



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 three 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 on-line 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
HPE Shadowbase Queue Manager (QMGR)	8
Native Calls	8
<i>Flexible, Heterogeneous</i>	9
<i>Heterogeneous Database Support</i>	9
<i>Selectivity</i>	9
<i>Sophisticated Data Transformation and Mapping</i>	9
Security, Data Encryption, Data Compression	10
One – One; One – Many; Many – One; Many – Many Architectures	10
Reliability, Availability, Scalability, and Manageability	10
<i>Reliability</i>	10
<i>Availability</i>	10
<i>Scalability</i>	11
<i>Manageability</i>	11
Uni-directional and Bi-directional Replication	12
<i>Uni-directional Replication</i>	12
<i>Bi-directional Replication</i>	12
<i>Bi-directional Replication and Asynchronous Replication Data Processing</i>	13
Zero Data Loss	13
Case Studies	14
<i>One Bank's March Towards Active/Active</i>	14
<i>Large Telco Deploys Shadowbase Continuous Availability Architecture for Scale-Out to Handle Massive Growth in Smart Phone Usage</i>	15
<i>Telecom Italia's Active/Active Mobile Services</i>	16
<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>	18
<i>HPE Shadowbase Helps a Major ISP Migrate from Sybase to HPE NonStop with No Downtime</i>	19
<i>Two Merged Retailers Integrate HPE NonStop SQL and Oracle RAC</i>	20
<i>Major Bank Uses Active/Active with User Partitioning to Avoid Data Collisions</i>	21
Summary: The HPE Shadowbase Product Suite	22
HPE Shadowbase Options	23
<i>HPE Shadowbase Audit Reader (SAR)</i>	23
<i>HPE Shadowbase Compare</i>	23
<i>HPE Shadowbase UNDO</i>	23
Outstanding Service and Support	23
International Distributors	23
International Partner Information	24
Gravic, Inc. Contact Information	24

Table of Figures

Figure 1 – HPE Shadowbase Platforms and Environments	6
Figure 2 – HPE Shadowbase Overview	7
Figure 3 – Queue Manager (QMGR) Flow Chart	8
Figure 4 – HPE Shadowbase Customized Processing	9
Figure 5 – Sample HPE Shadowbase Enterprise Manager (SEM) GUI Display.....	11
Figure 6 – Various HPE Shadowbase Topologies	12
Figure 7 – Disaster Tolerant Banking System	14
Figure 8 – The Telco Smart Phone Billing and Provisioning System	15
Figure 9 – TIM's INS Service.....	16
Figure 10 – The Bank's New Active/Active Architecture	17
Figure 11 – The Bank's Tri-node Active/Active ATM/POS Application During the Migration Phase.....	18
Figure 12 – The Challenge of Migration	19
Figure 13 – Integrated Refunds and Exchanges	20
Figure 14 – Active/Active Replication with User Partitioning to Avoid Data Collisions.....	21

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
- **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
- **Compliance Reporting and Resolution⁶**
 - Report on application transactional and event updates
 - Generate an archive database of the application's database change activity (recording what was done, and when, to your database)
 - Compare a target database to and source database, reporting on any inconsistencies and differences
 - *UNDO* erroneous database changes to roll a database (or a portion thereof) back to a previous state while the application remains active and the database remains online

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.

⁶For a complete list of Shadowbase products, refer to the web page: ShadowbaseSoftware.com/products/.

HPE Shadowbase Platforms and Environments

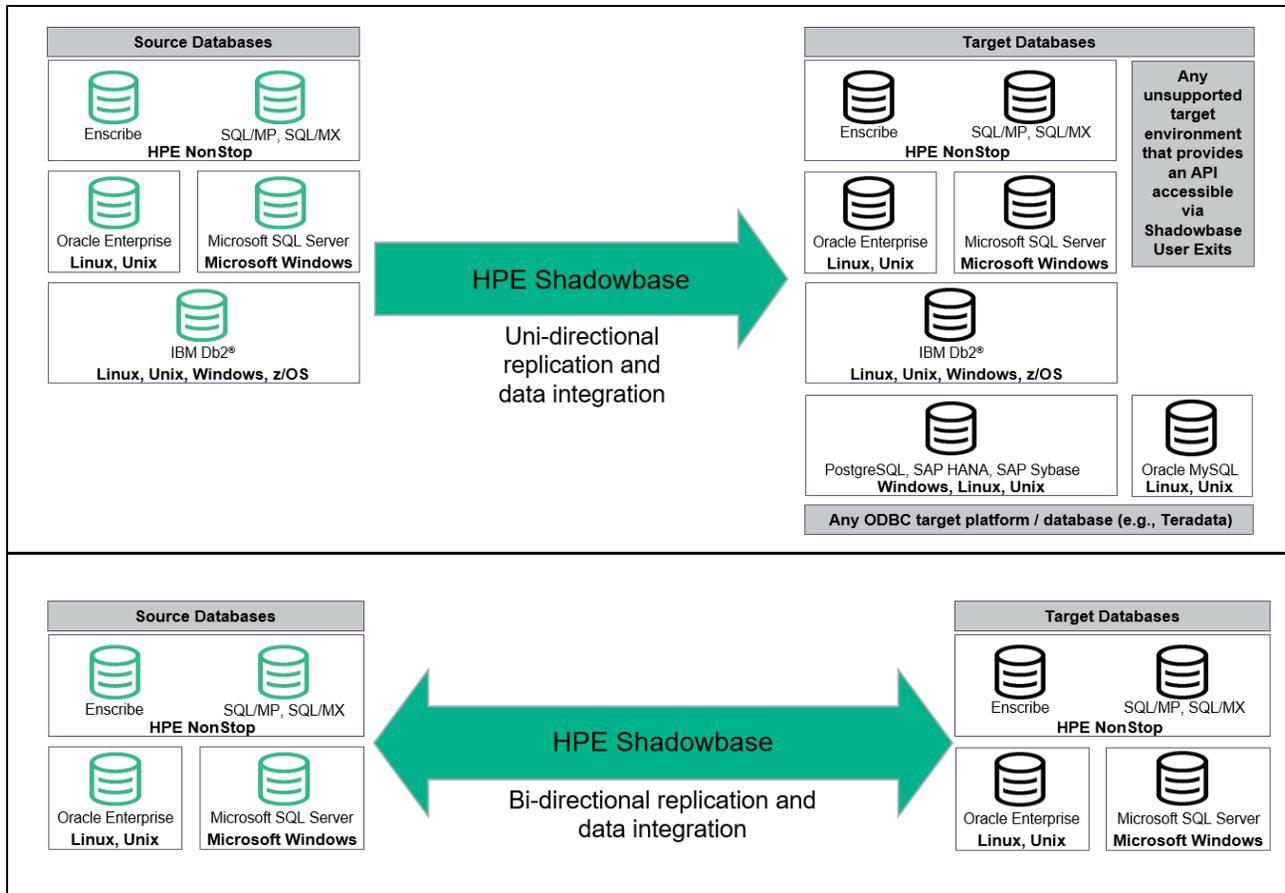


Figure 1 – HPE Shadowbase Platforms and Environments

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

For replication source environments, Shadowbase software supports⁷:

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

For replication target environments, Shadowbase software supports:

- NonStop Enscribe, SQL/MP, and SQL/MX
- Oracle, SQL Server, Sybase, Db2, 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.

