff happens. Unavoidable events such as fires, floods, power outages, hardware faults, etc., can take IT production systems offline. A recent study[1] shows that the average number of complete datacenter outages is one per year, with an average duration of 91 minutes. Other reports[2] estimate that the average business revenue lost per hour of downtime is $1.4M (USD). Thus, the average annual business costs from unplanned outages is about $2.1M. For some industries the costs will be significantly higher. Further, the U.S. Bureau of Labor found that 93% of companies that suffer a significant data loss are out of business within five years. The net of these chilling statistics is that serious IT outages are not rare, and their consequences to the business can be dire. They are also completely avoidable via the implementation of an appropriate business continuity (BC) plan.

BC encompasses those activities that an enterprise performs to maintain consistency and recoverability of its data, operations, and services. Application availability depends upon the ability of IT services to survive any fault, whether it is a server failure, a network fault, or a datacenter disaster. Data availability depends on the existence of up-to-date backup data copies. An enabling technology for achieving high or continuous availability for application services and the timely backup of important data is *data replication*. This is where the HPE Shadowbase product suite from Gravic, Inc. provides solutions.

By replicating data in real-time between systems, Shadowbase software enables recovery from unplanned outages in times ranging from minutes (high availability, disaster recovery) to immediate (continuous availability, disaster tolerance), with minimal data loss (measured in sub-seconds). Shadowbase solutions support active/passive, *sizzling-hot-standby* (SZT), and fully active/active BC architectures.[3]

These same Shadowbase features can also be used to minimize planned downtime. Planned downtime is traditionally used to perform maintenance activities such as database reloads, point-in-time production system snap-shop backups, or to install new system, platform, network infrastructure, database, or application versions. By keeping backup systems current with data updates performed by production systems, Shadowbase software allows backup systems to be quickly brought up to production status, minimizing outage downtime.

Shadowbase solutions take this protection a step further to totally eliminate planned downtime by allowing the maintenance activity to be performed in parallel with full production processing. In this mode, Shadowbase replication keeps each environment's data synchronized, regardless of the data formats in the existing and new environments. Once the maintenance activity is completed, Shadowbase technology supports both phased as well as instantaneous cut-over to a known-working and known-correct environment, allowing what were previously disruptive maintenance operations to be performed with no loss of business services.

Achieve the competitive advantages promised by eliminating planned and unplanned downtime.

HPE Shadowbase Real-Time Data Integration and Application Integration

Turbocharge your business through real-time, event-driven data integration and application integration.

[1]Source: Ponemon Institute.
[2]Sources: Network Computing, The Meta Group, Contingency Planning Research.
[3]For descriptions of these terms, please see our white paper: *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

In today's competitive environment with high consumer expectation, business decisions are based on the most current data available to the enterprise that will improve customer relationships, increase revenue, and maximize operational efficiencies.

The speed of today's processing systems has moved classical data integration into the realm of real-time. The result is *real-time (business) data integration (RTBDI)*, also referred to as *real-time data integration* (RTDI). For example, in a classic data warehousing example, the speed of data integration has moved operational analytical processing into *real-time business intelligence* (RTBI). Operational business transactions are fed as they occur to an RTBI system that maintains the current state of the enterprise. The RTBI system not only supports the classical strategic functions of data warehousing for deriving information and knowledge from past enterprise activity, but it also provides real-time tactical support to drive enterprise actions that react to immediate events. As such, it replaces both the classical data warehouse and *enterprise application integration* (EAI) functions. An RTBI system with a consolidated real-time view of the business across the enterprise enables the provision of new business services (applications) which were previously impossible, achieving competitive advantage.

Shadowbase solutions enable the provision of RTBI capability via reliable low-latency real-time data replication and distribution between heterogeneous systems, databases, and applications. Using Shadowbase technology, data is collected, transformed, and distributed wherever it is needed, to other databases or directly to applications. Changes made in any database are quickly and easily integrated in real-time to other data environments to keep that target environment synchronized. So-called data silos, or "islands of information," are eliminated. With Shadowbase replication, applications that were once isolated are now able to interoperate in an efficient event-driven fashion in real-time. Critical data generated by one application is distributed by Shadowbase replication to other applications where it can be analyzed and acted upon immediately.

