HPE Shadowbase Total Replication Solutions for Other Servers

A Gravic, Inc. White Paper

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. HPE Shadowbase Total Replication Solutions for HPE NonStop is Gravic's real-time data replication and data integration solution for the HPE NonStop server, and other operating system and database platforms ("Other Servers"). This white paper describes features and uses for Shadowbase in the Other Server market. For additional information about Shadowbase for HPE NonStop servers, please see the white paper, *HPE Shadowbase Total Replication Solutions for HPE NonStop*.

For over four decades, Gravic has built low-latency, highly reliable data replication products for the demanding 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.

HPE Shadowbase solutions for Other Servers provide data synchronization and integration across a wide variety of platforms and environments including Linux, Microsoft Windows, UNIX, and others, for a variety of databases including Microsoft SQL Server, NonStop SQL, NonStop Enscribe, Oracle, Sybase, SQL Server, IBM Db2®, MySQL, and others. Unlike many traditional data replication and data integration products, HPE 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.

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.

The HPE Shadowbase product suite provides the means to meet user 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.

Table of Contents

Table of Figures

HPE Shadowbase Solutions

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

- **Business Continuity**[1]
  - 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 downtime[2]
  - *Zero data loss* (ZDL) for eliminating data loss in the event of an outage, and eliminating data collisions in active/active architectures
- **Data Integration**[3] **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) environment[4]
- **Application Integration**[5]
  - 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 & Resolution**[6]
  - 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.

[1]For more information, please refer to the Gravic white paper, *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.
[2]For more information, please refer to the Gravic white paper, *Using HPE Shadowbase Software to Eliminate Planned Downtime via Zero Downtime Migration*.
[3]For more information, please refer to the Gravic white paper, *HPE Shadowbase Streams for Data Integration.*
[4]For more information, please refer to the Gravic white paper, *The Evolution of Real-Time Business Intelligence and How to Achieve it Using HPE Shadowbase Software*.
[5]For more information, please refer to the Gravic white paper, *HPE Shadowbase Streams for Application Integration.*
[6]For a complete list of Shadowbase products, refer to the web page: ShadowbaseSoftware.com/products/.

HPE Shadowbase Platforms and Environments

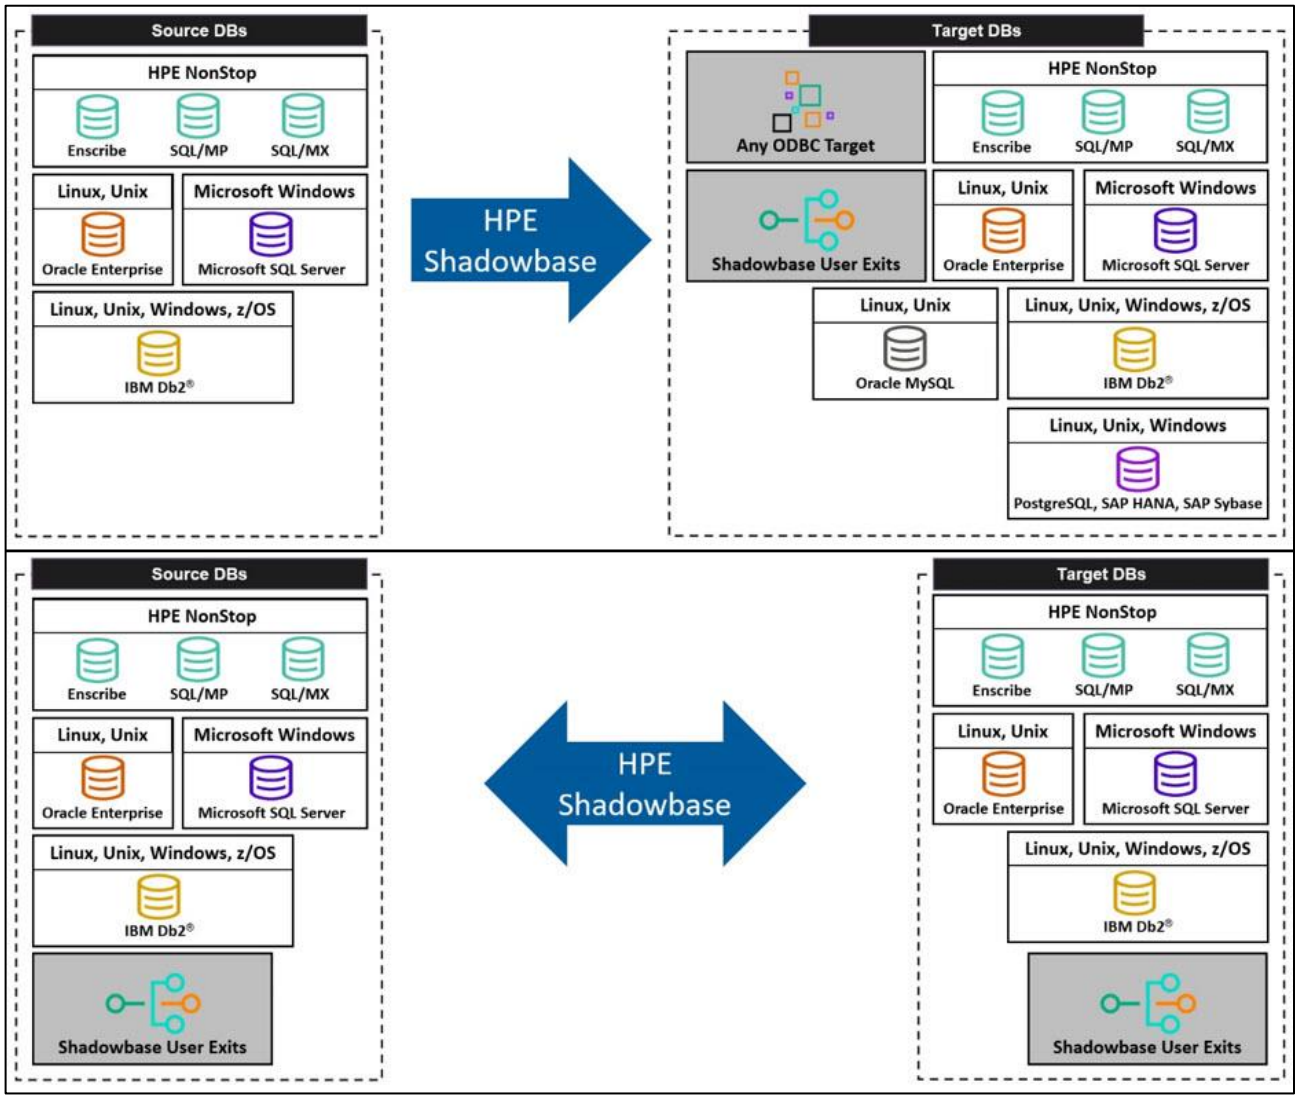

Figure 1 – HPE Shadowbase Platforms and Environments

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

For replication source environments, Shadowbase supports[7]:

- HPE 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 supports:

- HPE 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.