⁷All source environments are bi-directional.

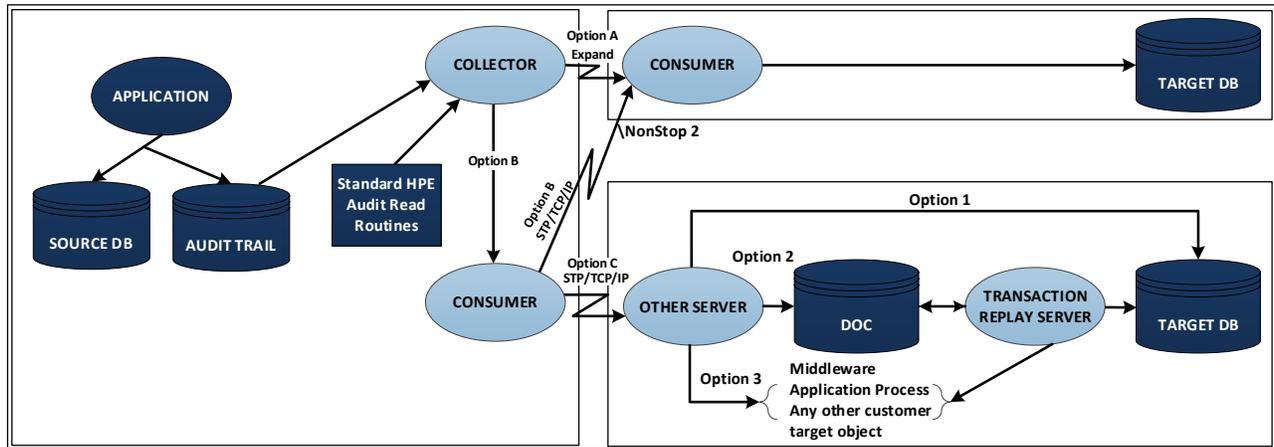


Figure 2 – HPE Shadowbase Overview

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 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.

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.

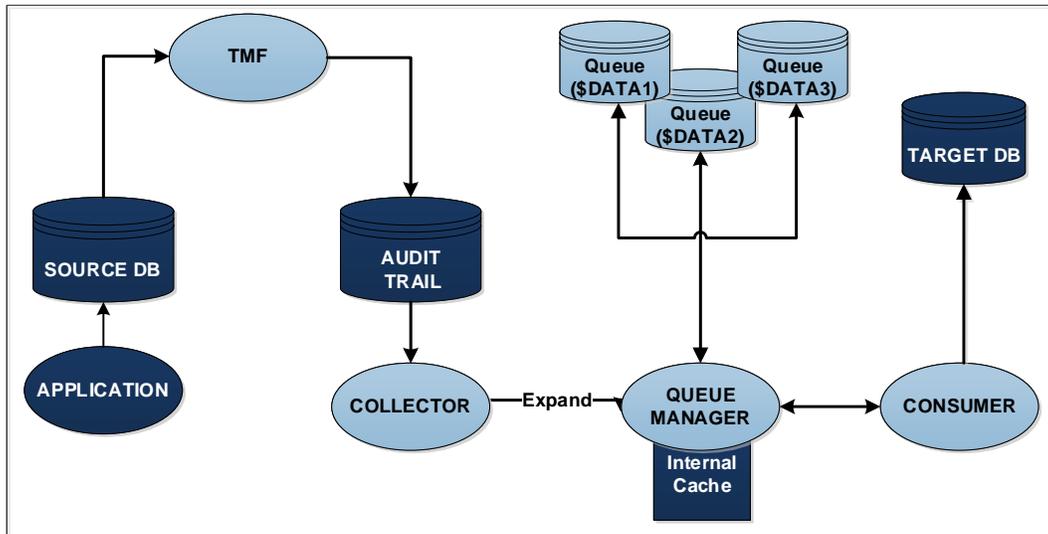


Figure 3 – Queue Manager (QMGR) Flow Chart

There are cases, however, where it may be desirable to queue the data being replicated either on the sending side or on the receiving side, for example if the target database goes offline for some reason. In these cases, Shadowbase replication provides an optional Queue Manager (QMGR) queuing subsystem (Figure 3). The QMGR works similarly to how the TMF audit trails work – it is 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 and in parallel with applying those changes into the target database. This action effectively divorces the delivery of the replicated data to the target environment from the applying of the replicated data into the target database. In other words, if the target database goes off-line during replication, the Shadowbase software continues to deliver the data into the Queue Files until the target database comes back on line again. When that occurs, Shadowbase replication automatically restarts applying the queued data into the target database using the Queue Files.

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 deblock 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 the Shadowbase for 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.