HPE Shadowbase Zero Downtime Migration (ZDM)

Enable significant changes or upgrades to be made to existing IT infrastructure (hardware and software), with no loss of business services.

Both high and continuous availability require the minimization or elimination of *unplanned downtime* due to unexpected failures such as those caused by natural disasters. They also require the elimination of *planned downtime* such as that needed for an operating system or hardware upgrade or for installation of new database or application versions. After all, users are unable to access business services and are considered "down" in either case. Whenever changes are made to a system – whether these changes affect hardware, software, data, networks, or operating procedures – there must be a process to make these upgrades without denying users access to their business services. When upgrades are undertaken without denying business services to users, it is called *zero downtime migration* (ZDM).

Shadowbase solutions provide the means to achieve ZDM and eliminate planned, as well as unplanned, downtime. When performing routine maintenance to existing systems, by using HPE Shadowbase data replication to keep backup systems current with data updates performed by production systems, backup systems can quickly be brought online allowing them to continue to provide business services while the maintenance is performed.

When performing a major hardware or software change or migration (e.g., to a new hardware platform, or a new version of an application, operating system, or database), Shadowbase technology can be used to efficiently copy the source (existing) database to the target (new) environment/database while the source environment is still undergoing active transaction processing, thus avoiding an outage in order to build the new environment/database. Shadowbase replication coordinates the database copy with the online updates to ensure that the data loaded is properly serialized with the data being replicated. Thus, at the end of the load cycle, the new database is up-to-date with the production database. Once this is achieved, the new system can be thoroughly tested while Shadowbase replication continues to keep the database copy up to date. When successfully tested, users can be incrementally phased over to the new system, all the while providing full application business services, and a ZDM system upgrade is achieved. When used in this way, HPE

Shadowbase ZDM eliminates the impact of disruptive upgrade activity, avoiding the all-or-nothing risk inherent in so-called "big bang" cutovers.[4]

HPE Shadowbase Data Utilities

Provide the tracking, reporting, comparison, and correction facilities necessary for audit compliance and ensuring the integrity of your data.

Audit compliance has become an increasingly important part of most businesses. Failure to meet compliance requirements can result in heavy fines, or even suspension of operations. Fraudulent activity can result in significant costs if left unchecked. It is therefore imperative that your business knows what data is changing, when, how, and by whom.

Shadowbase solutions provide the tools necessary to help your business meet these requirements. As transactions are processed on production systems, the data changes associated with them are typically written to a transaction audit trail (or log). However, the format of these audit trail files is proprietary. Shadowbase technology can mine these audit trails directly using a powerful query and filtering capability, and display transaction activity as it happens. Shadowbase software can also create an archival database from these audit trails of all transaction change data performed by the production systems. This archival database is searchable, making it easy to find specific transactions or suspicious data changes. Threshold trigger points can be configured to automatically notify applications when data is changed in specific ways (e.g., when a record value is increased beyond a certain amount).

When replicating data for BC or data integration purposes, it is important to know that the target database is consistent with the source database. HPE Shadowbase solutions provide tools to compare and validate that the target database matches the source database, and to provide insight into the specific differences found. These capabilities can also be used for example when an erroneous application is put into production, and all the data changes made by it are required to be backed out or undone, even while the application remains online to users.

