



HPE Shadowbase Total Replication Solutions for HPE NonStop

A Gravic, Inc. White Paper



HPE Shadowbase Total Replication Solutions for HPE NonStop

Executive Summary

In today's business world, access to real-time online transactional data is a powerful competitive advantage. To realize the advantage, this data must be available at any time, all the time, and it must be current. The corollary to this advantage 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. 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 Total Replication Solutions for HPE NonStop is Gravic's real-time data replication and data integration solution for the HPE NonStop platform. The HPE Shadowbase product suite provides the means to meet the above requirements, via reliable low-latency, real-time data replication and distribution across heterogeneous systems and applications. With these powerful capabilities, 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.



For over four decades, Gravic has built low-latency, highly reliable data replication products for the demanding mission-critical HPE NonStop (formerly Tandem) marketplace. Many Fortune 500 companies worldwide trust their priceless data to HPE Shadowbase software for solving needs that range from asynchronous and synchronous [business continuity](#) solutions to homogeneous and heterogeneous [data integration](#) (fast data) solutions. The HPE Shadowbase Products Group provides solutions and services for hundreds of global customers/licenses in the HPE NonStop space. Gravic's extensive patent portfolio, *Breaking the Availability Barrier* book series, and comprehensive published white papers, articles, and case studies reinforce the expertise of Gravic's staff and the value of its technology to solve a wide range of business challenges.

HPE and Gravic, Inc. are strategic partners and offer HPE Shadowbase global sales and support directly through the HPE organization. HPE now licenses, services, and supports the leading-edge Shadowbase product suite for NonStop and Other Servers.

The key to the success of the Shadowbase solutions is its flexibility – its ability to provide continued value to our customers across a wide range of projects, solving a diverse set of business problems. Shadowbase solutions include business continuity, from uni-directional active/passive disaster recovery architectures to continuous availability active/active disaster tolerant architectures; data integration for feeding data warehouses, business intelligence systems, and online query processing (OLQP) reporting systems; and application integration for integrating operational processing with ancillary applications in real-time, event-driven architectures.

Unlike many traditional data replication and data integration products, Shadowbase software not only provides extremely low-latency replication between homogeneous databases and systems, it also provides extensive flexibility in selectivity, sophisticated data transformation and mapping, one-to-many or many-to-one configurations, replication between heterogeneous sources and targets, and true bi-directional replication between two or more live, production systems and their databases. Shadowbase replication provides data synchronization and integration across a wide variety of platforms and environments including NonStop, Unix, Linux, Windows, OpenVMS ([please inquire](#)), and others; and for a variety of databases including NonStop SQL, NonStop Enscribe, Oracle, Sybase, SQL Server, IBM Db2®, MySQL, and others.

This white paper describes features and uses for Shadowbase software in the HPE NonStop market. For additional information about Shadowbase solutions for other platforms, please see the white paper, [Shadowbase Total Replication Solutions for Other Servers](#).

Table of Contents

Executive Summary	2
HPE Shadowbase Solutions	5
HPE Shadowbase Platforms and Environments	6
HPE Shadowbase Architecture for HPE NonStop	7
Zero Data Loss	8
<i>HPE Shadowbase Queue Manager (QMGR)</i>	9
Native Calls	9
<i>Flexible, Heterogeneous</i>	10
<i>Heterogeneous Database Support</i>	10
<i>Selectivity</i>	10
<i>Sophisticated Data Transformation and Mapping</i>	10
Security, Data Encryption, Data Compression	11
One – One; One – Many; Many – One; Many – Many Architectures	11
Reliability, Availability, Scalability, and Manageability	11
<i>Reliability</i>	11
<i>Availability</i>	11
<i>Scalability</i>	12
Uni-directional and Bi-directional Replication	12
<i>Uni-directional Replication</i>	12
<i>Bi-directional Replication</i>	13
<i>Bi-directional Replication and Asynchronous Replication Data Processing</i>	13
Data Validation	14
<i>HPE Shadowbase Compare</i>	14
<i>HPE Shadowbase Repair</i>	14
Data Transformation	14
<i>Shadowbase API's</i>	15
<i>SBMAP</i>	15
<i>User Exits</i>	15
<i>HPE Shadowbase Essentials Bundle</i>	16
Case Studies	17
<i>One Bank's March Towards Active/Active</i>	17
<i>Large Telco Deploys Shadowbase Continuous Availability Architecture for Scale-Out to Handle Massive Growth in Smart Phone Usage</i>	17
<i>Telecom Italia's Active/Active Mobile Services</i>	17
<i>A Large Canadian Bank Dramatically Improves its ATM/POS Availability with HPE Shadowbase Solutions</i>	17
<i>HPE Shadowbase ZDM Achieves Zero Downtime Migration for Large Bank Datacenter</i>	17
<i>HPE Shadowbase Helps a Major ISP Migrate from Sybase to HPE NonStop with No Downtime</i>	17
<i>Major Bank Uses Active/Active with User Partitioning to Avoid Data Collisions</i>	17
Summary: The HPE Shadowbase Product Suite	18
Outstanding Service and Support	18
International Distributors	19
International Partner Information	20
Gravic, Inc. Contact Information	20

Table of Figures

Figure 1 – HPE Shadowbase Platforms and Environments	6
Figure 2 – HPE Shadowbase Overview	7
Figure 3 – Queue Manager (QMGR) Flow Chart	9
Figure 4 – HPE Shadowbase Customized Processing	10
Figure 5 – Various HPE Shadowbase Topologies	13
Figure 6 – Automatically APIs, SBMAP to Write SQL-like scripts, and Shadowbase User Exit Code Inserts and Customizations Enable Data Aggregation, Disaggregation, and/or Transformation.....	15

HPE Shadowbase Total Replication Solutions for HPE NonStop

HPE Shadowbase Solutions

The HPE Shadowbase suite of data replication and data integration products provides the following solutions:

- **Business Continuity¹**
 - Uni-directional active/passive *disaster recovery* for high availability
 - Bi-directional active/almost-active *sizzling-hot-takeover* for higher availability
 - Bi-directional active/active (*hot-hot*) disaster tolerant architecture for continuous availability
 - *Zero downtime migration* (ZDM) for eliminating planned downtime²
 - *Zero data loss* (ZDL) for eliminating data loss in the event of an outage, and eliminating data collisions in active/active architectures
 - *Shadowbase Compare* (see below) is a *must-have* add-on for validating that the target environment matches the source and is ready to take over on a moment's notice
- **Data Integration³ and Data Synchronization**
 - Loading operational data into a data warehouse, data mart or other ETL environment
 - Replicating change data capture information into a data warehouse, data mart or other ETL Environment (to keep the target data current/not stale while the source data is being updated)
 - Offloading reporting from the host system to create *online query processing* (OLQP) environments
 - Feeding a *real-time business intelligence* (RTBI) environment⁴
- **Application Integration⁵**
 - Integrating operational processing with ancillary systems to improve value-add and upsell opportunities
 - Building real-time, event-driven architectures based on database change processing
- **Data Validation**
 - *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
 - *Shadowbase Repair* can repair the discrepancies uncovered by *Shadowbase Compare*, and 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
- **Shadowbase Essentials**
- Shadowbase Essentials provides a suite of important utilities and tools for creating, managing, and operating your Shadowbase environments.
- Shadowbase Essentials is a must-have for business continuity environments, as well as those performing data and application integration replication.