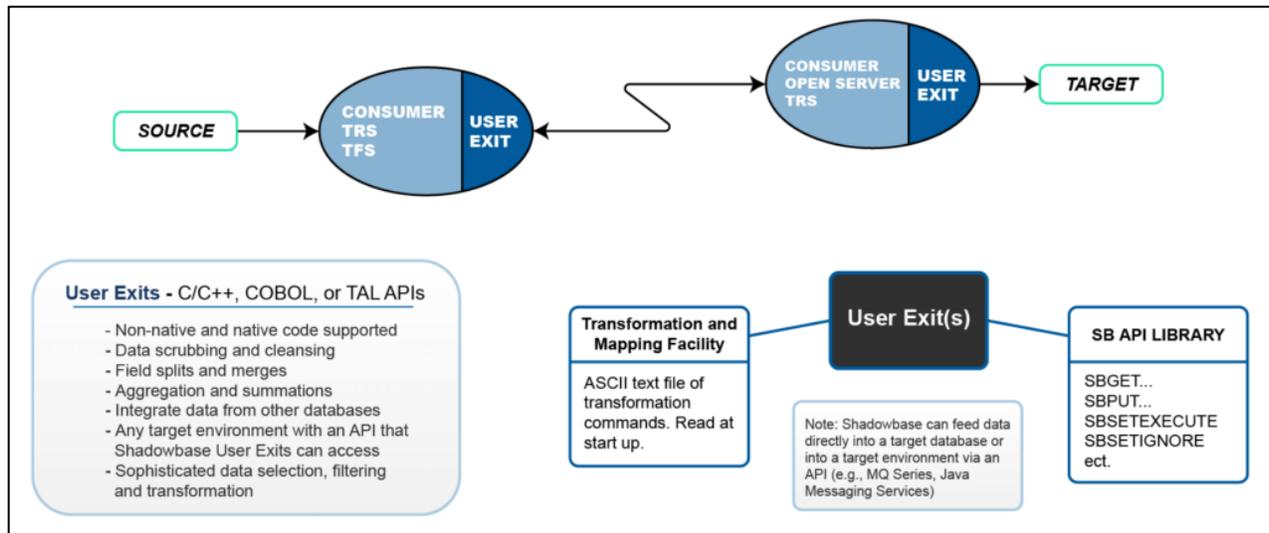


Figure 4 – HPE Shadowbase Customized Processing

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 off-line, 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 off-line, 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 *off-line*, 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.

Manageability

Shadowbase software is easy to configure, monitor, and operate. The [HPE 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 Shadowbase supports. Ease of use features such as simple graphical stoplights that are configured for each Shadowbase process or component make the product easy 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. A sample screen shot from the SEM is shown in Figure 5.

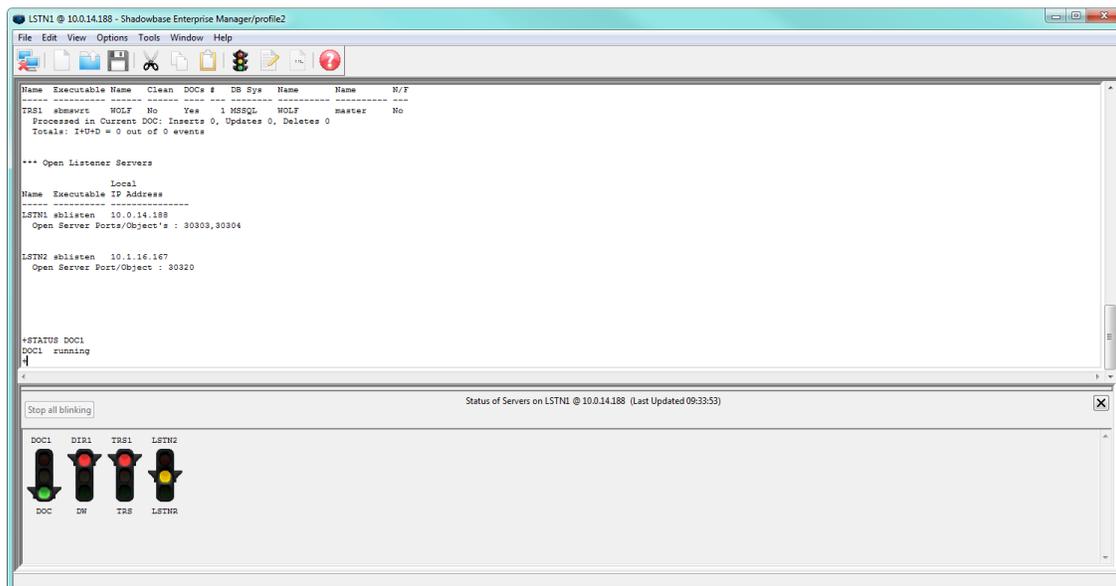


Figure 5 – Sample HPE Shadowbase Enterprise Manager (SEM) GUI Display

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 35 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 6 shows these various Shadowbase topologies.

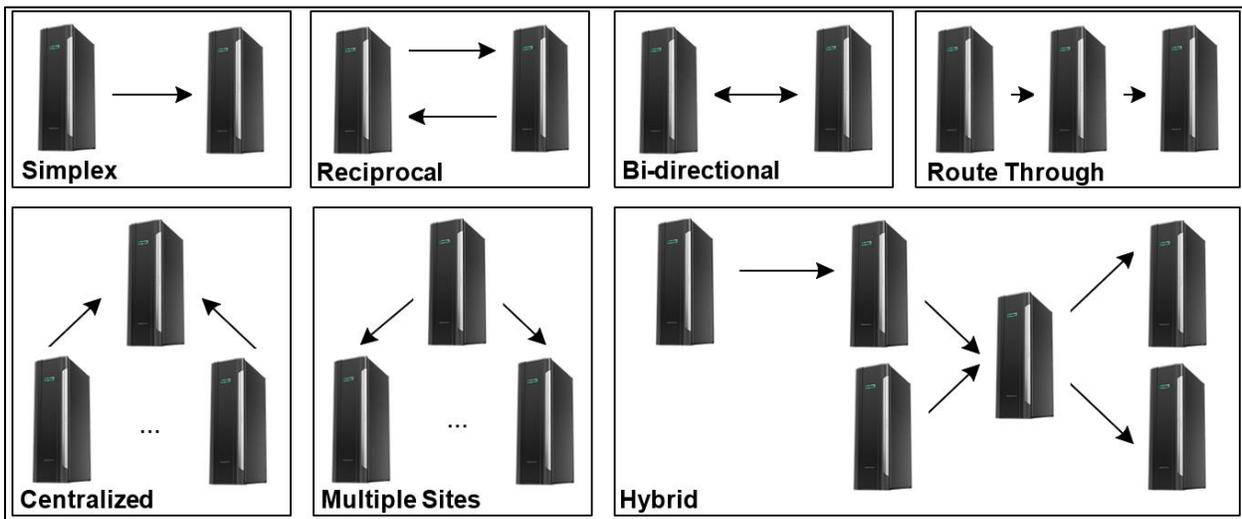


Figure 6 – Various HPE Shadowbase Topologies

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.

Zero Data Loss

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

⁸ For further information on Shadowbase ZDL and ZDL+, visit <https://www.shadowbasesoftware.com/solutions/business-continuity/zero-data-loss/>

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.

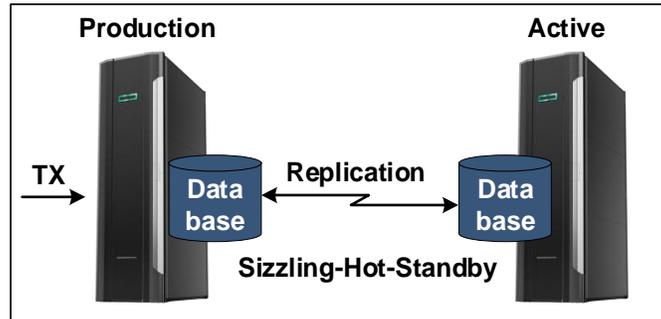


Figure 7 – Disaster Tolerant Banking System

One Bank's March Towards Active/Active

A bank located on the Pacific Rim *Ring of Fire* lost its data center for several hours during an earthquake, online banking services and ATMs were unavailable to customers during this time. When the failover to the backup system did not go as planned, the bank said, "Never again!" and began its march towards active/active.