HPE Shadowbase Zero Data Loss (ZDL)

Shadowbase software supports both asynchronous and synchronous replication. With asynchronous replication, change data is sent to the target system after the changes have been made on the source system. In rare circumstances, it is possible for data to be lost in the event of a failure. For some applications lost data is not a problem, but for others, the data is critical and must not be lost. Shadowbase Zero Data Loss (ZDL), a future technology, uses synchronous replication to solve this problem. No data is changed on the source system unless the data has been safe-stored on the target system, ensuring no data loss, no matter what the failure. Asynchronous replication also allows for the possibility of data collisions,[5] which may be unacceptable for some applications. Shadowbase synchronous replication also solves this problem with another future technology, Shadowbase ZDL+, preventing the data collision from occurring in the first place. Shadowbase software with synchronous replication is the solution for the most mission-critical applications, where data loss and/or data collisions cannot be tolerated.[6]

[4] For more information on ZDM, please read our white paper: *Using HPE Shadowbase Software to Eliminate Planned Downtime via Zero Downtime Migrations*.

[5] In an active/active architecture, a data collision is where the same data field is simultaneously updated on two (or more) copies of a database. When replication of that update between those copies occurs, each update will overwrite the other, resulting in an inconsistent database.

[6] For further information on Shadowbase ZDL and ZDL+, visit https://www.shadowbasesoftware.com/solutions/business-continuity/zero-data-loss/

The HPE Shadowbase Product Suite

1. Business Continuity

HPE Shadowbase Data Replication

The core of the product portfolio, HPE Shadowbase data replication provides low-latency uni-directional and bi-directional data replication between homogeneous and heterogeneous systems and databases with scalability, selectivity, and sophisticated data transformation and mapping facilities. Shadowbase data replication provides a BC solution to meet any business requirement:

- Uni-directional active/passive disaster recovery for *high availability*
- Bi-directional active/almost-active *sizzling-hot-takeover* (SZT) for *higher availability*
- Bi-directional active/active hot-hot disaster tolerant architecture for *continuous availability*

In active/active architectures, Shadowbase data replication provides automated data collision detection, notification, and customized resolution facilities.[7]