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.

¹Refer to the Gravic white paper, [Choosing a Business Continuity Solution to Match Your Business Availability Requirements](#) for more information.

²Refer to the Gravic white paper, [Using HPE Shadowbase to Eliminate Planned Downtime via Zero Downtime Migration](#) for more information.

³Refer to the Gravic white papers, [HPE Shadowbase Streams for Data Integration](#) and [HPE Shadowbase Streams for Application Integration](#) for more information.

⁴Refer to the Gravic white paper, [Evolution of Real-Time Business Intelligence](#) for more information.

⁵Refer to the Gravic white papers, [HPE Shadowbase Streams for Data Integration](#) and [HPE Shadowbase Streams for Application Integration](#) for more information.

HPE Shadowbase Platforms and Environments

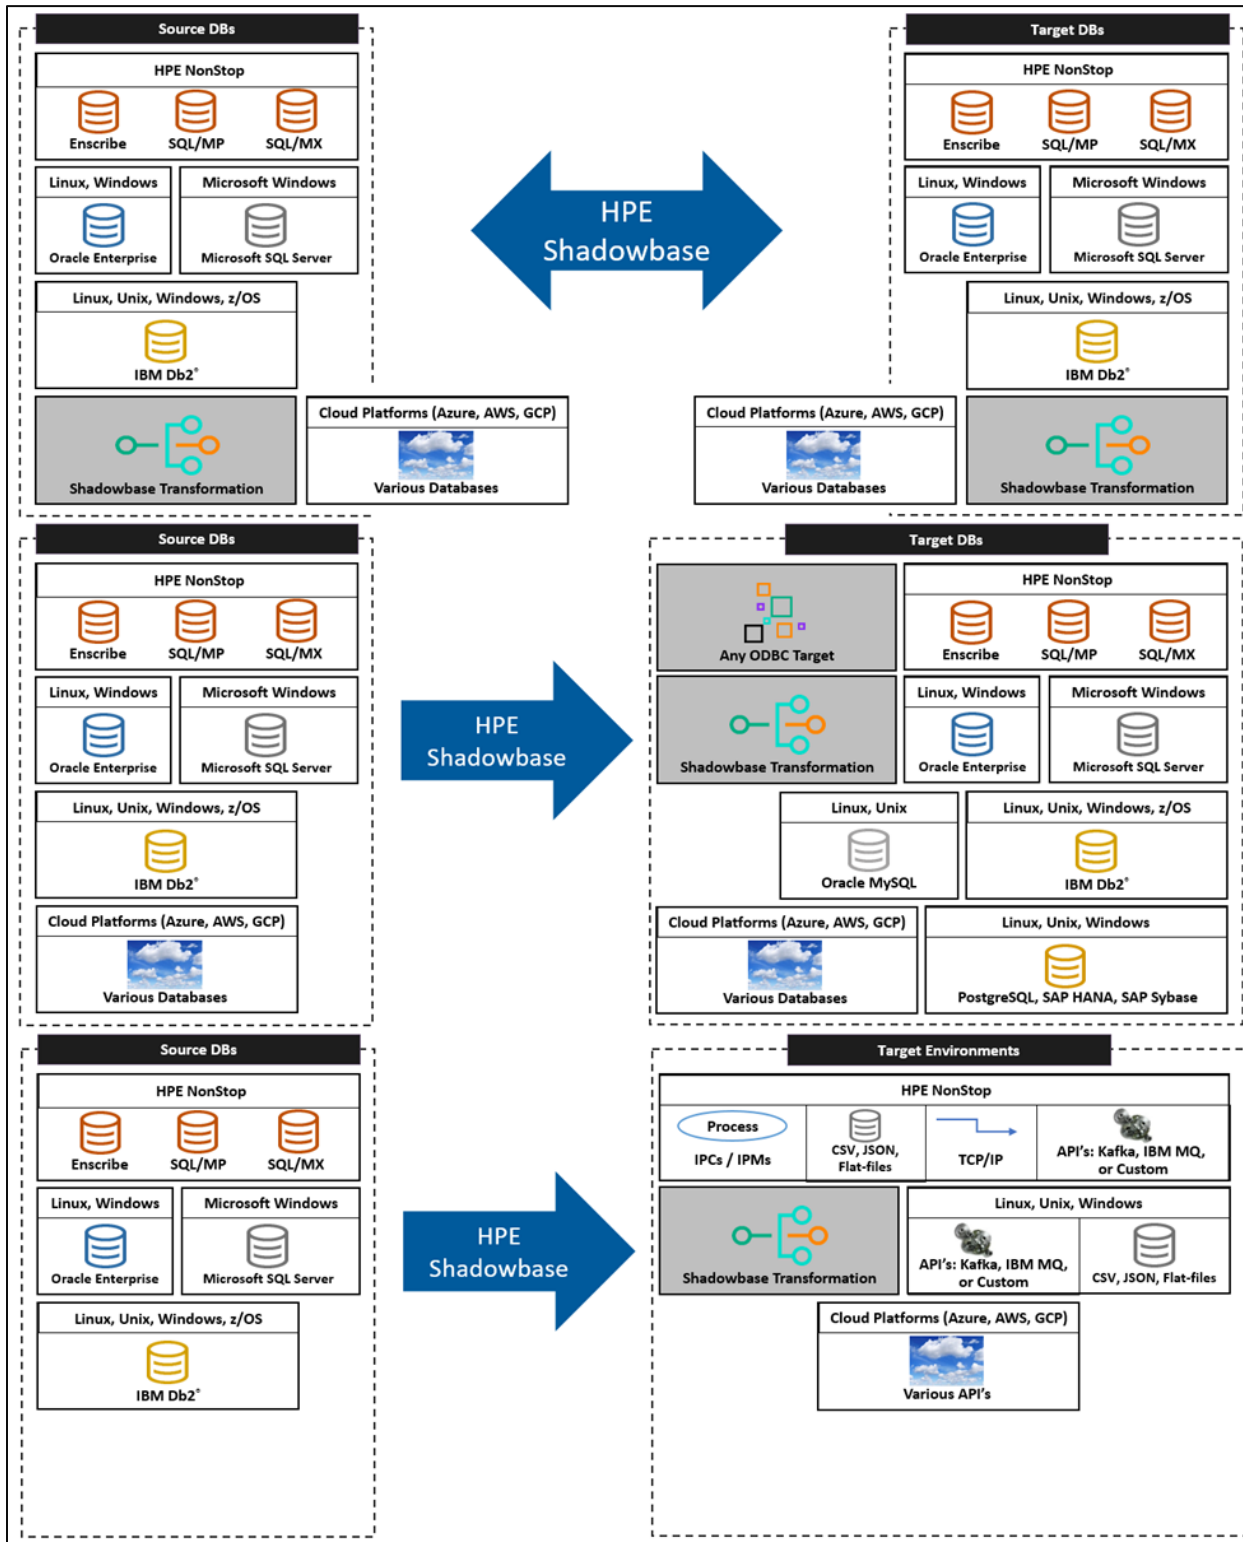


Figure 1 – HPE Shadowbase Platforms and Environments

Figure 1 depicts the source and target platforms, databases, and environments that HPE Shadowbase software supports for either uni-directional or bi-directional replication. (Note that all combinations above are tested and supported. [See our website](#) for the most up-to-date list of supported environments.)

For replication source environments, Shadowbase software supports⁶:

- HPE NonStop Enscribe, SQL/MP, and SQL/MX
- Oracle, SQL Server, Db2, and Sybase when running on Linux, Unix, or Windows environments

For replication target environments, Shadowbase software supports:

- HPE NonStop Enscribe, SQL/MP, and SQL/MX
- Oracle, SQL Server, Sybase, Db2, PostgreSQL, and MySQL when running on Linux, Unix, or Windows environments
- Any ODBC-compliant target database ([contact Gravic](#) for specific requests as a minor port may be needed depending on the ODBC client API/version available)

The source/target platform, database, and environment can be the same, or vastly different, as Shadowbase technology handles the mapping for homogeneous and heterogeneous data replication and data integration. This functionality is supported for both uni-directional as well as bi-directional replication.

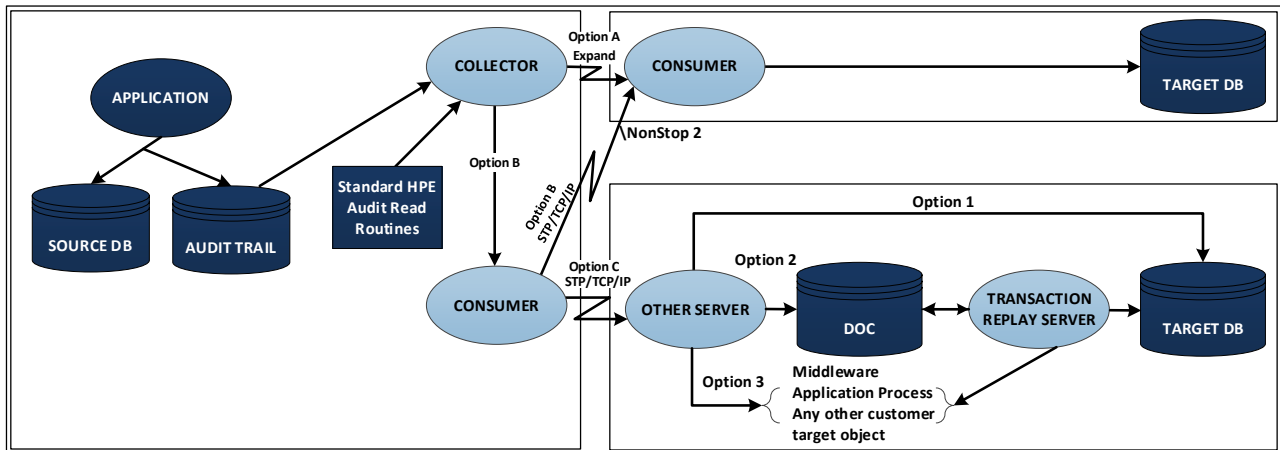


Figure 2 – HPE Shadowbase Overview

HPE Shadowbase Architecture for HPE NonStop

As shown in Figure 2, Shadowbase replication captures changes to NonStop SQL/MP, SQL/MX, and Enscribe databases from the NonStop TMF audit trail files through the HPE-provided audit read routines. This interface has been optimized to collect, filter, and batch the database change activity that is to be replicated to the target environment.