The bank decided that it had to move from an unreliable *disaster recovery* architecture to a known, working *disaster tolerant* architecture. Disaster recovery means that the IT systems can *recover* from a disastrous event and continue operating, even if that means hours or days of downtime. Disaster tolerance means that recovery is so fast that no one notices the outage or at least is not inconvenienced by it.

The bank evaluated the various replication alternatives that were available for NonStop systems, and chose the Shadowbase replication engine as the one that best satisfied its requirements, to implement an active/passive, disaster tolerant business continuity architecture, as shown in Figure 7. In the process of implementing this solution, the bank also took advantage of HPE Shadowbase Zero Downtime Migration (ZDM) to upgrade its NonStop S-series servers to NonStop NS servers with little if any downtime.

The bank's next step was to extend to Shadowbase bi-directional replication, making failover testing simpler. If backup applications are not already running, then the bank starts them, switches the network, and tests the backup system. The production database is maintained in a current state by bi-directional replication. Therefore, fallback is simply a matter of rerouting the transaction stream back to the production system. This process ensures the backup system is in a *known-working* state at all times.

With the installation of Shadowbase software and some system reconfiguration, the bank moved from multi-hour unreliable failover to multi-second reliable failover. It has achieved its goal of continuous availability with no change in its hardware configuration. With a known working backup system in place, the bank will never again experience a loss of business services because of failover faults. The extensive Shadowbase capabilities position the bank to move towards faster failover using a sizzling-hot-standby architecture, and ultimately putting both of its systems into active production in a fully active/active configuration, with zero downtime.

Large Telco Deploys Shadowbase Continuous Availability Architecture for Scale-Out to Handle Massive Growth in Smart Phone Usage

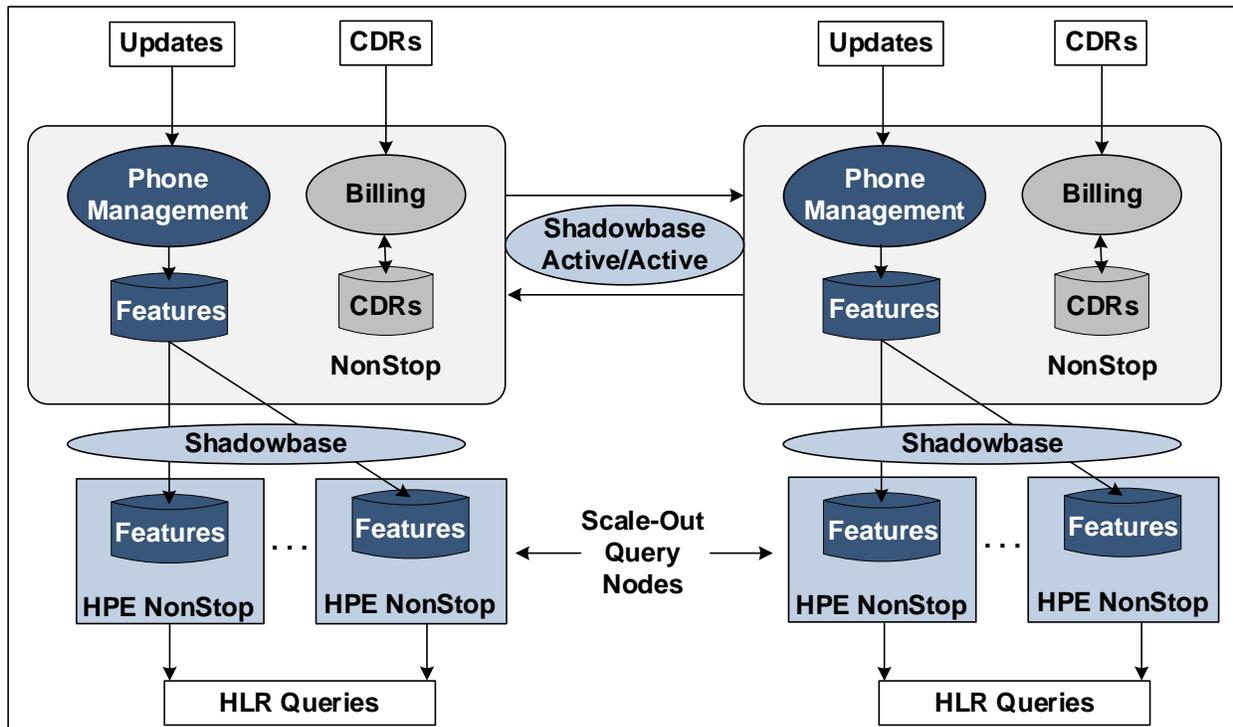


Figure 8 – The Telco Smart Phone Billing and Provisioning System

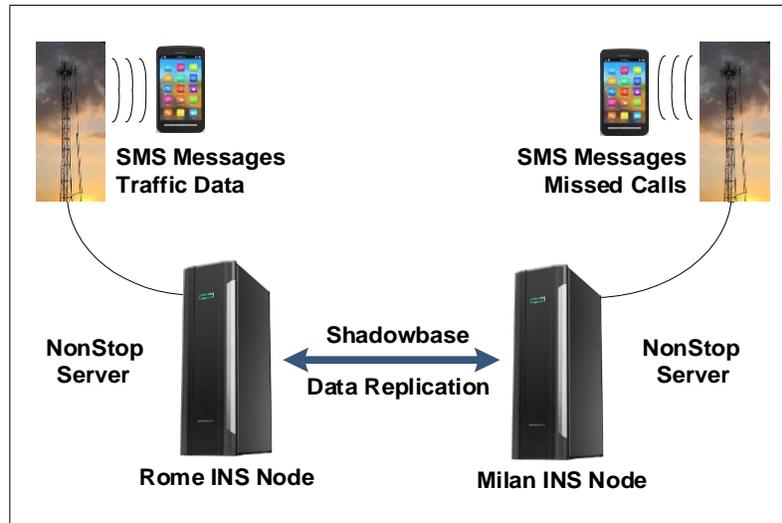
A major international telephone company (“Telco”) based in the United States traces its roots to over one-hundred-years ago, when it was founded to deploy landline telephone services to rural areas in the U.S. It ultimately grew to be one of the major long-distance providers as measured by number of subscribers.