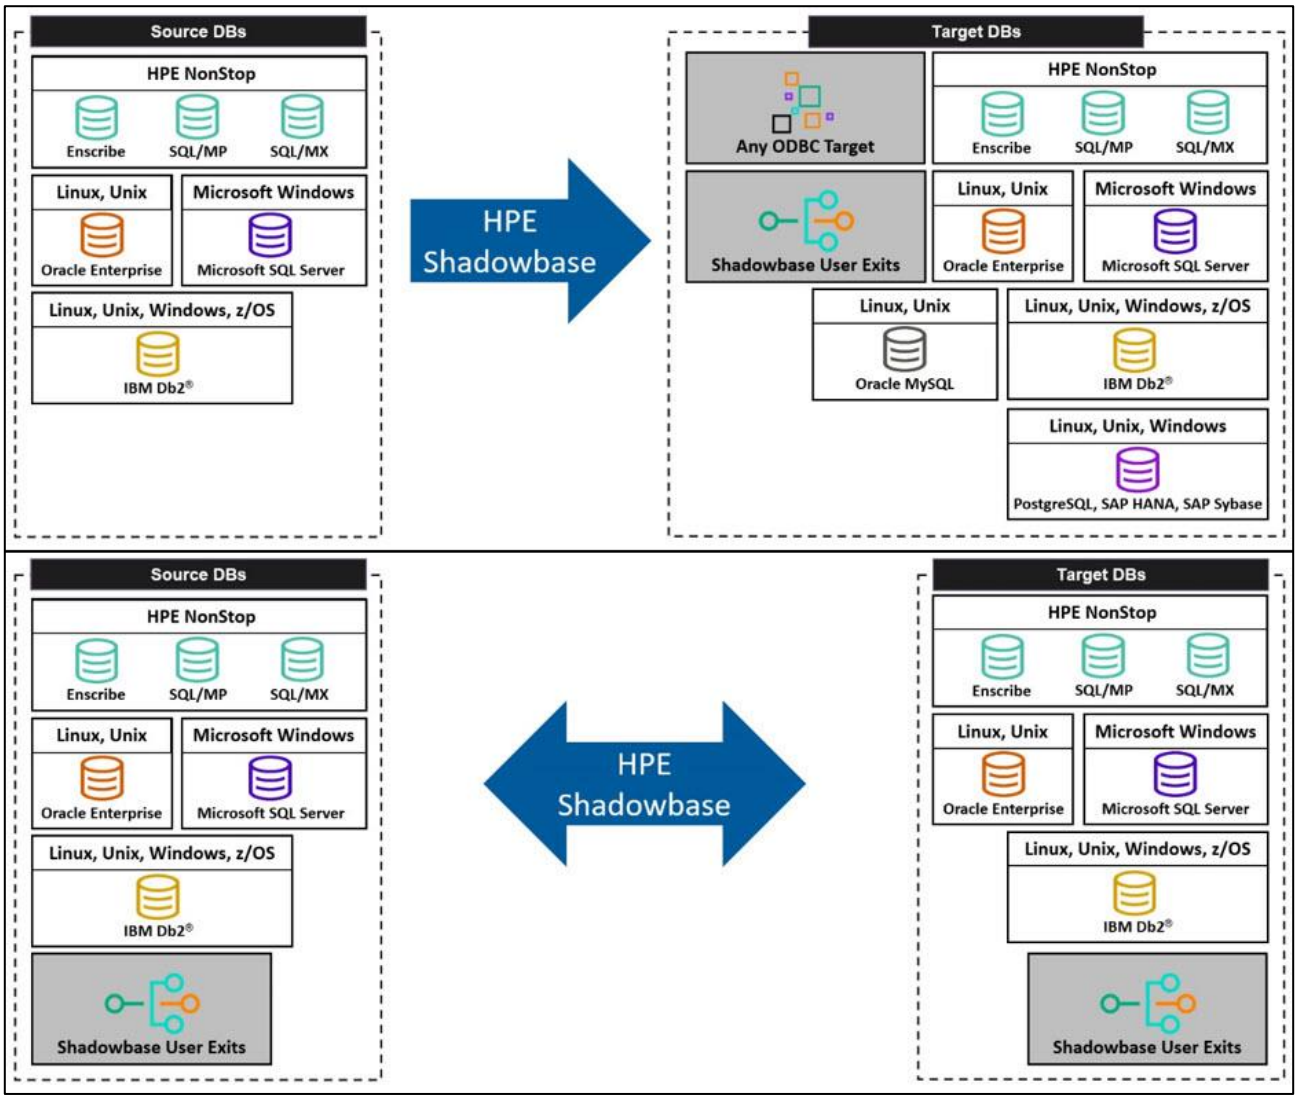

Figure 1 – HPE Shadowbase Supported Platforms, Databases, and Operating Environments

[7]For a description of data collisions in active/active replication environments, please see our white paper: *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

Shadowbase data replication supports replication between heterogeneous systems. Figure 1 depicts the source and target platforms, databases, and operating environments that Shadowbase data replication supports for uni-directional replication (top), and bi-directional replication (bottom).[8] Please note that all combinations are tested as part of a standard QA cycle.

Shadowbase data replication also comprises these additional features:

HPE Shadowbase Queue Manager (QMGR)

Shadowbase data replication supports extracting the data from the source environment, transferring it to the target environment, and applying it directly into the target database, or the data can be queued to disk while it is being replicated to separate the extraction and delivery of the data from the replaying of the data being replicated. 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 failure. This configurable option is provided by the Shadowbase Queue Manager (QMGR). In parallel, Shadowbase replication will read the data from the queue file and apply it to the target database.

HPE Shadowbase File Chaser

Shadowbase data replication uses the source database transactional audit logs in order to capture and replicate the database change data. However, not all data updates are audited. HPE Shadowbase File Chaser reads non-audited *application-maintained* log files/tables and injects those events into the replication stream to be applied into the target environment, thereby avoiding the need for the source database to be audited, while ensuring that such changes are also replicated.