[7] All source environments are bi-directional.

HPE Shadowbase Architecture for Other Servers

Source Data Extraction

Depending on the database, HPE Shadowbase Total Replication Solutions for Other Servers supports different methods for extracting the source database changes to be replicated. In some instances, Shadowbase software uses an available DBMS change log file approach (for example a redo log), either directly or indirectly, for the source of the database changes. In others, Shadowbase software uses a trigger-based approach to capture the changes as they are made to the source database.

Each approach has various advantages. For example, using log-based reading generally means that replication is totally decoupled from the source application as it changes the database. In this case, the replication engine runs independently of the application. However, the changes to the source database are usually not made available to replication until after that transaction has committed, which generally increases the replication latency (the time from when the I/O completed on the source database to when it is replicated and applied into the target).

On the other hand, using triggers to capture the source data changes has very low replication latency, as the database changes are made available to replication as they are being applied into the source database. However, the replication engine data capture is now part of the source transaction, and both committed and aborted transactions are captured (however, only the committed transactions are replicated/applied).

Which approach is used depends on the particular platform/database where HPE Shadowbase replication is installed as a source. At the time of this writing, the following are supported:

- Oracle Source Databases – Trigger-Based or Log/Redo File Capture
- Microsoft SQL Server Source Databases – Trigger-Based Capture
- Sybase Databases – Log-Based Capture, in this instance, Sybase Replication Server to capture/forward the database change events (Sybase Replication Server reads these events from the database log and only forwards committed transactions.)

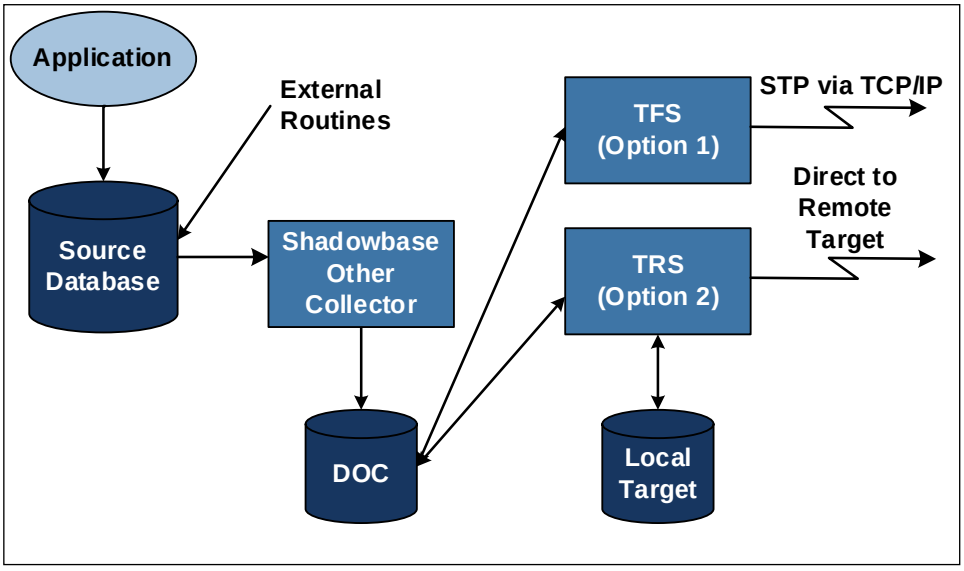

Figure 2 – HPE Shadowbase Other Server Trigger-Based Data Extraction (Oracle, SQL Server)

Trigger-Based Source Database Change Extraction

For trigger-based data capture extraction, as shown in Figure 2, Shadowbase solutions for Other Servers uses highly efficient "C" language routines as part of the Shadowbase Other Collector Process to extract change data from the source databases. The external routines are called by short triggers that minimally extend the path length of the database's transaction. The external routines are designed to quickly and efficiently copy the change data out of the database instance and into stable Shadowbase disk queues called *database of change files*, or DOCs, for replication to the target database(s). These DOC files are ISAM disk queues of SQL statement data that reside outside of the database instance.

From the DOC files, the user selects one of two Shadowbase processes to propagate the change data to the target database(s). The first option is to use an HPE Shadowbase Transaction Forwarding Server (TFS). The TFS reads the data in the DOC file in transaction commit or event order (at the user's discretion) and sends blocks of change data over TCP/IP using the proprietary HPE Shadowbase Transport Protocol (STP) to another Shadowbase process, which could be a DOC WRITER or a DIRECT WRITER (same options as were available on this target). STP optimizes the transmission blocks and ensures that all data is successfully received and processed by the target Shadowbase process.

The second option for the user is to use the HPE Shadowbase Transaction Replay Server (TRS). The TRS process sends data in commit order directly into a target database either on the same system or on a remote system over the database vendors' transport (e.g., Oracle's SQLNet to an Oracle target).

Often, customers use the TFS when replicating in heterogeneous environments, e.g., Oracle to HPE NonStop or Microsoft SQL Server, and the TRS when replicating in homogeneous environments, e.g., Oracle to Oracle. If requested, HPE Shadowbase Technical Support will work with the user to determine which process, TFS or TRS, is the right choice for each replication project.

Log-Based Source Database Change Extraction

For log-based source data capture extraction, as shown in Figure 3, HPE Shadowbase solutions for Other Servers extracts the source database changes (or has them extracted on its behalf) from the source database log files; hence the Shadowbase extract processing is running independently of the source application/database.