However, as it entered the new millennium, it found that its cellular traffic was quickly outgrowing its landline traffic. It decided to focus on mobile communications and spun off its landline business. It is now one of the largest U.S. cellular network operators, providing a full range of mobile phone and internet services.

In the early 1990s, the Telco implemented a NonStop Server system to manage cell phone billing and fraud detection. HPE Shadowbase software solutions were selected to interconnect the carrier’s NonStop Home Location Registers (HLRs) with the new central billing system to keep the billing system up-to-date with subscriber call activity.

Fast forward to the present: the marketplace is no longer comprised of just cell phones, but rather, *smart* phones. The management of smart phone features is far more complex than it is for the older, simpler cell phones. The HLRs could no longer support all that is required to provision and manage smart phones. Therefore, the Telco implemented a new distributed multi-server NonStop system to provision smart phones and manage their more complex billing requirements. The use of Shadowbase data replication was so successful over the years that the Telco again chose Shadowbase technology to integrate its new smart phone Billing and Provisioning System to build a continuously available, scale-out solution (as shown in Figure 8).

With its choice of HPE Shadowbase solutions to provide active/active geographical continuous availability and system integration, the Telco is positioned to support its rapidly expanding smart phone services well into the future.

Telecom Italia's Active/Active Mobile Services**Figure 9 – TIM's INS Service**

Telecom Italia is the largest cell phone service operator in Italy and provides coverage to over 95% of the country via its TIM-branded mobile services. When Telecom Italia first implemented its cell phone messaging services on HPE's Open Call Intelligent Network Server (INS), running on a NonStop S74000 server, it found that the server was handling one thousand transactions per second; and this volume was growing. Clearly, the company had to prepare for future significant expansion. Furthermore, though the NonStop servers were fault-tolerant, the company had to be prepared to recover from a technical, human, or natural disaster that might take down its processing center.

Therefore, Telecom Italia decided to expand its INS system to a two-node active/active system. One INS node was installed in Milan, and the other in Rome (as shown in Figure 9). These locations provide sufficient separation for disaster tolerance and offered an efficient network topology to support cell tower networking. Each system normally processes half of the transaction load and sends updates via data replication to its companion system. Telecom Italia chose HPE Shadowbase software to perform this bi-directional data replication, as changes are made to a data item in one database, that change is replicated to the other database so that the two are kept in synchronism. Shadowbase data replication is asynchronous so that it does not affect the responsiveness of the application. With this configuration, the TIM network can now withstand the loss of one of its data centers by simply rerouting all transactions to the surviving system. Also, capacity can easily be expanded by adding nodes to the active/active application network.

In an active/active environment with duplicate databases on each node, it is possible that the same row is updated on each node simultaneously, when each update is then replicated to the other node a data collision occurs, and the databases are now inconsistent. To solve this potential problem for Telecom Italia, Shadowbase technology uses *relative* replication rather than *absolute* row replication. If the change is numeric (such as adding five minutes to the subscriber's used time), the call time is incremented by five in the local database. Then, rather than sending the modified record to the other system, only the relative change to the numeric field is sent. In this case, the other system would be directed to add five to that data field for that subscriber, thereby avoiding a data collision and keeping the databases consistent.

Telecom Italia is planning to upgrade its INS systems to the latest versions of the NonStop operating system with no interruption to subscriber services. This upgrade will be done by switching all traffic to one node while upgrading the other node, taking advantage of the HPE Shadowbase Zero Downtime (ZDM) solution.

A Large Canadian Bank Dramatically Improves its ATM/POS Availability with HPE Shadowbase Solutions

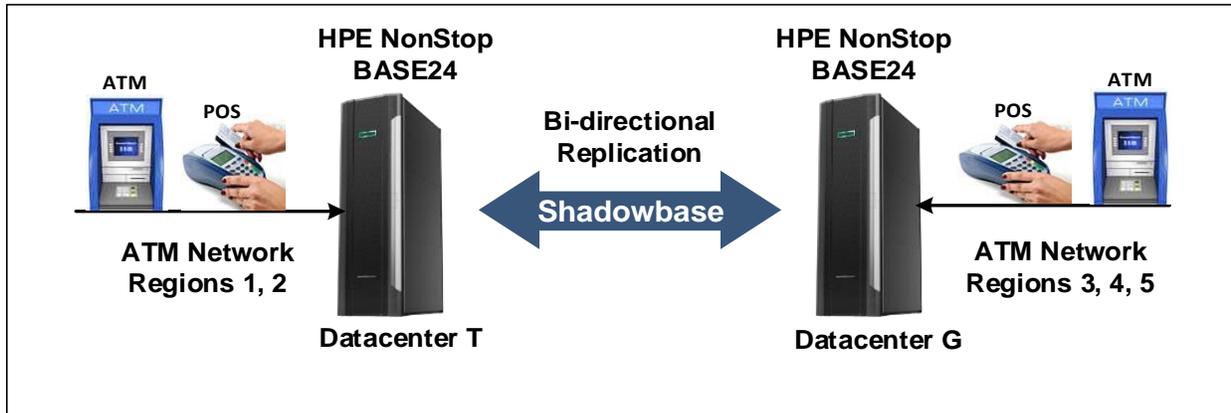


Figure 10 – The Bank's New Active/Active Architecture

The bank operates the largest ATM/POS network in Canada. Should this network go down, much of Canadian retail commerce would come to a halt. The bank uses the BASE24™ product from ACI to manage its ATM/POS network. BASE24 running on NonStop servers is a major application used by banks globally for this purpose.

To ensure continuity of service, the bank operates two geographically-distributed data centers. In the original active/backup configuration, one site was active while the other was a passive standby. Recovery of the application, regardless if for a planned or unplanned outage, was time consuming and complex and typically took about four hours.

The bank decided to dramatically improve on this architecture and recovery time by upgrading its network capabilities to handle an active/active architecture, and then to actively run its application across both sites (Figure 10). Unfortunately, the bank's original active/backup replication product did not have necessary features and was not as flexible as the bank needed. The bank thus performed an extensive evaluation of available solutions, and turned to HPE Shadowbase technology to keep the databases synchronized.

Previously, when an outage of the primary site occurred, all users were affected and would be down for several hours. Now when an outage occurs at one of the sites, fewer users are affected (only those connected to that site), and the recovery takes a significantly shorter amount of time. More importantly, failover is always to a known working system/environment as that site is actively running the application already. The bank no longer has to worry about whether or not the DR target environment will "come up." The new architecture also avoids paying for idled system capacity because there is no "standby" node as all nodes are performing productive work, and all databases are available for application processing work.