HPE Shadowbase Enterprise Manager (SEM)

Shadowbase software is easy to configure, monitor, and operate. Shadowbase Enterprise Manager (SEM), a Windows-based GUI command and monitoring interface, is available to control and monitor Shadowbase replication on all of the platforms that Shadowbase software supports. Ease of use features such as simple graphical stoplights that are configured for each Shadowbase process or component make the product easy to use for new users and operators. Error messages from all processes and platforms are viewable from the SEM interface, and SEM provides an email/warning interface to pagers to alert operators of impending problems.

2. Real-Time Data and Application Integration

Of all the various technologies that may be used to deliver real-time data and application integration, data replication is by far the most flexible, least disruptive to existing applications, and easiest to implement. The HPE Shadowbase product portfolio provides the replication and online copying facilities needed for implementation of effective real-time data integration systems.

HPE Shadowbase Streams for Data Integration

HPE Shadowbase Data Integration 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 facilities, 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 performed automatically (and may be easily customized).[9]

HPE Shadowbase Streams for Application Integration

Shadowbase Application Integration sends or receives database change events as they occur in real-time via a direct interface or publish/subscribe mechanism to other applications in other systems, avoiding the need to modify your application code. As the application data changes, Shadowbase replication picks up the change

[8] Email Shadowbase Product Management for the latest list of supported source and target platforms and environments at SBProductManagement@gravic.com or check the Supported Databases and Platforms page.

[9] For more information, please read our white paper, *HPE Shadowbase Streams for Data Integration*.

data events, filters, cleanses, and transforms them as necessary, and replicates them to the target environment for delivery into the target applications using an efficient, event-driven architecture.[10]

HPE Shadowbase Online Loading and Verification (SOLV)

SOLV is a utility that provides online loading of a source database into a target database, meaning that 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, when using SOLV there is no need to take either the source or target databases offline (and make them inaccessible to applications) while the load occurs. SOLV can load audited and non-audited NonStop Enscribe source files and NonStop SQL tables into any target environment and database combination supported by the Shadowbase line of data replication products.[11]

HPE Shadowbase SOLVUTIL for HPE NonStop Open System Services (OSS)

SOLV for OSS supports loading/copying OSS *regular* files and validating the data consistency, and can create a backup copy of the data and synchronize it with the source. It leverages the SOLVUTIL process (a low-level Shadowbase replication process that runs in the Guardian space) to copy a particular file (or set of files) from the source to the target.

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

ETL is used to extract database changes or initial load data into flat-files for subsequent extract, transform, and load (ETL) loading into a data warehouse using common comma-separated value (CSV), fixed-position, and tab-delimited formats. More specifically, for certain applications, a customer may want to use an existing vendor's ETL utility to load either initial data, or change data, into a data warehouse. ETL enables a Shadowbase user to perform these tasks. For example, when coupled with SOLV's loading process, ETL 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 a user's ETL utility can understand and load into a target database environment not directly supported by Shadowbase replication. Using ETL, the customer can also periodically extract source database change data from the source database's audit trail (change log), and process it 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 Data Transformation

As well as providing the distribution fabric to deliver change data in real-time to other databases and applications, HPE Shadowbase includes powerful capabilities to transform that change data into whatever format is required by the target database/application, data may be aggregated, disaggregated, and transformed. The following methods are available:

- ***SBMAP*** – a scripting "language" that can be used to tell Shadowbase how to transform source data into target data formats. Powerful, sophisticated and extensible.
- ***SBCREATP*** – a utility that converts the schema structure of an SQL/MP table into another SQL variant. At this time, it reads the SQL/MP catalog to extract an existing table schema, and converts the SQL/MP columns/data types into Oracle or Microsoft SQL server CREATE TABLE equivalent syntax.
- ***SBDDLUTL*** – a utility that reads an HPE Enscribe DDL record definition, simplifying the replication of unstructured Enscribe data into structured SQL databases. SBDDLUTL includes features to allow manipulation of the source fields when creating the target columns, including dropping fields, renaming fields, and normalization of the non-normalized Enscribe data.
- ***User Exits*** – enable the execution of customized user logic at various points in the Shadowbase replication stream. More complex than a scripting language, but extremely flexible, enabling almost any kind of data transformation to be performed.
- ***DBS Mapping*** – a scripting "language" for target-side other-server platforms. With capabilities such as drop all events for a target table, drop certain events for a target table, convert updates to inserts, drop column(s), concatenate (text) columns, formatting, and character conversion/replacement capability.