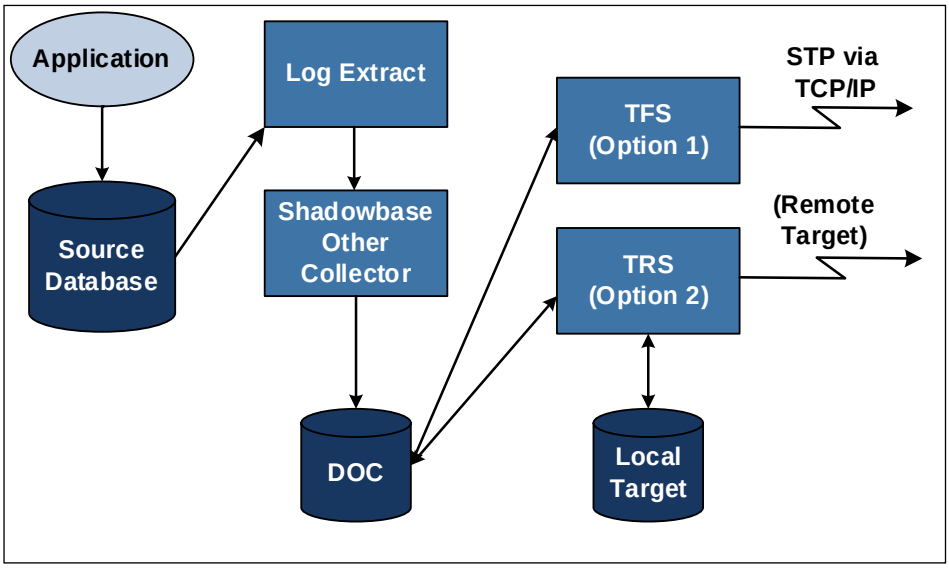

Figure 3 – HPE Shadowbase Other Server Log-Based Data Extraction

In the case of Sybase, HPE Shadowbase replication uses the Sybase "Other Server" interface to Sybase Replication Server to collect the database changes from the Sybase log files, and to forward them to the HPE Shadowbase Other Collector. Sybase Replication Server extracts the data to be replicated from the Sybase log files and provides a guaranteed delivery mechanism to ensure that no source database changes are lost.

Upon receiving the database changes from Sybase, HPE Shadowbase replication uses the Sybase guaranteed delivery protocol to safely store the changes to disk in a DOC file, similar to the architecture used

for trigger source collection. From there, these changes are replicated by HPE Shadowbase technology using the same TRS and TFS options as described above.

In the case of Oracle, HPE Shadowbase replication uses Oracle's Log Miner feature to extract and return the database events for replication.

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

Flexible, Heterogeneous

Unlike other replication solutions, HPE Shadowbase replication is an extensible tool. Over 35 years of experience in designing low latency, real-time replication solutions have gone into the HPE Shadowbase for Other Servers solution, producing a highly optimized, very efficient replication tool with high volume sub-second source database to target database latency. Where HPE Shadowbase solutions excel, however, are in their 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.

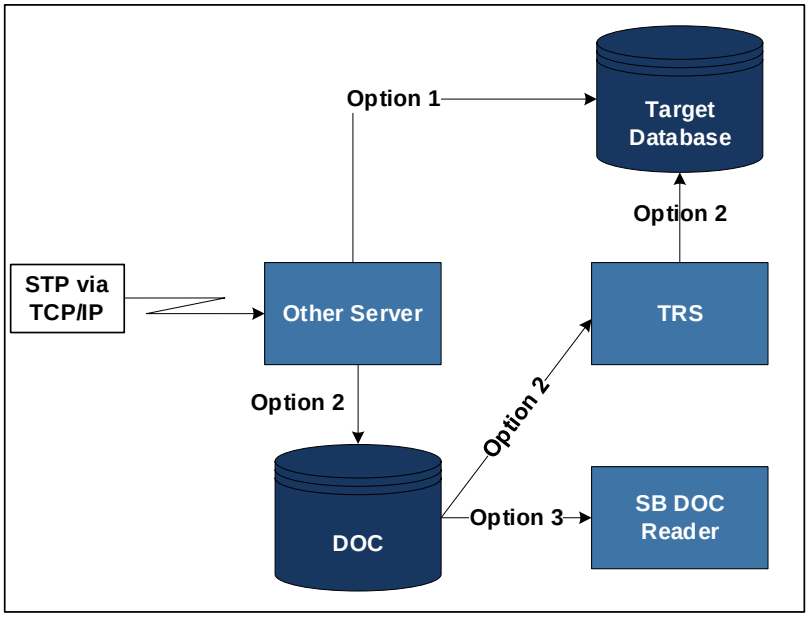

Figure 4 – Target Database Interfaces

HPE Shadowbase software offers the user multiple options for feeding a target database. For some of our users, low latency and high throughput are the more important criteria in their replication process. For these users, HPE Shadowbase replication provides a direct update process (called a DIRECT WRITER) on the target system that receives blocks of change data via STP over TCP/IP, optionally filters and transforms the data, and applies the changes directly to the target database (see Figure 4, Option 1). In the event of an aborted transaction on the source system, Shadowbase technology intelligently aborts/reverses those events in the target database.

Other users have the need to apply only committed transactions to the target database or want to further massage data using pre-written routines prior to feeding the target database. For these users, HPE Shadowbase replication on the target system is configured with a DOC WRITER process that applies the change data arriving via the TCP/IP feed into a DOC file on the target system. From the DOC file, the user has the option of using a TRS process to apply the changes in transaction commit order into the target database (see Figure 4, Option 2), using a transaction forwarding server (TFS) to send it to another environment, or writing their own interface to the DOC file using the provided "C" language shell program, SBDOCRD (see

Figure 4, Option 3). Either way, Shadowbase technology can be tuned to process the DOC files at desired intervals according to the user's needs in order to provide operational flexibility.