The changes captured by the Shadowbase Collector process are sent through an efficient interprocess message channel to one or more Shadowbase Consumer processes. These Consumer processes can be resident on the same node or on a remote node from the Collector.

When replicating directly from one NonStop node to another, the Collector can use EXPAND to communicate directly with the remote Consumer (Option A in Figure 2). Alternatively, when using TCP/IP for communications, the Collector communicates first with a local Consumer using interprocess messages, and this source-side Consumer then uses the STP Shadowbase protocol over TCP/IP socket connections to communicate with the remote Consumer process (Option B in Figure 2). In either case, the target Consumer will then apply the change data into the target database. When using TCP/IP, the NonStop nodes can have the same or different node names and node numbers.

When replicating from a NonStop source to an Other Server target (Option C in Figure 2) the source database changes captured by the Collector may be sent to one or more Consumer processes on the source system for replication over TCP/IP to any of the databases supported by Shadowbase architecture on another system. On these systems the user has the option of directly updating the target database with the I/O events as soon as they occur on the source database (Option 1 in Figure 2) or queuing the changes to disk in a target Database of Change (DOC) file and applying only committed transactions (Option 2 in Figure 2). Additionally, instead of

⁶All source environments are bi-directional.

having Shadowbase technology replay the replicated events into the target file, the Shadowbase infrastructure provides a method to consume the events by sending them to non-Shadowbase processes, feeding middleware adapters, or consuming them via customer-provided application(s) outside of the Shadowbase architecture (Option 3 in Figure 2.)

Note: For a *direct* configuration (Option 1), the source database changes are replicated and applied as soon as they occur, meaning both committed and aborted transactions are replayed. However, this configuration provides the lowest overall replication latency. In a DOC configuration (Option 2 and Option 3), the source database changes are replicated to the target system and stored in the DOC as soon as they occur, but are then only replayed into the target database after the commit occurs. In this case, only committed transactions are replayed.