[10]For more information, please read our white paper, *HPE Shadowbase Streams for Application Integration*.

[11]For the most up-to-date list, please visit our Supported Databases and Platforms page.

- ***Miscellaneous*** – there are a number of parameter settings that can be set in Shadowbase to assist with data transformation, for example, to convert non-printable binary data in a character field to spaces. However, these work on the data record at an aggregate, not field level; SBMAP and the other methods described above allow for much finer data transformation granularity.

3. Zero Downtime Migration (ZDM)

HPE Shadowbase ZDM

By online replication of data from production source to backup target systems, Shadowbase ZDM enables production systems to be taken offline for maintenance while backup systems continue to provide the business services.[12] Shadowbase ZDM provides zero down-time migration support for continuous application services by eliminating planned downtime across complex migrations, upgrades, and conversions. It also avoids the risk of classic "big bang" conversions during outage windows.

Shadowbase SOLV also has a role to play in ZDM. In the case of a migration to a new system or database, Shadowbase SOLV is used to load the database on the new system from the database on the existing system, while that existing database remains online. Shadowbase SOLV thereby enables the new database to be built and kept synchronized with no outage of the existing production database while the conversion/migration effort occurs. The new system can be thoroughly tested with production data while the existing system remains in full operations. After testing is completed, Shadowbase ZDM allows the users to be phased over, or brought over all together, to a known-working new environment.

4. Audit Compliance and Data Utilities

Data replication is at the heart of the HPE Shadowbase product portfolio. However, it is also important to be able to monitor and if necessary correct that data in order to detect anomalous behavior, ensure continuation of proper business operations, and for audit compliance purposes. The Shadowbase product suite includes additional products to meet these needs.

HPE Shadowbase Audit Log (SAL)

Available for the HPE NonStop Server platform, Shadowbase Audit Log (SAL) creates a searchable archival database of transactional activity (e.g., inserts, updates, and deletes) on a reporting database for application change data auditing purposes. The reporting database can be created on the HPE NonStop Server system in either Enscribe (for Enscribe source data changes) or SQL (for SQL source data changes) formats, as well as on an off-platform target database such as Oracle or SQL Server.

When configured, SAL will replicate the HPE NonStop TMF audit trail database change activity to the Enscribe or SQL (NonStop target) or SQL (other platform target) target reporting database, storing the change information as a series of records or rows, one per each of the insert, update, and delete statements that were originally issued against the source database. This report allows the user to review, historically, what data was changed, and when, in their source database.

HPE Shadowbase Audit Reader (SAR)

HPE Shadowbase Audit Reader (SAR) analyzes and displays all audited database activities on HPE NonStop Server systems. It shows what application transactions did, and when they did it, to database files and tables, enabling investigation of how data is being changed. 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 system). SAR supports both Enscribe and SQL I/O events. It also shows both before and after I/O images, as well as "undo" for aborted transactions. SAR 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).

[12]For more information, please read our white paper, *Using HPE Shadowbase Software to Eliminate Planned Downtime via Zero Downtime Migrations*.

HPE Shadowbase Compare and Repair

Available for the HPE NonStop platform, Shadowbase Compare allows the user to compare a target Enscribe file or SQL table to its source, and reports on all discrepancies found between the two. (Discrepancies may occur, for example, if the user accidentally purges necessary audit trails, etc.) Shadowbase Compare is helpful for validating that a target database matches its source, and for satisfying regulatory requirements. It is meant for comparing HPE NonStop to HPE NonStop, like-to-like source/target environments where Expand is available between the nodes (for example, for BC environments). HPE Shadowbase Compare also contains a repair function to assist with correcting or repairing SQL/MP or SQL/MX database discrepancies between two tables. Shadowbase Repair works on any two tables in these databases, provided they are both of the same type (MP or MX) and does not require the Shadowbase replication engine. Repair is particularly useful in business continuity data replication environments to provide confidence (or certification) that the source and target data matches, and to repair the discrepancies when it does not.