Heterogeneous Database Support

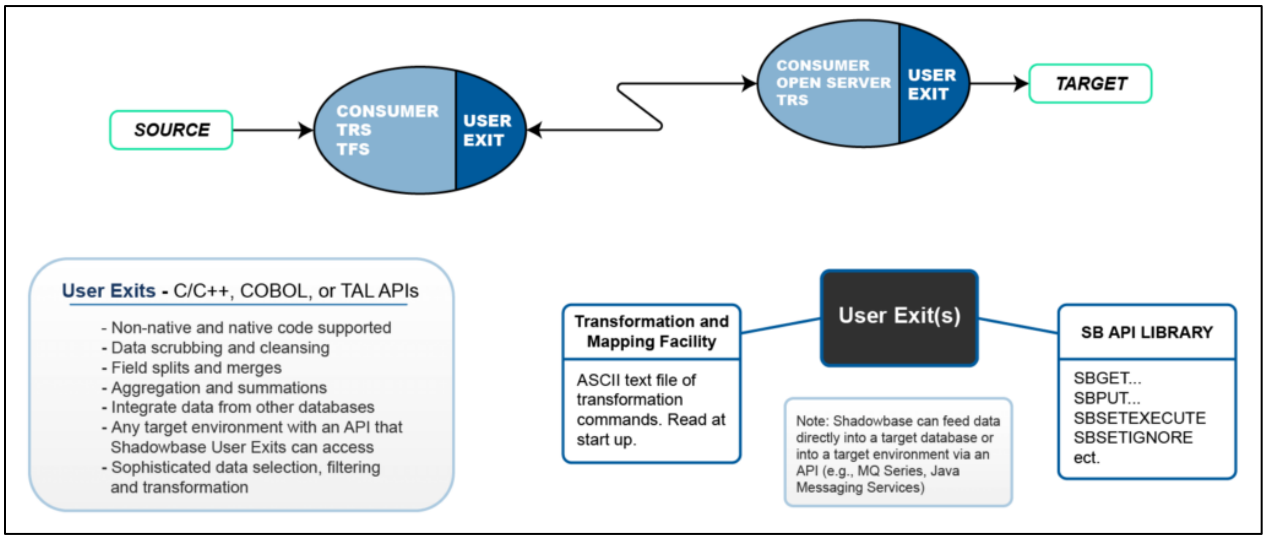

Figure 5 – HPE Shadowbase Customized Processing

Out of the box, HPE Shadowbase for Other Servers supports replication from Oracle, SQL Server, and Sybase databases on most Windows and many UNIX and Linux platforms to the following target platforms and databases:

- Oracle and Sybase on UNIX or Linux,
- Oracle, Sybase, and Microsoft SQL Server on Windows,
- Db2 on AS400 (please inquire), and,
- HPE NonStop SQL and Enscribe on the HPE NonStop Server platform.

Additional source databases or data feeds can easily be supported through custom input APIs built into Shadowbase architecture, and additional targets can be supported through customized user exits or the SBDOCRD shell program (see Figure 4, Option 3). SBDOCRD provides a shell program for reading the Shadowbase DOC files and replicating those events into other environments not directly supported by Shadowbase software. The software can also support any ODBC-compliant target database. HPE Shadowbase Product Group Support is available to assist customers interested in these additional/extensible approaches. Figure 5 illustrates where the input APIs and user exits are driven in the data flow process.

Selectivity

HPE 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

The HPE 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 HPE 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, HPE Shadowbase replication provides APIs in source and target processes for users to add custom written code called *user exits* into the replication process (see Figure 5). These compiled code modules allow the user to otherwise extend the replication engine's default processing to encapsulate additional logic beyond what is otherwise 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.) 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 and compression could always be performed externally using encrypting routers.

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

In HPE 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 – HPE Shadowbase replication handles getting the right data into the right place, in the right order, fast.

Reliability, Availability, Scalability, and Manageability

All of the characteristics of reliability, availability, and scalability of the HPE Shadowbase NonStop products are utilized in the Other Server products, as are the full-featured flexible capabilities such as sophisticated data transformation, filtering, cleansing, mapping, and the HPE Shadowbase Enterprise Manager (SEM) GUI command, control, and monitoring module. Data replication and some of these features are supported from Oracle, SQL Server, and Db2 source databases to Oracle, Sybase, Microsoft SQL Server, Db2, and NonStop target databases.

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 fault tolerant process principles, using a fault-tolerant process monitor to maintain process persistence for the rest of the Shadowbase environment.

Availability

Simply, 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 or fault-tolerant processes 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. 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 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, 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 like 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 6.

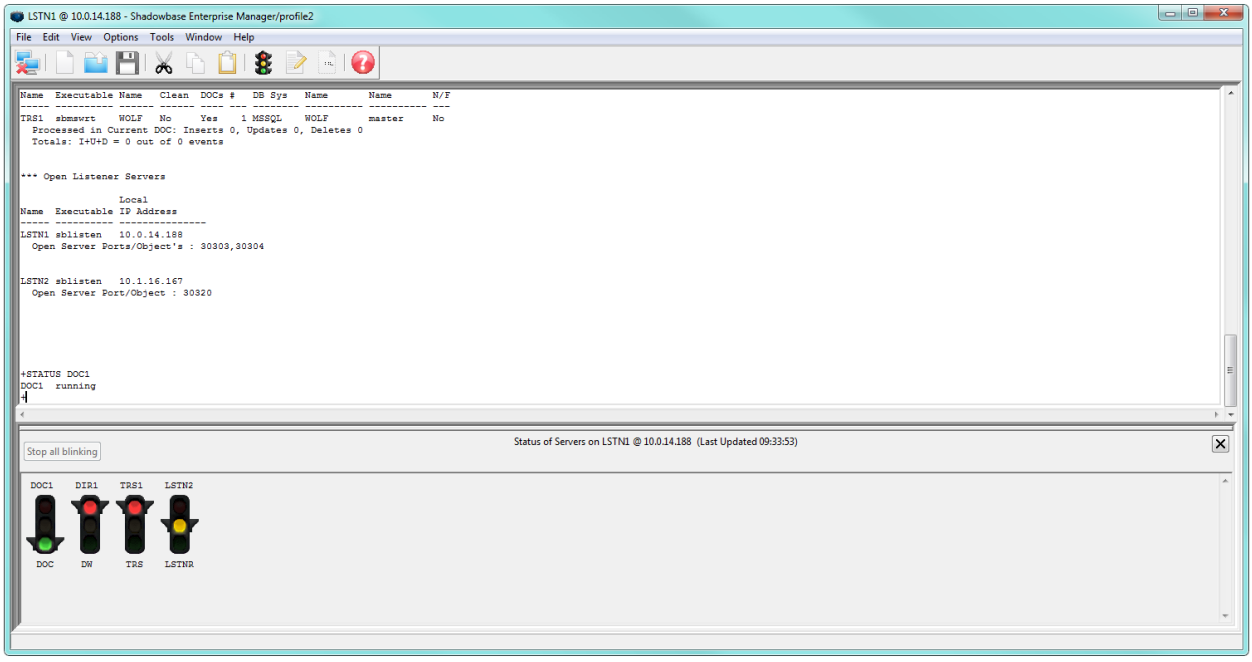

Figure 6 – Sample HPE Shadowbase Enterprise Manager (SEM) GUI

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 7 shows these various Shadowbase topologies.