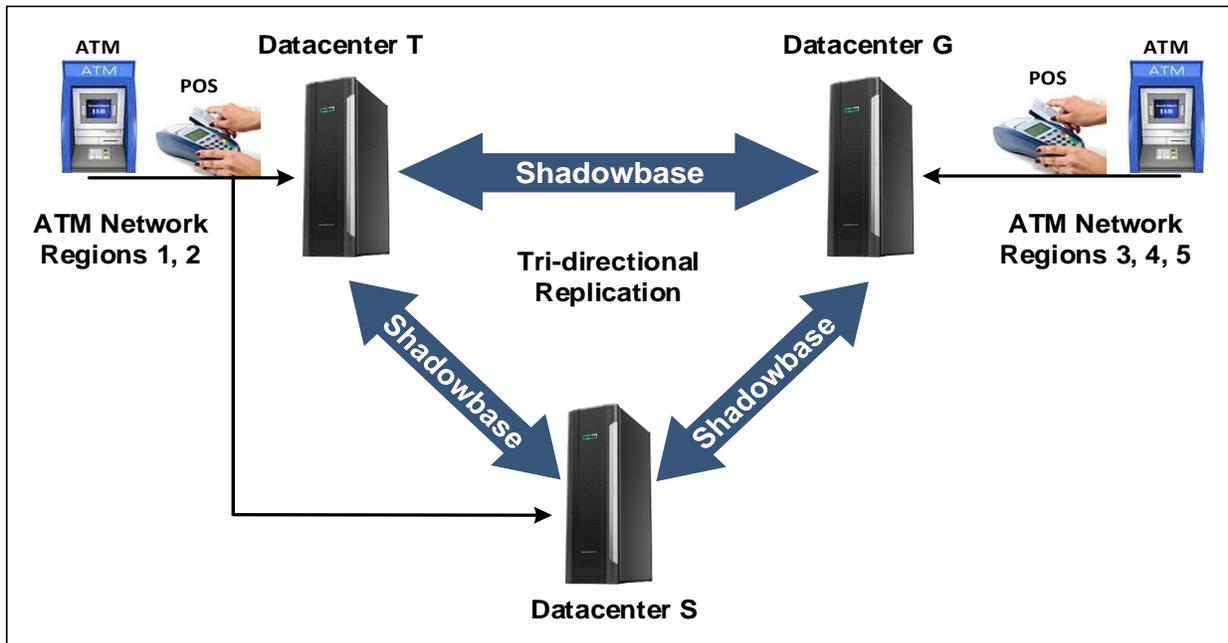
HPE Shadowbase ZDM Achieves Zero Downtime Migration for Large Bank Datacenter

Figure 11 – The Bank's Tri-node Active/Active ATM/POS Application During the Migration Phase

One of the datacenters for this active/active ATM/POS application was located downtown in a major metropolitan city (Figure 11). The bank found that this real estate was becoming very expensive to maintain, and there were other access and security issues. Consequently, the bank decided to relocate this datacenter to a brand new, state-of-the-art (and less-expensive to maintain) facility in a smaller town a few hundred miles away, and close the original datacenter. Given the critical availability requirements of this application, this datacenter migration had to be performed with the least amount of downtime.

In order to avoid outages, it was not feasible and risky to simply establish the new datacenter and instantly switch all users from the old to the new (the so-called *big bang* approach). Rather, it was decided that as users were gradually migrated to the new datacenter, an incremental migration was required with no decrease of availability protection, and with the option to quickly fallback to the original datacenter should problems arise. To achieve this migration, the bank turned to HPE Shadowbase Zero Downtime Migration (ZDM).

The bank first established the new datacenter (including systems, applications, and network). Using Shadowbase ZDM, it then set up tri-directional replication between each of the three datacenters (Figure 11). The new datacenter is shown as Datacenter S and the datacenter planned for shut-down is shown as Datacenter T. This tri-directional replication configuration achieves several important objectives:

- It brings the new database into synchronization with the database to be replaced while the existing database remains online.
- It enables the incremental migration of users from the old datacenter to the new datacenter (all three systems are actively processing transactions), while still providing the required availability levels should a failure occur with any system. It especially provides a fast and reliable fallback position should any problems arise with the new system.
- It enables Datacenters G and S to provide active/active backup for each other ready for the time that Datacenter T is retired.

Once the new database was synchronized, users from Regions 1 and 2 connected to Datacenter T were incrementally migrated to the new datacenter while users in Regions 3, 4, and 5 continued to be serviced by Datacenter G. If any problems arose, it was easy to move the migrated users back to the old datacenter with no loss of service. Since the original Datacenters T and G were still current and backing each other up with active/active bi-directional replication during this interim phase, there was no loss in overall availability protection should one of them fail before the new datacenter went into full production. Finally, all users from Regions 1 and 2 were connected to and serviced by the new datacenter, and the old Datacenter T was shut-down. Datacenters G and S now serviced all users, replicating between them with Shadowbase bi-directional

active/active replication as had been the case with the original two datacenters. The final configuration is the same as shown in Figure 10, except one of the datacenters is now in a different physical location.

By using a Shadowbase ZDM solution, the bank was able to successfully and incrementally migrate one datacenter in an active/active ATM/POS application to a different physical location, with no intermediate decrease in availability protection. Because of the continuously available ZDM architecture employed, the maximum outage period experienced by any user during this whole process was four minutes. No unplanned outages occurred at all. If any serious issues had arisen during this migration, the bank had exactly the same levels of availability protection as it did for the original configuration, and could easily have recovered to an alternate system with little to no user outage.

HPE Shadowbase Helps a Major ISP Migrate from Sybase to HPE NonStop with No Downtime