HPE Shadowbase Database Recovery Software – Shadowbase UNDO

This Shadowbase utility allows the user to selectively "undo," or rollback, all of the changes made to a file or table (or set thereof) to a previous date/time. This option is useful, for example, to restore a database to an earlier/previous set of values using a selective "as of" approach. Shadowbase UNDO restores a corrupted database by undoing corrupting database changes while retaining correct database changes, thus leaving the database in a known, consistent, and current state. By maintaining an Undo Queue of changes that have been made to the database, Shadowbase UNDO can follow the Undo Queue in reverse time order to the initial point of corruption and can reverse any corrupting changes that have been made. The rollback of corrupted data can be accomplished by Shadowbase UNDO while the application continues its processing functions.[13]

HPE Shadowbase Database Recovery Software – Shadowbase REDO

Shadowbase REDO maintains a REDO Queue of all changes that have been made to the database. If a new version of an application is to be run, 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 before executing the system modification is created and saved. This step is often accomplished by un-mirroring the application database and saving one of the mirrors. The upgraded system is then run with the remaining mirror (and can be re-mirrored if sufficient additional disk space is available). If serious problems result in database corruption, the saved mirror can be restored quickly. Shadowbase REDO deletes corrupting changes from its REDO Queue and applies only the valid database changes from its REDO Queue to the saved mirror in order to roll it forward to a current and correct state.

5. Zero Data Loss (ZDL)

Reduce Outage Costs with Shadowbase 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 the other significant impacts such as loss of customer loyalty, negative publicity, regulatory violations, brand reputation and even threats to human health and safety.

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.

[13]For more information on Shadowbase Data Recovery Software, please read our white paper: *HPE Shadowbase Data Recovery Software – Shadowbase REDO and Shadowbase UNDO*.

Key Zero Data Loss Highlights:

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

Sales, Service, and Support

One of the hallmarks of the HPE Shadowbase Total Replication Solutions product suite is the support and service provided by Shadowbase Support, a team of experienced data replication and data integration experts working to meet your needs on any project, including business continuity, disaster recovery, and data and application integration. With many years of experience, Gravic professionals have accumulated a significant amount of detailed knowledge and capability, especially in complex areas such as configuring and managing HPE NonStop TMF, the Guardian File System, Enscribe, and NonStop SQL. Access to this deep knowledge base is a key part of the Shadowbase service offering. Support packages can be tailored to your needs, whether that is for local business hours-only support, or full 24x7x365 support.

HPE and Gravic, Inc. are strategic partners and offer Gravic Shadowbase global sales and support directly through the HPE organization. HPE licenses, services, and supports the award-winning HPE Shadowbase product suite for HPE NonStop and other servers. By providing a single point of purchase, HPE and Gravic are improving the overall customer experience. Customers benefit from the worldwide reach, industry expertise, and 24×7 support available from HPE, while HPE customers benefit from the wide range of unique features available with Shadowbase software.

The product suite is sold by HPE under the name, HPE Shadowbase. For more information, please contact your local HPE account team or visit our website.

International Partner Information

Global

Hewlett Packard Enterprise
6280 America Center Drive
San Jose, CA 95002
USA
Tel: +1.800.607.3567
www.hpe.com

Japan

High Availability Systems Co. Ltd
MS Shibaura Bldg.
4-13-23 Shibaura
Minato-ku, Tokyo 108-0023
Japan
Tel: +81 3 5730 8870
Fax: +81 3 5730 8629
www.ha-sys.co.jp

Gravic, Inc. Contact Information

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

Hewlett Packard Enterprise Business Partner Information

Hewlett Packard Enterprise directly sells and supports Shadowbase Solutions under the name ***HPE Shadowbase***. For more information, please contact your local HPE account team or visit our website.