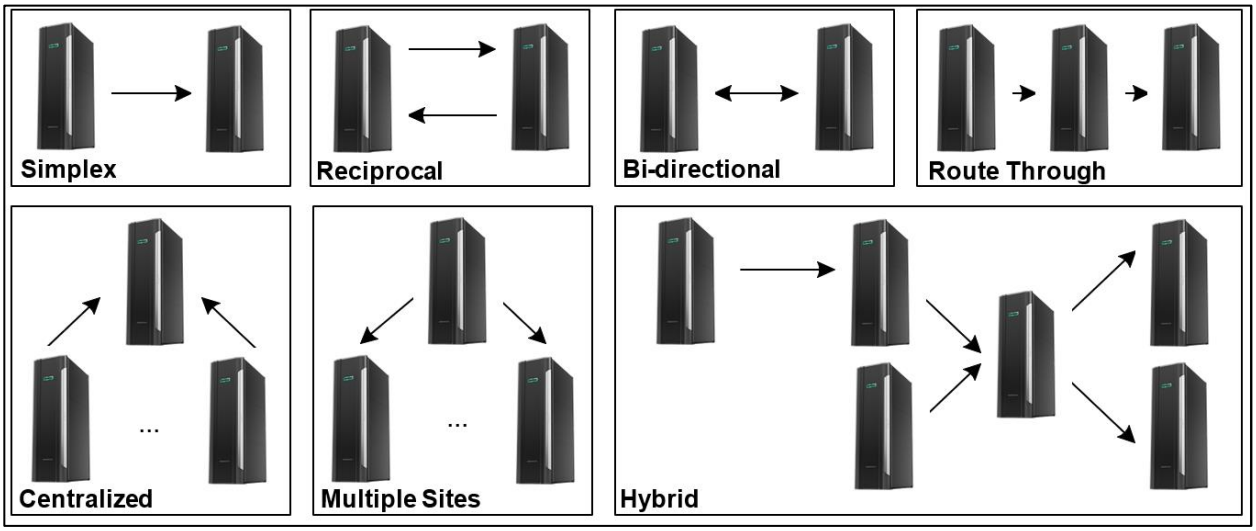

Figure 7 – Various HPE Shadowbase Topologies

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 Other Server customers. Any of the databases listed above that Shadowbase solutions support as a source can be combined with any other in these configurations.

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. Bi-directional replication is used to support

active/active business continuity architectures, which provide continuous service availability. In the event of a system outage, half of the users see no outage at all (in a two-node configuration), and the others are quickly re-routed to an available system.

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-ponging out-of-the-box using patented Gravic technology.

The more difficult component to bi-directional replication occurs when data collisions arise, 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, these solutions are not viable options. For example, they don’t 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*. Contact Gravic for more information on the availability of this powerful solution.

Zero Downtime Migration

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

HPE Shadowbase ZDM software provides the facilities needed for zero downtime migration for both active/active and active/passive architectures. They include the HPE Shadowbase data replication engine, the SOLV online-load facility, and the SOLV validation and verification utility. Taken together, these products offer the means to eliminate planned downtime for system, application, site, or database upgrades and to verify that the upgrade was successfully and accurately performed, thereby eliminating all the associated business costs and risks of an IT service outage.[8]

Zero Data Loss

Shadowbase software supports both asynchronous and synchronous replication. With asynchronous replication, change data is sent to the target system after the changes have been made on the source system.

[8] For further information on Shadowbase ZDM, visit https://www.shadowbasesoftware.com/solutions/zero-downtime-migration/.

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 software with synchronous replication is the solution for the most mission-critical applications, where data loss and/or data collisions cannot be tolerated.[9]

Case Studies

In this section we provide examples of the many ways in which users are taking advantage of HPE Shadowbase for Other Servers to benefit their business and customers.

Stock Exchange and Clearinghouse Gain Operational and Availability Benefits with HPE Shadowbase Solutions

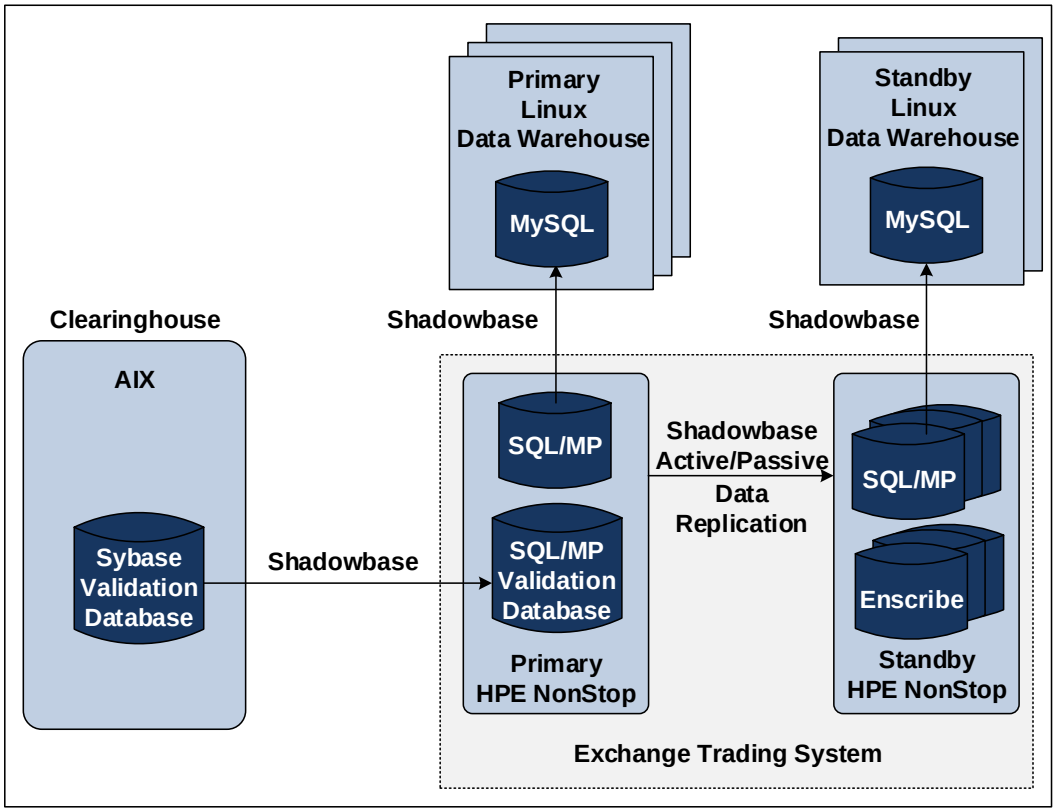

Figure 8 – The Exchange's Many Uses of HPE Shadowbase Data Replication

A regional South American stock exchange embarked on a project to decrease settlement time and improve the accuracy of its trading operations. Commonly referred to as "T+3 Settlement" in the industry, the stock exchange often ran into issues when it tried to settle the trades in the three allotted days. Additionally, with the industry heading towards the "T+0" immediate settlement paradigm, also referred to as Straight-Through Processing (STP) in the industry, the exchange found that it needed to fully automate trading operations to avoid manual re-entry of settlement information. The exchange re-architected its interaction with the clearinghouse to make it fully automated using an HPE Shadowbase data replication solution.