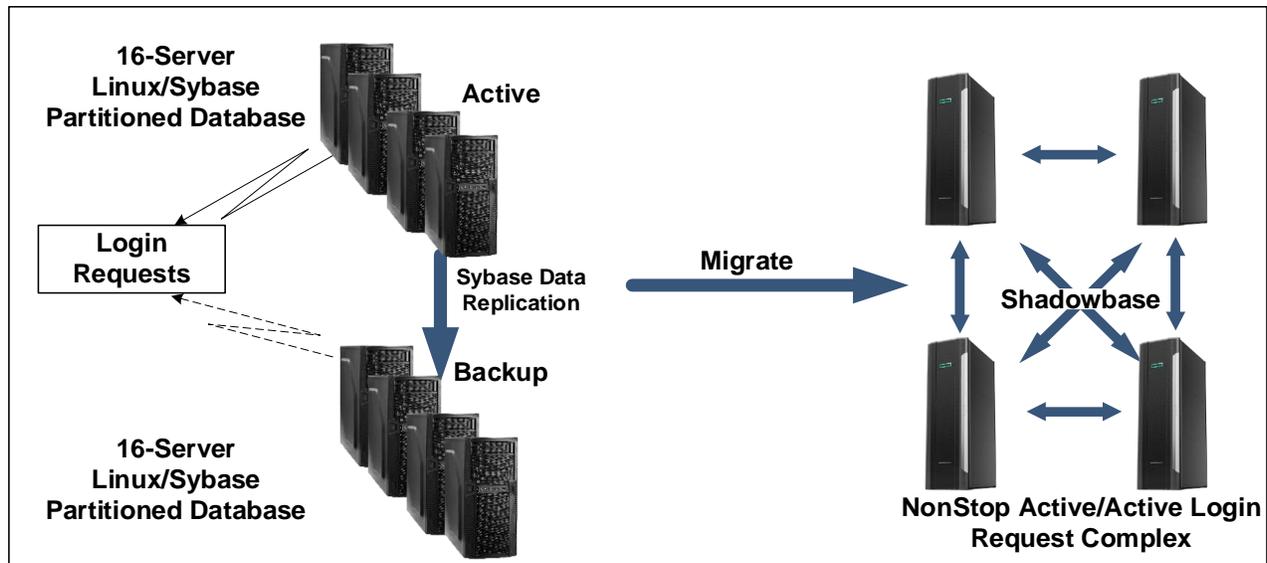


Figure 12 – The Challenge of Migration

A major international Internet Service Provider (ISP) offers email services, internet access, and other functions to millions of global customers. Several million of these customers may be logged in at any one time. Maintaining continuity of service to its customers is a mandatory requirement for the ISP.

Rapid growth led to capacity strains in its IT infrastructure. One such issue occurred in its login subsystem. Though comprising a large farm of redundant servers running many instances of Sybase on Linux, this login subsystem had reached the limits of its capacity. Further additions to capacity were going to be very expensive. The ISP therefore decided to architect and build an entirely new login subsystem using a NonStop SQL/MP database, running on HPE Integrity NonStop hardware in a 4 node active/active configuration.

The problem then became how to migrate from the old Login Request Complex to the new NonStop system without impacting the ISP's customers. The goal was to perform an online migration of the application with zero (or minimal) application downtime, often referred to as zero downtime migration (ZDM). By the judicious use of the Shadowbase data replication engine the ISP was able to gracefully migrate all of its customers over a period of time to the new NonStop system with no interruption in service (Figure 12).

Working with Shadowbase software engineers, a ZDM solution was formulated. Temporary servers were set up to capture all changes to the Sybase databases in the Login Request Complex for a brief period of time. This change capture function was then activated while the Sybase database contents were loaded into the NonStop databases via an extract, transform, and load (ETL) utility. Once the load was completed, the changes that accumulated during the load were sent to the NonStop servers to synchronize them with the Sybase servers. At this time, the users were switched to the NonStop servers for login service.

To minimize any potential negative impact to the users, this process was carried out in phases. There was always a fallback path to the Sybase Login Complex if a migration step failed. The migration proceeded over a period of several months, after which the Sybase servers were decommissioned.

When companies have user communities measured in the tens of millions of active users, any consideration of migrating a core system to another platform is daunting, especially if it is to be achieved with no application service outage. By using Shadowbase data replication technology, this ISP was able to successfully migrate several hundred million user accounts to a new NonStop system with no impact on user service, as well as provide continuous availability using a 4-node active/active login complex.

Two Merged Retailers Integrate HPE NonStop SQL and Oracle RAC

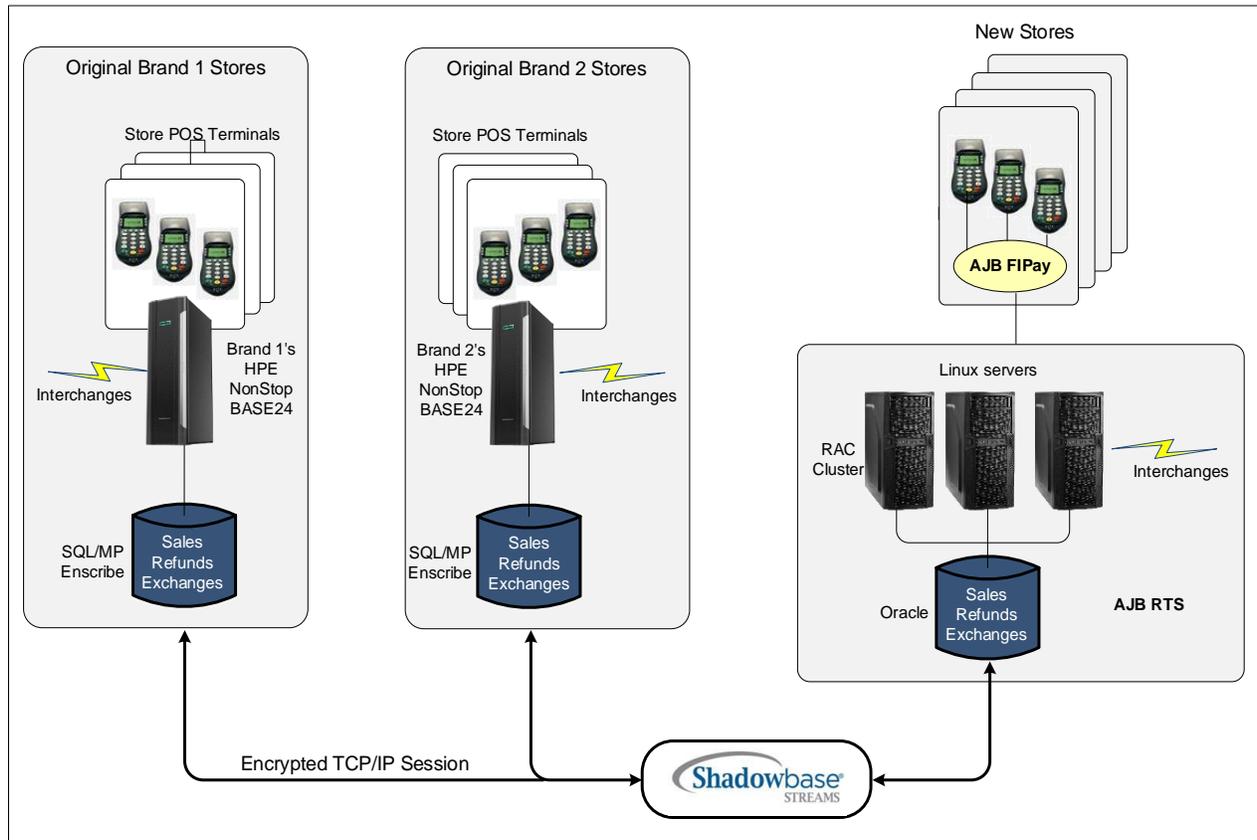


Figure 13 – Integrated Refunds and Exchanges

Company mergers always present a host of challenges, and prominent among them is integrating the disparate IT systems of the merged parties. Two major retailers that recently joined forces faced this exact predicament. Each retailer had thousands of stores, and they wanted to allow a shopper to return or to exchange goods at any of its stores regardless of the store where the original purchase took place. This challenge required transforming NonStop SQL (HPE NSSQL) applications and Linux-based Oracle RAC applications so that the disparate databases could be integrated into a common framework.