Note: The Shadowbase architecture provides data manipulation language (DML), such as insert, update, and delete and data definition language (DDL), such as create capture for TMF-audited source files and tables. If the source files and tables are non-audited (i.e., the database change activity is not captured by TMF), Shadowbase architecture cannot capture the individual DML or DDL activity unless you implement the NonStop AutoTMF capability to convert the non-audited application into a TMF-audited application. Alternatively, for replicating non-audited source files and tables, the technology provides two options: SOLV “Snap-Shot” Loads/Refreshes, and the *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. Using the *Shadowbase File Chaser* technology, Shadowbase software can “chase” the EOF of non-audited log files/tables and inject those events into the replication stream. [Contact Shadowbase Product Management](#) for additional information about these non-audited solutions.

Zero Data Loss⁷

Shadowbase 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, 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 with synchronous replication is the solution for the most mission-critical applications, where data loss and/or data collisions cannot be tolerated.⁸

⁷ Disclaimer: This product does not guarantee that you will not lose any data; all user warranties are provided solely in accordance with the terms of the product License Agreement. Each user's experiences will vary depending on its system configuration, hardware and other software compatibility, operator capability, data integrity, user procedures, backups and verification, network integrity, third party products and services, modifications and updates to this product and others, as well as other factors. Please consult with your supplier and review our License Agreement for more information.

⁸ Please visit [the website](#) for further information on Shadowbase ZDL and ZDL+.

HPE Shadowbase Queue Manager (QMGR)

As previously shown in Figure 2 when replicating data from a NonStop source to a NonStop target, the change data is replicated directly to the target environment and is applied directly into the target database. In other words, after it is read from the source system's TMF audit trails, the changes do not hit disk again until actually being applied into the target database. This approach is the most efficient and has the lowest overall replication latency.

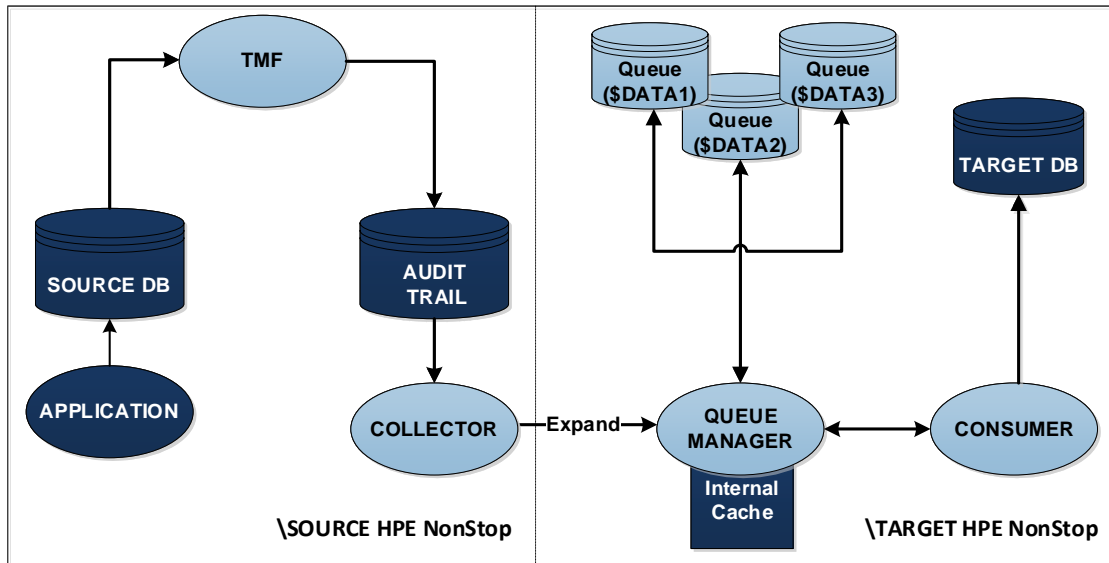


Figure 3 – Queue Manager (QMGR) Flow Chart

There are cases, however, where it is clearly preferred to queue the data being replicated either on the source system (sending side) or on the target system (receiving side); the Shadowbase Queue Manager (QMGR) supports both. For example, if the customer's network is slow or unreliable, the queuing can be done on the sending side. In other cases, if it is important for business continuity to have as much data as possible on the target side, then Shadowbase replication provides an optional target QMGR subsystem (Figure 3). The QMGR works similarly to the TMF audit trails. It produces a series of unstructured disk files that hold the sequence of data to be replicated. When the QMGR is configured, Shadowbase replication flushes the change data being replicated to disk in specially formatted *Queue Files* at the same time as it applies those changes to the target database. This action allows Shadowbase software to separate the delivery of the replicated data to the target environment from applying the replicated data into the target database. In other words, if the target database goes offline during replication, then Shadowbase software continues to deliver the data into the Queue Files until the target database comes back online again. When that occurs, Shadowbase replication automatically restarts applying the queued data into the target database using the Queue Files while continuing to receive data from the source. This approach is very common when performing a [Shadowbase Zero Downtime Migration](#).

Note that the QMGR subsystem can be selectively enabled on a replication "channel" (replication thread), meaning that some channels can have QMGR processing while others do not have it. Also note that the QMGR works in parallel with the Consumer applying the events to the target database and it is optimized for bulk I/O operations. Adding the QMGR onto a replication channel does not typically increase that channel's replication latency. In other words, it is faster for the QMGR to flush the bulk blocks of data into the Queue Files than it is for the Consumer to extract the events and position/apply them into the target database. This fact means that adding in QMGR processing does not "slow down" the rate of applying the events into the target database. However, adding the QMGR does add overall system overhead as the blocks of data are written, and possibly re-read, from the Queue Files during processing.

Native Calls

Each and every Shadowbase interaction with a target database is written using efficient native database calls specific to that DBMS. The focus is on speed and efficiency. Note that for some databases, however, a generic ODBC interface is also provided (for example, ODBC is now a native interface for Microsoft SQL Server). In

these cases, Shadowbase replication has enhanced its patented optimization layer (called *statement caching*) to make using that interface as efficient as possible.

Flexible, Heterogeneous

Unlike other NonStop-based replication solutions, Shadowbase replication is an extensible tool. Over 35 years of experience in designing low latency, real-time replication solutions have gone into Shadowbase NonStop solutions, producing a highly optimized, very efficient replication tool with high volume sub-second source database to target database latency. Where Shadowbase solutions excel, however, is in its flexibility and platform support. Our years of experience have taught us that most replication projects demand more than a solution that simply keeps two identical databases in sync.