The exchange also manages several Linux/MySQL data warehouses that provide trading history to its traders. The exchange was previously updating these data warehouses using a periodic batch update process, which unfortunately led to the data in the warehouse being stale and out of date. Consequently, the exchange decided to replace this process with real-time Shadowbase data replication to keep the data warehouses current.

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

As part of its IT re-architecture process, the exchange also decided to standardize its data replication and business continuity solutions to a common replication vendor. This standardization required heterogeneous uni-directional data replication solutions for the clearinghouse interface and for feeding the data warehouses, as well as bi-directional homogeneous data replication solutions for the exchange's new business continuity architecture. The exchange chose the HPE Shadowbase data replication engine to satisfy these current and future needs.

Shadowbase software now plays a major role in integrating the many systems in the exchange's IT infrastructure, as well as providing continuous availability for its mission-critical business services. These many Shadowbase technology uses by the exchange are shown in Figure 8, and are summarized here:

- Provide active/passive replication between the exchange trading system's HPE NonStop servers, for disaster recovery. Both NonStop SQL/MP and Enscribe databases are replicated.

- Replicate the validation data from the clearinghouse's Sybase database running on an IBM AIX to the NonStop SQL/MP validation database running on NonStop servers (the original requirement to solve the exchange's manual data entry problem).

- Replicate real-time trading data from the exchange's NonStop SQL/MP databases to the data warehouse MySQL databases running on Linux systems (keeping the data warehouse information synchronized with the current state of the exchange's primary trading database).

It is the flexibility of the Shadowbase homogeneous and heterogeneous data replication product capabilities that allow it to provide this wide variety of functions, in this case demonstrating its suitability for both business continuity and data and application integration purposes.

Large International Tour Operator Utilizes Commodity Servers to Optimize Look-to-Book Processing

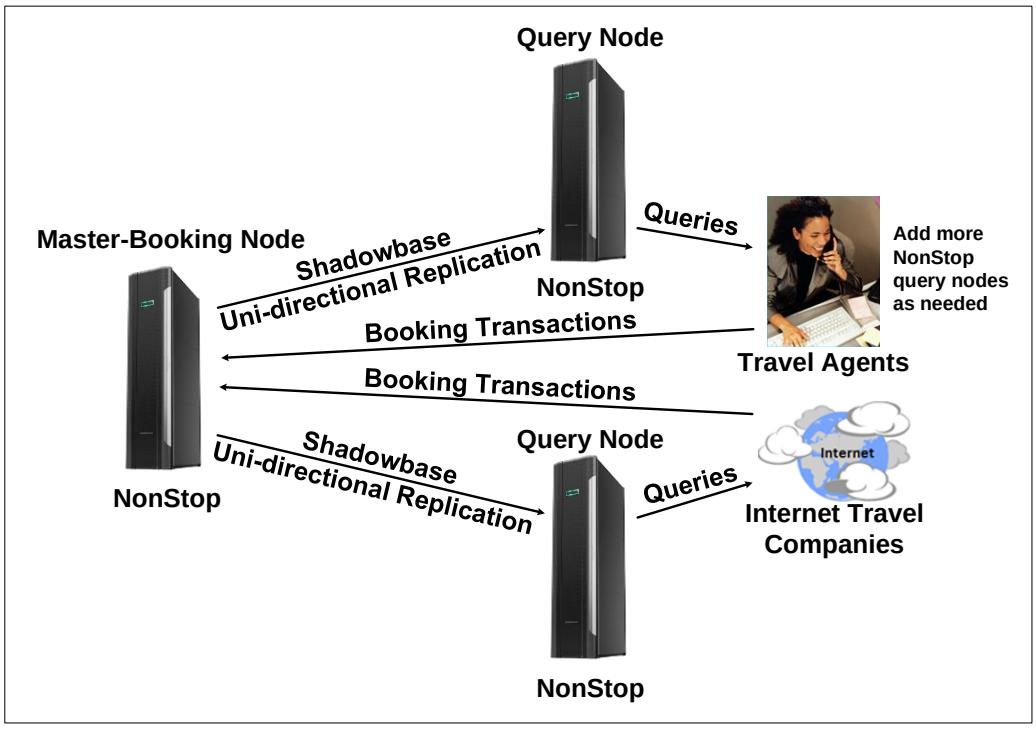

Figure 9 – Tour Operator Multi-node Distributed Heterogeneous Query and Booking System

A large, international packaged tour operator depends upon its HPE NonStop systems to provide around-the-clock services to its travel agent customers as well as to groups and individuals who book reservations on its tours.

A characteristic of travel reservation systems is that a customer makes a lot of queries searching for availability and comparing rates for airline seats, hotels, cars, and other services needed to complete a vacation. Only after a great deal of searching does the customer make a decision and book reservations. The ratio of queries

to the actual booking transaction is known as the *look-to-book* ratio. Before the age of the Internet, when travel agents were doing the booking, the booking activity involved 80% to 90% complex queries and 10% to 20% booking transactions. The Internet changed all that. Now users search many sites to find the best deals before booking. The look-to-book ratio has dramatically climbed to perhaps 500 queries or more for every booking.

The company installed its NonStop system to provide these services before the advent of the Internet, and for years the system performed its functions flawlessly. Then came the Internet with its exploding look-to-book ratios. The load ultimately grew to the point at which the booking system was in peril of running out of capacity. The tour operator faced the problem of expanding the booking system's capacity beyond the capabilities of a single system.

The company decided to move to a multi-node, distributed system by placing the query processing and the actual booking transaction processing onto different systems. Its plan was to have one parent booking system that handled all booking transactions. This system maintained the parent booking database, which was to be replicated to multiple read-only query systems that would handle the heavy query load. The tour operator chose NonStop for its parent system. The query systems are Windows systems using Microsoft SQL Server (MS SQL). Shadowbase provides the heterogeneous replication between the parent booking NonStop SQL/MP database and the Windows MS SQL query databases. To provide the reliability that the tour operator needs, some of the query nodes are also NonStop systems. In this way, if the parent NonStop system fails, one of the NonStop query nodes can be promoted to be the new parent (Figure 9).