The two retailers turned to HPE Shadowbase Streams to integrate the distinctly different applications by mirroring the HPE NSSQL and Oracle databases using bi-directional, active/active data replication (Figure 13). The goal was for each store to have local access to all sales, refund, and exchange activities so that any customer could easily be serviced at any store regardless of the retailer. The integration also needed to occur with little or no application downtime.

The databases of the three payment systems are kept synchronized via bi-directional data replication provided by Shadowbase Streams. In this way, any system can properly process any transaction, regardless of where it originated. In particular, the retailer is able to extend its customer service by allowing a customer to return or exchange a product at any store regardless of the store brand at which it was purchased (with some restrictions).

Integrating applications and databases via data replication is a powerful technique for managing the differences between diverse IT infrastructures and for extending their functionalities. Shadowbase Streams successfully

integrated the distinctly different applications of this retailer by mirroring the HPE NSSQL and Oracle databases using bi-directional, active/active data replication.

Major Bank Uses Active/Active with User Partitioning to Avoid Data Collisions

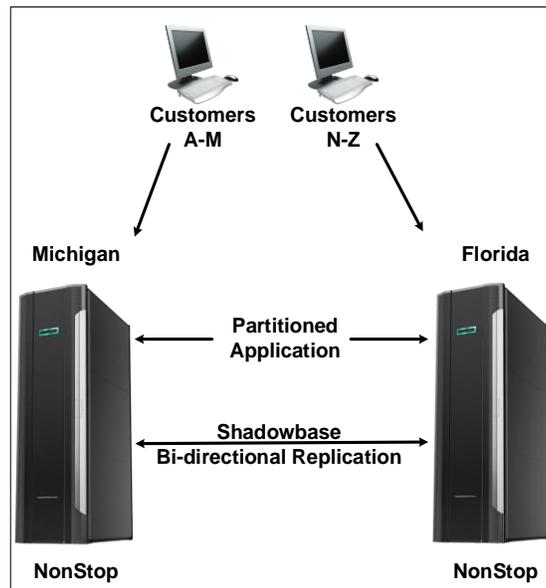


Figure 14 – Active/Active Replication with User Partitioning to Avoid Data Collisions

Headquartered in the Midwest, a major U.S. bank serves much of the eastern United States. As part of the bank's processing solutions business, it provides credit and debit card processing for its merchant customers. If these services fail, users would be denied the use of their credit or debit cards for the duration of the outage, so they must survive any system failure, no matter the cause, with rapid failover time.

The bank decided to choose highly reliable NonStop servers in a two-node active/active configuration to provide these services. One node is located in St. Petersburg, Florida, and the other is located in Grand Rapids, Michigan. This geographical separation ensures that no single environmental disaster, manmade disaster, or system failure will take down both nodes. The Shadowbase asynchronous bi-directional replication engine is used to replicate data between the nodes, keeping them synchronized and consistent with each other, such that active transaction processing occurs concurrently on both nodes (Figure 14).

In active/active systems, data collisions can occur if two users at two different nodes attempt to update the same row or record at the same time. To prevent data collisions, the bank splits its merchants between the two nodes, roughly half of the merchants are assigned to the Florida node and half to the Michigan node. Because the work of each customer merchant is being done on only one node, no data collisions occur as a result of the asynchronous data replication. That is, in no case will one node be making a change to a row that is also being changed by the other node during the replication latency interval. Also, the network is configured so that each merchant can switch to the other node in the event that the merchant's primary node fails, with no interruption to service.

This banking system demonstrates how active/active systems provide extraordinary availabilities – even in the face of frequent natural disasters – and provide an approach to avoiding data collisions by user partitioning with the Shadowbase bi-directional data replication engine.

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** software offers data integration and data replication services for business continuity, disaster recovery, continuous availability, and zero downtime migration. The Shadowbase replication engine forms the data replication backbone used by the other products in the Shadowbase product suite.
- **HPE Shadowbase Streams** feeds database events to other applications or target database environments for data warehousing, real-time business intelligence, data synchronization, data integration, and application integration at the data level or at the event-driven service level.
- **HPE Shadowbase ETL** feeds ETL (extract, transform, and load) facilities by generating the formatted input files these facilities use to load databases, to load data warehouses, or to refresh stale data in downstream environments with current information. It can also accept data generated by other ETL loaders and inject it into the replication engine for subsequent delivery to downstream applications or target database environments.

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.

HPE Shadowbase Options

A number of optional HPE Shadowbase components are available for handling additional data replication and data integration tasks.

HPE Shadowbase Audit Reader (SAR)

This optional Shadowbase utility allows the user to interactively display the contents of the NonStop audit trails using a powerful SQL-like syntax. When reading an event from the audit trails, the Shadowbase utility locates the corresponding DDL definition (for Enscribe files) or schema (for SQL/MP and SQL/MX tables) and format the audit trail event data accordingly. SAR supports selective (i.e., WHERE CLAUSE) filtering, event-type filtering, and file/table data formatting/filtering.

HPE Shadowbase Compare

This optional Shadowbase utility allows the user to compare a target file or table to its source, reporting on all discrepancies found between the two (discrepancies may occur, for example, if the user accidentally purges necessary audit trails, etc.). This utility is helpful for validating that a target matches a source and for satisfying regulatory requirements.

At the present time, the Shadowbase Compare utilities are meant for comparing NonStop to NonStop, like-to-like source/target environments where Expand is available between the nodes (for example, for disaster recovery environments). For additional capabilities, contact Shadowbase Product Management as additional heterogeneous options are under consideration.

HPE Shadowbase UNDO

This optional Shadowbase utility allows the user to selectively “undo,” or roll-back, 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 an “as of” approach. It is also useful to undo the changes made by a buggy program that was accidentally deployed. This utility is particularly useful as it does not require the source file or table (or application) to be taken off-line for the UNDO to execute.

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 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.