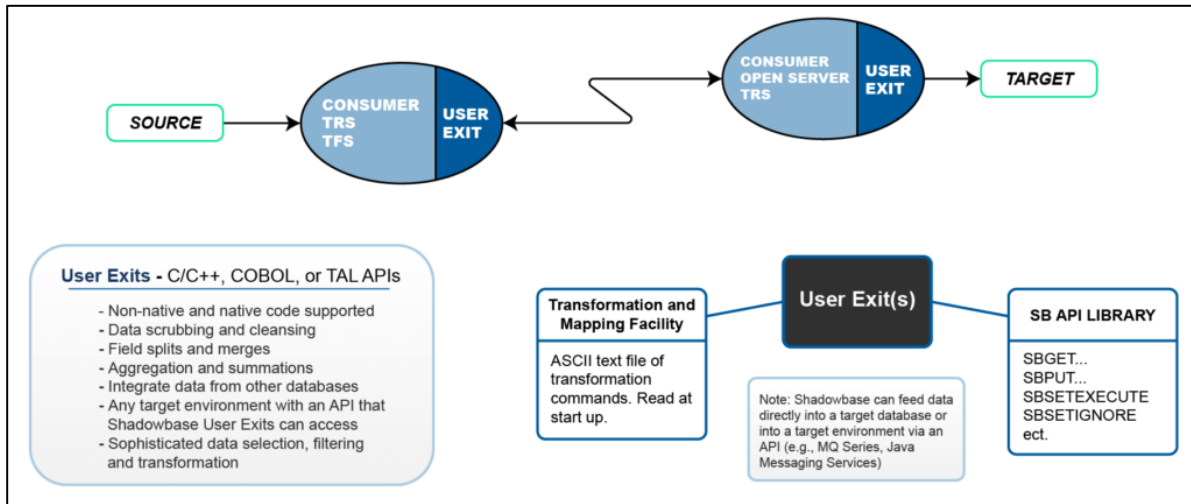


Figure 4 – HPE Shadowbase Customized Processing

Heterogeneous Database Support

Out of the box, in addition to supporting all of the NonStop databases, Shadowbase for NonStop provides replication to a variety of additional database management systems, including Oracle, SQL Server, Sybase, Db2, and MySQL. In all cases, the source and target schemas and data formats may be the same, or vastly different.

Additional or custom source databases or data feeds are supported through custom input APIs built into Shadowbase software, and additional or custom targets are supported through Shadowbase user exits or through the provided SBDOCRD shell program (SBDOCRD provides a shell program for reading the Shadowbase DOC files and replicating those events into other environments not directly supported by Shadowbase technology). Shadowbase Product Group Support is available to assist customers interested in these additional/extensible approaches. Figure 4 illustrates where the input APIs and user exits are driven in the data flow process.

Selectivity

Shadowbase replication supports replicating entire files or tables or only subsets of the files or tables – specific columns, rows, or even fields. Some subsets can be replicated to one target with other subsets replicated to other targets simultaneously.

Sophisticated Data Transformation and Mapping

Shadowbase solutions empower the user with multiple options for transforming data in-flight or performing sophisticated mapping of data between heterogeneous sources and targets. The Shadowbase Transformation and Mapping Utility provides a scripting tool that enables users to perform many unique data mapping functions and some simple data transformations. For more powerful transformations, Shadowbase replication provides APIs in source and target processes for users to add custom written code called *user exits* into the replication process. These compiled code modules allow the user to otherwise extend the replication engine's default processing to encapsulate additional logic beyond what otherwise is available.

User exits written in “C” are supported on all platforms, and on some platforms, users may choose to write the user exits in Cobol, C++, Java, TAL/pTAL, or other languages. Through the use of these user exits, users can perform routines like splitting fields, merging fields, adding data from other tables in-flight, or performing aggregations or summations, data content filtering, and/or data obfuscation (e.g., removing plain-text sensitive data and replacing it with obfuscated alternative values). Generally, the programmer has the power to perform any/all sophisticated data manipulation and other I/O operations that the host language provides to manipulate the replicated data in-flight.

Security, Data Encryption, Data Compression

Shadowbase software supports encrypting the session traffic for TCP/IP connections by using proxy servers. The proxy servers must be procured from HPE or another third-party security vendor. (They come standard in most current NonStop releases or can be downloaded from the OpenSSL site.) For NonStop to NonStop environments using the newer NonStop IP CLIM architectures, Shadowbase replication also supports the use of the IP-SEC layer for CLIM to CLIM communication.

Additionally, through the use of the Shadowbase user exit extensibility feature, the user could encrypt/decrypt, and perform sophisticated data compression, on the data being replicated. Of course, data encryption or compression could always be performed external to the NonStop system via encrypting routers.

One – One; One – Many; Many – One; Many – Many Architectures

In Shadowbase replication, the above configurations apply not only at the database level, but also at the database, table, column, row, and field levels. In other words, the Shadowbase solutions provide the user the power to not only replicate one database to multiple target copies or merge multiple source databases into one target database, but to break out one field in the source database into multiple fields in the target or merge data from multiple tables in the source into one table in the target. The user is empowered to design target databases that are best suited to the needs of the project – Shadowbase replication handles getting the right data into the right place, in the right order.

Reliability, Availability, Scalability, and Manageability

Born in the 24x7, mission-critical world of NonStop computing, reliability, availability, and scalability are the fundamental principles to which Gravic-developed software adheres. This commitment to excellence in the software we develop makes the HPE Shadowbase Total Replication Solutions for HPE NonStop a technical leader in the NonStop marketplace. On NonStop, Shadowbase software implements an architecture logically similar to NonStop TS/MP, using process pairs and persistence monitoring of replicated Shadowbase system processes with workload distributed across CPUs to achieve availability and scalability.

Reliability

Reliability is paramount in a replication product. It is absolutely vital to the customer that the data in the target is accurate. Many of the world’s preeminent exchanges, banks, securities trading firms, financial switches, and telecommunications companies trust their data to Shadowbase software each and every day, 24x7x365. Shadowbase replication is designed to not lose transactions or data, and to recover fully and automatically when serious faults occur, by keeping a persistent copy of key information such as the replication restart position that is accessible across shutdowns, restarts, and failure recovery (e.g., if a CPU crashes). Shadowbase replication is built on NonStop fault tolerant process principles, using a fault-tolerant process monitor to maintain process persistence for the rest of the Shadowbase environment.

Availability

Simply stated, uptime is vital to a replication product, and when a target is down or offline, it is of paramount importance that the product recovers seamlessly where it left off and quickly catches up when that target becomes available again. Upon the loss of a critical component, Shadowbase processes can be configured to be automatically restarted, and techniques like clustering, fault-tolerant processes, or persistent process pairs in backup CPUs are employed to ensure that the product is always available and functioning. When the network or a target database is down or offline, Shadowbase replication continually monitors the network and/or the target, remembering where it left off so that when the network or target is once again available, replication will catch up from the point it left off with minimal or no user involvement.

All of the key Shadowbase components support multiple processes or paths, such that if any key component becomes unavailable or unusable, Shadowbase replication will restart or route around the problem and continue as soon as possible. On the NonStop, this action is accomplished using fault-tolerant process pairs for the Shadowbase monitors, persistent processes for the core replication components, and multiple path support for resolving network issues.