This multi-node, heterogeneous system architecture has several other advantages. Because it is distributed and replicated, the availability characteristics are very similar to those of active/active systems. It provides, in effect, continuous availability. Planned downtime is also eliminated. To upgrade a query node, queries directed to it are directed to other query nodes. The query node is then taken down, upgraded, and returned to service. In this way, an update is rolled through all query nodes one node at a time, with no service outage. Another advantage is that query systems are located in each of the tour operator's fifteen worldwide offices. Providing this level of data locality significantly decreases the response time of queries and reduces communication costs. The tour operator also took advantage of the new system's data replication facilities to integrate an Operational Data Store (ODS). Shadowbase software replicates data from the parent booking node to Windows servers running MS SQL that support a large ODS application used to improve the quality and appeal of the operator's tours.

Real-Time Credit and Debit Card Fraud Detection: An HPE Shadowbase Real-Time Business Intelligence Solution

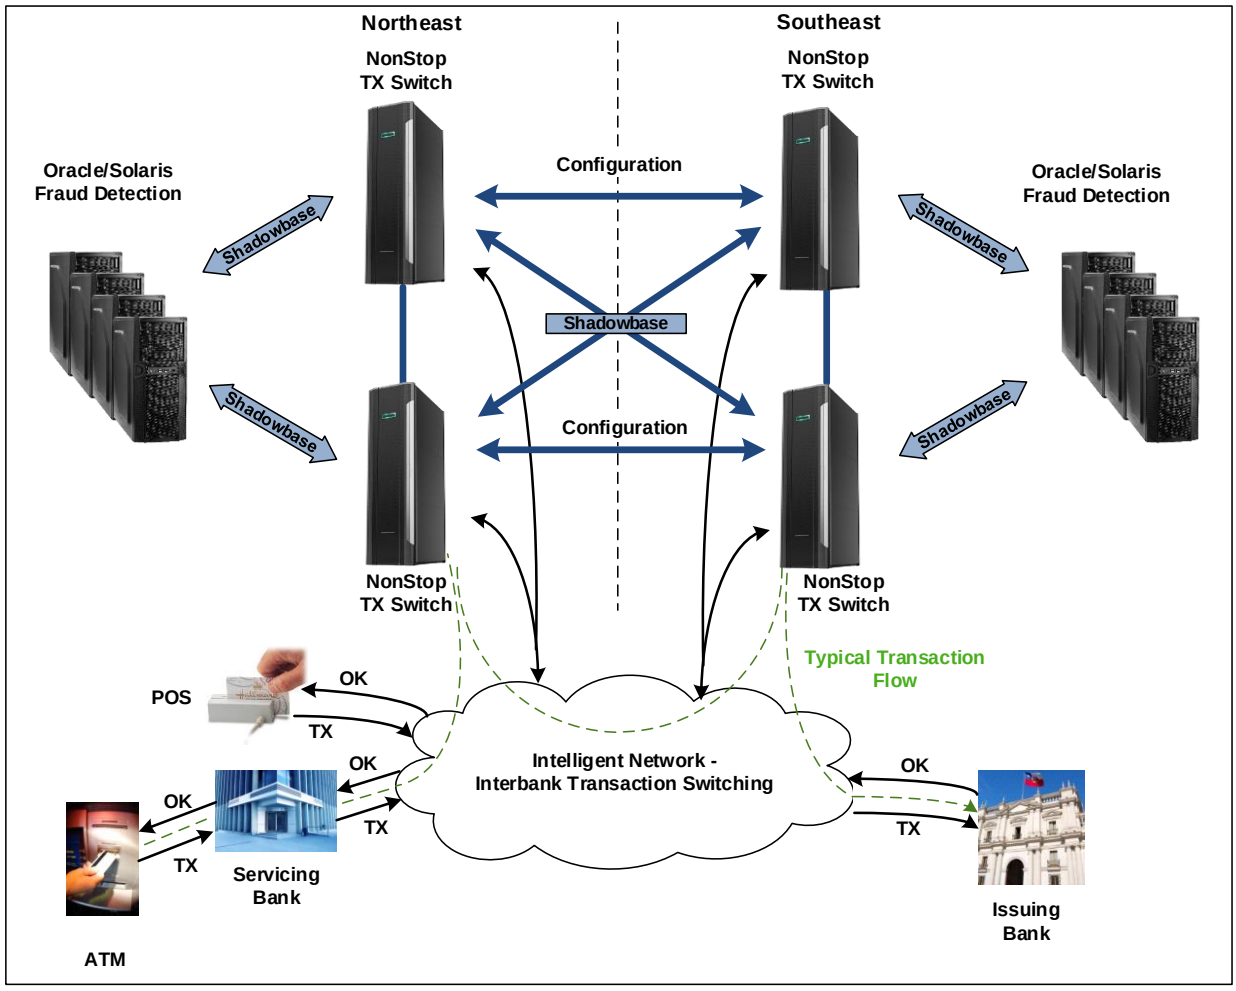

Figure 10 – Real-Time Credit and Debit Card Fraud Detection

A major provider offers interbank transaction-switching services for the authorization of credit and debit card transactions. It uses a redundant network of HPE NonStop computers to implement a credit-card authorization and message switch that gathers the customer transactions from the servicing network, routes them to the appropriate issuing banks for authorization, and then returns the authorization or rejection responses back to the servicing network for delivery to the origination point. Shadowbase active/active replication is used between the NonStop servers to provide for continuous service availability.

A critical problem faced by the issuing banks is that of fraudulent transactions. The transaction-switching service provider realized that there was an opportunity to provide a unique and important service to the issuing banks. If it could detect suspicious or fraudulent activity in real-time, it could stop fraudulent transactions at the retail counter or at the ATM much sooner, or in some cases, even before they were authorized. This service would be a value-added service that would distinguish it from other ATM/POS switching networks.

To implement this system, the switching provider installed multiple high-performance servers that could quickly analyze transactions on-the-fly to determine if they were suspicious. The selected servers were large Sun Solaris servers running Oracle databases. Now, when a transaction is received by a switch node, it is sent not only to the issuing bank for authorization, but it is also replicated in real-time to a fraud detection server via an HPE Shadowbase replication engine. Shadowbase engine routes the transaction to the particular fraud detection server that is monitoring that card or account (Figure 10). The powerful fraud detection system rapidly analyzes the transaction on-the-fly and, if suspicious, notifies the switch node via reverse replication using the Shadowbase replication engine.