An additional indicator of availability is the ability of the replication subsystem to coexist non-intrusively with the source application environment. It is unacceptable for an asynchronous replication subsystem to take the source application or database environment *offline*, or prevent the source system from satisfying user requests – locking up the source database or environment, freezing all access ports, or other intrusive action in the face of a replication subsystem error. Shadowbase solutions are designed to avoid such devastating source outages, and to recover the target automatically from most faults when possible.

Scalability

Some of the world's largest, most active production databases are replicated using Shadowbase software, made possible by its ability to scale to meet the throughput demands of the project. Shadowbase functions scale to support multiple components at each step of the replication processing, from extract, to distribution/transmission, to delivery/application of the changes into the target environment. For example, Shadowbase replication supports parallel source database change data extraction paths, parallel network transmission paths, and parallel appliers into the target database. Data transmission is optimized and blocked for windowing to send fewer, fuller packets, using a guaranteed delivery protocol. Shadowbase system processes are replicated across CPUs on the NonStop platform, and workload is balanced between them for scalability and optimum CPU utilization and throughput. Shadowbase processing is based on a real-time event-driven architecture, and avoids polling whenever possible.

Uni-directional and Bi-directional Replication

Shadowbase replication has the flexibility to be configured for uni-directional (one-way) replication (often used for active/passive business continuity or for feeding operational data into a data warehouse), as well as bi- and multi-directional replication (often used for active/active business continuity and many-way application scale-out). In uni-directional replication, Shadowbase software replicates the source data to one or more target environments. In this configuration, the target is generally read-only (i.e., is not being updated by application processes) and is often useful for offloading reporting and query processing from the host.

Shadowbase software also supports bi-directional and multi-directional replication, including route through architectures (where all nodes are not directly connected to each other) so that replication environments match the users' communications topology. In these architectures, all environments are typically actively being updated, with Shadowbase replication keeping all of the database environments synchronized with each other.

Uni-directional Replication

Gravic has been deploying uni-directional replication solutions for over 45 years. Shadowbase software supports all of the possible uni-directional replication configurations, including: parent/subordinate (simplex or peer-to-peer), one – many, many – one, cascade (route through), multiple contingency, ring, and reciprocal. Though support for all of these various configurations is important, the keys to uni-directional replication are simple – *replication latency* and *replication throughput*. Figure 5 shows these various Shadowbase topologies.

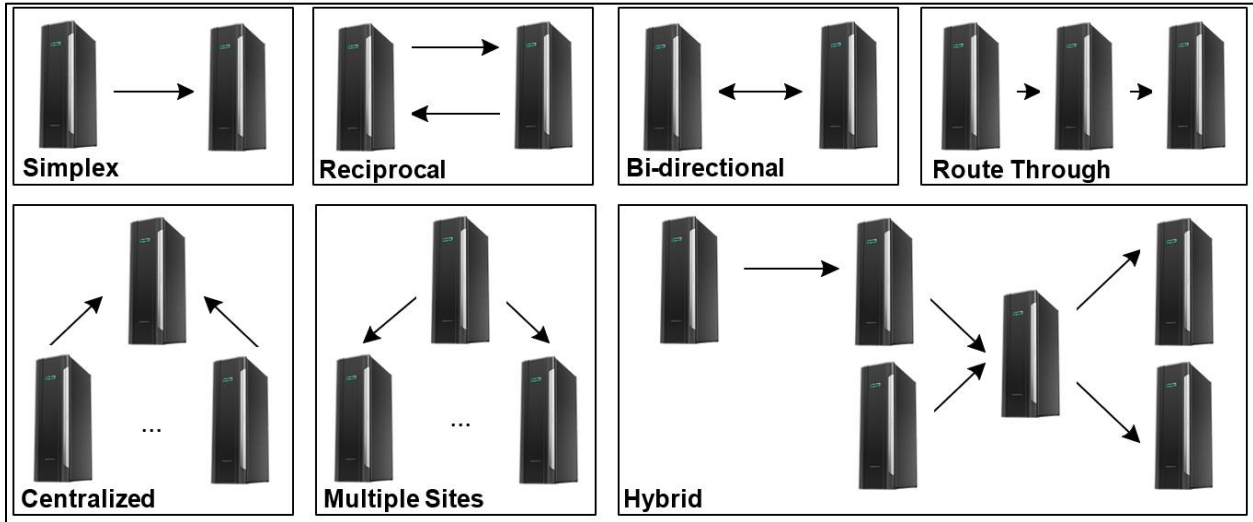


Figure 5 – Various HPE Shadowbase Topologies

Replication latency is defined as the lag time between the time the production application writes a database event (e.g., an insert, update, or delete) to disk in the source database and the time the replication engine applies the same event into the target database. Shadowbase replication excels at providing extremely low latency, with customers often reporting latency times measured in milliseconds.

Different from latency, replication throughput is measured by how much data a replication product pushes from the source to the target at sustainable speeds. The use of parallelism, and advanced replication algorithms like the Shadowbase patent-pending *statement caching* algorithm, enable it to scale to meet the throughput demands of the largest NonStop customers. Any of the databases listed above that Shadowbase solutions support as a source can be combined with any other in this configuration.

Bi-directional Replication

Shadowbase replication utilizes patented and patent-pending technology to perform bi-directional replication. Not to be confused with *reciprocal replication* (where two systems each have an active application running with the databases being replicated but not being shared across the two applications), bi-directional replication is defined as replication between two or more live databases on the same or different nodes. Multi-directional replication is an extension of bi-directional replication where there are more than two replication environments that are interconnected with the databases being shared.

Shadowbase technology takes the bi-directional replication concept one step further by enabling replication not only between homogeneous databases (e.g., for business continuity), but also between heterogeneous databases (e.g., for integrating an operational system with an ancillary system). Any of the databases listed above that Shadowbase solutions support as a source can be combined with any other in this configuration.

Bi-directional Replication and Asynchronous Replication Data Processing

There are two important elements to bi-directional replication, and Gravic won patents in both areas. These elements are *ping-pong* (data oscillation) avoidance and data collision avoidance versus data collision detection and resolution.

Ping-pong avoidance, also referred to as data or transaction oscillation or looping avoidance in bi-directional replication, occurs when a replication engine replicates an application database I/O event from system 1 to system 2, and then errantly replicates it back from system 2 to system 1 and so on. Shadowbase software avoids ping-pong out-of-the-box using patented Gravic technology.

The more difficult component to bi-directional replication occurs when data collisions are possible, i.e., the exact same row in both databases is updated at exactly or nearly the same time by the application. For some applications, this occurrence is not possible and hence is not an issue. For others, however, it is innate, particularly for the more sophisticated load-balanced active/active architectures using the “route anywhere” model. The route anywhere model allows any transaction to be routed to any node for processing.

Gravic works with customers to help them partition database feeds and/or the databases themselves to avoid these potential collisions. If this option is not possible, Shadowbase solutions also exist to resolve collisions after they have occurred by identifying when they occur and using business rules to resolve them.

Unfortunately, for some businesses, however, these solutions are not viable options. For example, they do not prevent an account from being simultaneously closed and all funds withdrawn at two separate branches at the same time, which is why Gravic invented *cooperative processing*. Cooperative processing is based on patented Gravic technology that enables Shadowbase replication to detect and avoid potential collisions before they occur using a high-performance form of synchronous replication called *coordinated commits*. Development on this solution is underway, and a cooperative processing version of Shadowbase solutions for NonStop (as well as other databases) should be available in the future. [Contact Gravic](#) for more information on this powerful solution.

Data Validation

HPE Shadowbase Compare

HPE Shadowbase Compare is a solution that compares two HPE NonStop databases (Enscribe and/or SQL) and reports whether or not they match. Mismatches, called discrepancies, may occur for example, if either a disaster corrupts one of the databases, or a user misconfigures the application or replication environment. Applications may remain active on both the source and target, even for updating, while the comparison takes place. Shadowbase Compare validates that a target database matches its source database, and may be used to help satisfy regulatory and audit requirements for data validation and verification. Shadowbase Compare can compare both TMF-audited and non-audited files and tables. It can perform the comparison while any data replication product is replicating between the databases, for example to validate an online business continuity replication configuration. It is designed for comparing like-to-like environments, such as HPE NonStop to HPE NonStop source/target replication environments (e.g., for Enscribe-to-Enscribe or SQL-to-SQL business continuity). The files or tables can be located on the same system, or when on separate systems, Shadowbase Compare works with both TCP/IP and Expand connectivity.

HPE Shadowbase Repair

HPE Shadowbase Repair is a function in the HPE Shadowbase Compare for SQL NonStop solution 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). It does not require a Shadowbase replication configuration and can be used in any environment where a source and target table need to be verified. 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.

Data Transformation

HPE Shadowbase software can transform data between source and target database formats automatically, via scripts, or custom coding.

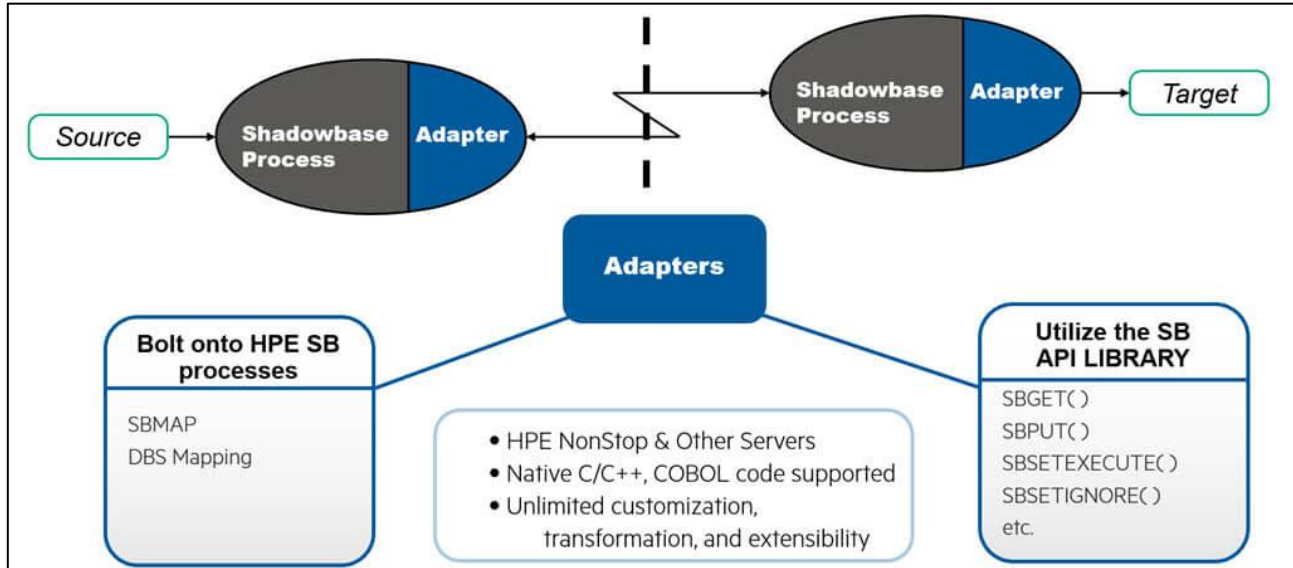


Figure 6 – Automatically APIs, SBMAP to Write SQL-like scripts, and Shadowbase User Exit Code Inserts and Customizations Enable Data Aggregation, Disaggregation, and/or Transformation

Shadowbase API's

Sophisticated transformation, scrubbing, and innumerable filtering options are available and configurable out-of-the-box.

SBMAP

Enables the user to write SQL-like scripts for sophisticated data transformation, scrubbing, and filtering.

User Exits

Enables the user to program scripts in C/C++ along with other languages and embed them in Shadowbase processes.

HPE Shadowbase Essentials Bundle

HPE Shadowbase Essentials Software Bundle includes several solutions for data transformation, management, compliance, monitoring, and controlling, and is sold separately from the main change data capture (CDC) replication engine. Data can be aggregated, disaggregated, and transformed as follows:

- **SBMAP (Shadowbase Map)**
 - Provides a powerful, sophisticated and extensible SQL-like scripting language that informs the Shadowbase engine how to transform source data into target data formats
- **SBCREATP (Shadowbase Generate CREATE Table Statements from SQL/MP DDL)**
 - Converts the schema structure of an SQL/MP table into another SQL variant
 - 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 (Shadowbase Generate CREATE Table Statements from Enscribe DDL)**
 - Reads an HPE Enscribe DDL record definition, simplifying the replication of unstructured Enscribe data into structured SQL databases
 - Includes features to allow manipulation of the source fields when creating the target columns, including dropping and renaming fields and normalizing the non-normalized Enscribe data
- **SB DCR (Shadowbase DDL Command Replication)**
 - Enables data definition language (DDL) command replication by replicating DDL command changes (e.g., source table schema changes) and applies them to the target database to keep the its structure consistent with the source database structure
- **SAL (Shadowbase Audit Log)**
 - Permanent archive to research and review what application data was changed and when
- **SAR (Shadowbase Audit Reader)**
 - Optional query tool that reads TMF Audit Trails using an SQL-like syntax
 - Interactively displays the contents of the NonStop audit trails using a powerful SQL-like syntax
 - Reads an event from the audit trails, then locates the corresponding DDL definition (for Enscribe files) or schema (for SQL/MP and SQL/MX tables) and accordingly formats the audit trail event data
 - Supports selective (i.e., WHERE CLAUSE) filtering, event-type filtering, and file/table data formatting/filtering
- **ETL Toolkit (Extract, Transform, and Load Toolkit)**
 - Converts NonStop Enscribe or SQL database changes or data loads into a variety of target formats, including flat files that are available for subsequent loading and transformation into another environment
 - Easily shares NonStop data across the enterprise
 - Provides adapters to load data directly into KAFKA and IBM MQ
- **SB UNDO & REDO**
 - NonStop-based utilities that selectively correct erroneous database changes
 - Maintains business continuity by correcting major errors or attacks without an application outage
 - Selectively undoes or roll-backs all of the changes made to a file or table (or set) to a previous date/time
 - Restores a database to an earlier/previous set of values using an “as of” approach
 - Reverts the changes made by a malware program which was accidentally deployed
 - UNDO does not require the source file/table (or application) to be offline
- **SBLSLOCK (Shadowbase Lock Collision Report)**
 - Enables investigation into what process or app owns the blocking transaction’s locks by immediately analyzing each lock’s base file/table partitions and all alternate key/index partitions, and then lists each lock’s object type, partition, transaction ID, and key values
 - Reports issues Shadowbase encounters when trying to update a file/table
 - Removes impediments to replication updating the target database, helping ensure the source and target databases remain synchronized
 - Enables remediating potential issues interrupting Business Continuity protection
- **SBLSTRAN (Shadowbase Transaction Report)**
 - Simplifies investigating long-running transactions, i.e., who started them and how long they were running
 - Reports restart point for ensuring all change data is properly captured and replicated from that point on
 - Lists relevant transaction information

Case Studies

In this section we provide examples of the many ways in which customers are taking advantage of Shadowbase capabilities to benefit their business, including business continuity, zero downtime migration, and application and data integration:

[*One Bank's March Towards Active/Active*](#)

[*Large Telco Deploys Shadowbase Continuous Availability Architecture for Scale-Out to Handle Massive Growth in Smart Phone Usage*](#)

[*Telecom Italia's Active/Active Mobile Services*](#)

[*A Large Canadian Bank Dramatically Improves its ATM/POS Availability with HPE Shadowbase Solutions*](#)

[*HPE Shadowbase ZDM Achieves Zero Downtime Migration for Large Bank Datacenter*](#)

[*HPE Shadowbase Helps a Major ISP Migrate from Sybase to HPE NonStop with No Downtime*](#)

[*Two Merged Retailers Integrate HPE NonStop SQL and Oracle RAC*](#)

[*Major Bank Uses Active/Active with User Partitioning to Avoid Data Collisions*](#)

Summary: The HPE Shadowbase Product Suite

HPE Shadowbase software solutions provide the facilities to track changes to a source database in real-time. It acts as an agent for remote systems or databases to distribute these changes to them. Shadowbase replication monitors database changes via several mechanisms, such as transaction logs, application logs, and database triggers.

Shadowbase replication supplies several means for distributing data:

- It replicates database changes to other heterogeneous databases via change data capture (CDC) in real-time, or on a scheduled snap-shot or micro-batch updating basis.
- It replicates point-in-time snapshots (such as for a key range) to other heterogeneous databases.
- It sends data changes as they occur to an application or a server or other target database environment.
- It makes data changes available to other applications that otherwise would have to poll for the data.
- It supports publish/subscribe architectures to send data to only those applications that have subscribed to the data.

The Shadowbase product suite includes the following families of products:

HPE Shadowbase Business Continuity allows enterprises to run advanced active/active solutions to maintain application uptime and continuous data processing (CDP), which requires minimizing downtime — both planned and unplanned. HPE Shadowbase data replication software 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.

HPE Shadowbase Streams for Data and 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, and 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.

HPE Shadowbase Compare – verifies replication software products to ensure consistency between databases even while they are actively running.

HPE Shadowbase Essentials Software Bundle includes several solutions (listed above) for data transformation, management, compliance, monitoring, and controlling, and is sold separately from the main change data capture (CDC) replication engine.

Shadowbase software is customizable by embedding any business processing logic that may be needed to satisfy an application need. The innovative Shadowbase technology has been awarded over two dozen patents, and many patent applications are pending.

In today's business world, access to real-time online transactional data is a powerful competitive advantage. To realize the advantage, this data must be available at any time, all the time, and it must be current. The corollary to this advantage 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. 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.

The HPE Shadowbase product suite provides the means to meet these requirements, via reliable low-latency real-time data replication and distribution across heterogeneous systems and applications. With these powerful capabilities, 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.

Outstanding Service and Support

One of the hallmarks of the HPE Shadowbase Product Suite is the support and service provided by the Shadowbase Support organization, a team of specialists who are available to help all Shadowbase customers.

Support packages can be tailored to the customer's needs, whether that is for local business hours only support, or full 24x7x365 support.

International Distributors

HPE and Gravic, Inc. are strategic partners and offer HPE Shadowbase global sales and support, directly through the HPE organization. HPE licenses, services, and supports the leading-edge Shadowbase product suite for NonStop and Other Servers. By providing a single point of purchase, HPE and Gravic are improving the overall customer experience. Our customers will benefit from the worldwide reach, industry expertise, and 24x7 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, [visit our website](#), and/or see the Shadowbase international partner list and contact information on the next page. Local time-zone distributors are available around the world to provide additional hands-on service and support.

We are also continuously interested in licensing our technology to more resellers and OEMs that want to embed Shadowbase components into their products. We have a long and highly successful track record of embedding our technology into ISV and customer applications to provide customized replication services. To discuss this approach, or for more information on our technological advancements, please [contact us](#).

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 **Shadowbase**. For more information, please contact your local HPE account team or [visit our website](#).

Copyright and Trademark Information

This document is Copyright © 2007, 2018, 2020, 2024 by Gravic, Inc. Gravic, Shadowbase and Total Replication Solutions are registered trademarks of Gravic, Inc. All other brand and product names are the trademarks or registered trademarks of their respective owners. Specifications subject to change without notice.

NOTICE: Shadowbase Zero Data Loss (ZDL) does not guarantee that you will not lose any data; all user warranties are provided solely in accordance with the terms of the product License Agreement. Each user's experiences will vary depending on its system configuration, hardware and other software compatibility, operator capability, data integrity, user procedures, backups and verification, network integrity, third party products and services, modifications and updates to this product and others, as well as other factors. Please consult with your supplier and review our License Agreement for more information.