This real-time fraud detection system is an excellent example of real-time business intelligence (RTBI), in which events as they happen can control the operational actions of an enterprise. Consequently, real-time business intelligence is often referred to as event-driven business intelligence. A fundamental benefit of RTBI systems is that they can integrate in real-time the independent results of diverse heterogeneous systems and consequently affect the actions of an operational system.

In the transaction switch described above, RTBI is made possible by the high-performance, heterogeneous HPE Shadowbase bi-directional data replication engine. Shadowbase technology can replicate data between a wide variety of databases and platforms, changing the data as it is replicated to meet the needs of the target application or of the target database's schema. The HPE Shadowbase engine is a high-performance, low-latency replication engine that can typically replicate between platforms in tens to hundreds of milliseconds. It is easily scalable to match any needed replication load and is configurable so that capabilities such as, in this case, routing transactions to the proper fraud detection server are simply added.

Cellular Provider Integrates Ancillary Systems to Improve Services

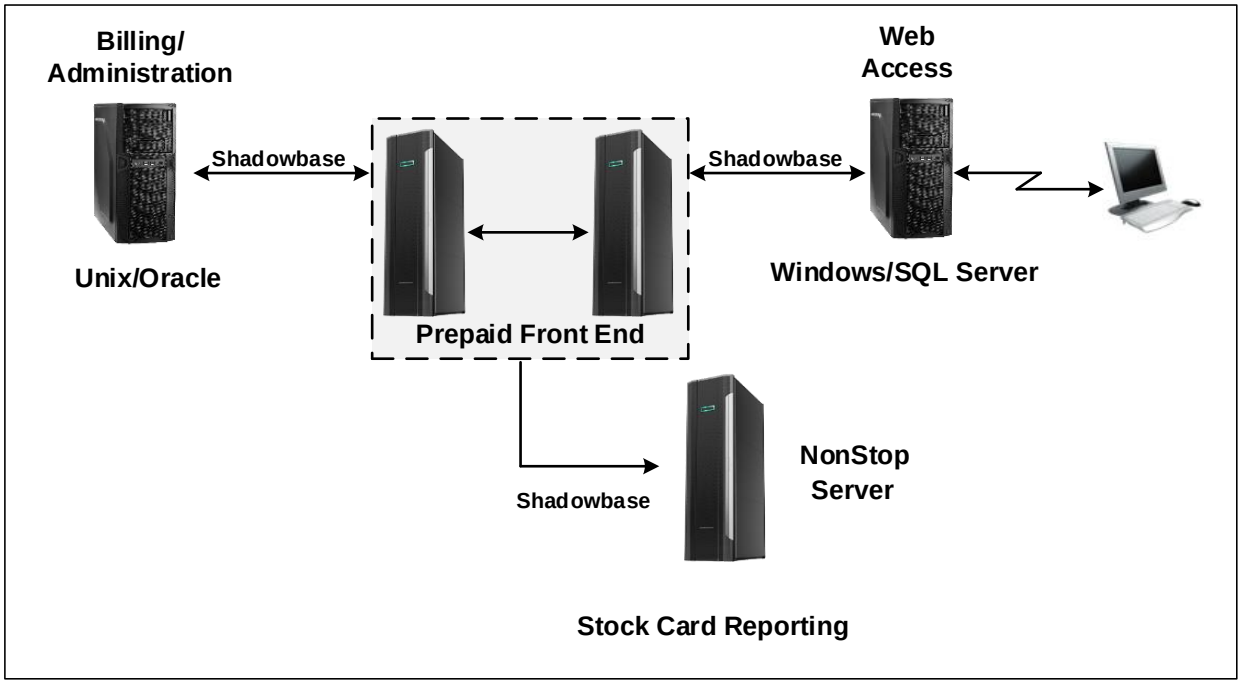

Figure 11 – Ancillary System Integration

One of the largest cellular telephone service providers in South Africa holds more than 50% of the cell-phone market share in that country. It manages a cell network made up of thousands of cell sites, which provide coverage to over 95% of the South African population. The company also provides cell service to several other African countries, including Tanzania, Mozambique, Lesotho, and the Democratic Republic of the Congo. In total, over 20 million customers in these countries are serviced by this provider.

The prepaid calling card market is the fastest growing cellular service segment in Africa. To handle the volume, the company uses in its production environment three HPE NonStop servers with XP storage along with some other ancillary systems. If the prepaid calling card service is not available, much of Africa's cellular service comes to a halt. To ensure the continuous availability of prepaid card service, the heart of this production system is run as an active/active configuration, using HPE Shadowbase data replication between the NonStop servers.

To improve customer service and business operations, the cellular provider also wanted to integrate the prepaid calling card system (PPFE), with other heterogeneous systems. To achieve this, the PPFE feeds several ancillary systems – Shadowbase replication is used to transfer information between these systems by replicating database transactions from one system to another (Figure 11). Subscriber web access to prepaid card data is provided by a web server running on a Windows/SQL Server platform. By using a browser, a subscriber can view the available money left in their account, add credit to a card, and transfer money from one cell number to another. HPE Shadowbase bidirectional replication is used by the web server to access

data from the PPFE to support a subscriber's requests, and to replicate updated SQL Server data back to the PPFE database, thereby keeping the two databases synchronized.

The Billing and Administration system is resident on a Unix system that runs under Oracle. Shadowbase software is used to replicate the Recharge Logs from the PPFE database to the Oracle database to charge the users' accounts for time and other services purchased, and to report on the recharge activity. The cellular provider's Prepaid Recharge System is an excellent example of a heterogeneous system mixing active/active technologies with high availability systems, all connected by the Shadowbase 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 HPE Shadowbase replication engine forms the data replication backbone used by the other products in the HPE 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 extract, transform, and load (ETL) 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.

Outstanding Service and Support

One of the hallmarks of the HPE Shadowbase Product Suite is the support and service provided by the HPE 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 HPE Shadowbase product suite for HPE NonStop and Other Servers. By providing a single point of purchase, HPE and Gravic are improving the overall customer experience. Our customers will benefit from the worldwide reach, industry expertise, and 24×7 support available from HPE, while HPE customers benefit from the wide range of unique features available with Shadowbase software. The product suite is sold by HPE under the name, HPE Shadowbase. For more information, please contact your local HPE account team, 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 ***HPE Shadowbase***. For more information, please contact your local HPE account team or visit our website.