HPE Shadowbase Solutions in a Big Data World

A Gravic, Inc. White Paper

Executive Summary

The amount of information being generated each year is exploding at an unprecedented rate. It is estimated that 90% of the world's data was generated in the last two years, and this rate is increasing. Social media such as Twitter and Facebook, articles and news stories posted online, blogs, emails, YouTube and other videos – they are all contributing to what is now called *big data*.

Big data offers the opportunity for companies to obtain real-time business intelligence that they could never reach in the past from their typical internal systems. Think of the customer-sentiment analysis that can be obtained simply from tweets. However, big data is a collection of data sets so large and so complex, comprising both structured and unstructured data, that it becomes impossible to process with current database management tools and data processing applications. Sometimes the term *open data* is used to describe big data that is publicly available.

Most of the content of big data is noise. It has no value to an organization. However, buried in this noise are tidbits of invaluable data which may be used to determine what customers are thinking, to plan new products, to find the strengths and weaknesses of competitors, to monitor for fraud and cyber-attacks, to defend against terrorism, and for many other purposes. The challenge is extracting the meaningful data from the noise. This is the task of the big data analytics engine.

A big data analytics engine requires a large network of tens, hundreds, or even thousands of heterogeneous, purpose-built servers, each performing its own portion of the task. All of these systems must intercommunicate with each other in real-time. They must be integrated with a high-speed, flexible, and reliable data distribution and sharing backbone. The Shadowbase product suite provides the ideal solution to this problem.

The HPE Shadowbase data replication engine is the glue that binds the servers comprising a big data analytics system. With its ability to rapidly and reliably reformat and transfer large amounts of data between heterogeneous databases and applications in real-time, Shadowbase replication is ready to play a major role in the big data explosion. In this paper, we look at several technologies that interact together to extract valuable business information from the noise of big data, and the role that HPE Shadowbase software solutions have to integrate these technologies.

> ***"Big Data: The next frontier for innovation, competition, and productivity."***
> — McKinsey Global Institute

Table of Contents

Table of Figures

HPE Shadowbase Solutions in a Big Data World

Big Data

Events are no longer sufficient. What do we mean by this statement? After all, business processes and business intelligence are based on events. What did a customer purchase? When was a call made? When was an order delivered? Who logged on to our system and when?

For decades, businesses have captured these events and stored them in transactional databases managed by highly reliable systems.[1] Events drive the business. They determine production schedules, product deliveries, banking, fraud detection, corporate financial statements, and a myriad of other business functions. A business would be paralyzed without its mission-critical online transaction processing systems.

However, the world has evolved. The amount of information being generated that can be valuable to a company has rapidly expanded. Tweets, Facebook postings, news articles and newscasts, YouTube videos, the email and customer service calls that a business receives – all of them may contain information invaluable to a company for making informed decisions and enhancing competitiveness.[2] All of this is big data, no matter the source or format, structured or unstructured.

The data stored in transactional databases represents high-density information. Every element is important. However, transactional data is a tiny fraction of the total data generated worldwide. The data contained in big data is low-density. Most of the data is noise and has no value to a company. But some of the data can be extremely important. How is the valuable data identified and extracted from the noise and put to use? We explore the solution to this challenge.

The Three Vs of Big Data

Big data is commonly characterized by the three Vs – volume, velocity, and variety:

kilobyte	10^3
megabyte	10^6
gigabyte	10^9
terabyte	10^{12}
petabyte	10^{15}
exabyte	10^{18}
zettabyte	10^{21}
yottabyte	10^{24}
brontobyte	10^{27}

Figure 1 – Measuring Data

1. **<u>Volume</u>** is the amount of data that is generated that *may* need to be analyzed by an analytics system.

   The amount of data available to businesses today is truly staggering. Every day, over two exabytes (10^{18}) are generated worldwide.

   Twelve terabytes of tweets per day can be mined for customer-sentiment analysis (what do people really think of you?). Nearly 350 million electric meter readings per year can be used to better predict power consumption and set variable rates to soften peak loads.

   Only a small portion of this data is of interest to any given business. Even so, petabyte and exabyte databases are about to become old news. We are now looking at zettabyte and yottabyte databases. Brontobyte databases are just around the corner (Figure 1).

2. **<u>Velocity</u>** is the rate at which data is entering the analytics system.

   Every minute, there are 98,000 tweets, 695,000 Facebook postings, 700,000 Google searches, and 166 million emails. All of these data streams may contain information that can be extracted in real-time and can be correlated with information from other sources to provide immediate and valuable insights for an enterprise.

[1]These databases typically contain *structured data*, fixed fields or records in a file conforming to some data model, as in for example a relational database or spreadsheet.
[2]Such data is typically known as *unstructured data*, information that either does not have a pre-defined data model and/or is not organized in a predefined manner.

A large brokerage firm may want to scrutinize in real-time five million trades per day as it looks for excessively risky trading by rogue brokers. A large mobile phone operator may want to analyze in real-time 500 million call records per day to predict channel utilization.

Real-time information delivery is one of the defining characteristics of big data analytics. Latency in the information pipeline, from data acquisition to data transfer, data storage, data analysis, and disseminating analytic results, must be kept to a minimum.

3. **Variety** is the range of data types and sources that an analytics system will need to understand.

   Data may arrive in different ways. Some data may be streaming, pushing its way into the big data analytics engine. Tweets and emails are examples of streaming data sources. Other data, such as Facebook postings and news stories, may have to be pulled from their sources by the big data analytics engine.

   Until the advent of big data, corporate data was typically structured and stored in relational databases such as SQL databases. These databases are organized according to well-defined metadata. This paradigm changes with big data. Most big data information is unstructured or loosely structured. Information sources vary widely in the presentation of their data. Sources of information that an organization might utilize above and beyond the structured data of today include web-click streams, social media (Facebook, Twitter, blogs), customer behavior (PayPal, Google Wallet), geo-location (mobile phone GPS), audio (broadcast news), video (YouTube and live video feeds), digital pictures, and sensor readings for a variety of purposes. Though some of these data sources may have some useful structure, e.g., tweets and web clicks, others such as blogs and broadcast news have no structure.

4. **Other Vs**

   Some practitioners have added additional Vs:

   A. **Value** is the worth of the data to an organization for such real-time functions as fraud detection, risk management, and targeted marketing. Many times, there is a temporal attribute added, as in operational data (considered more recent), which is useful for tactical/analytic activity, whereas archival data (considered more historical) is useful for strategic/analytic activity. Which is more valuable? It depends on your goal.

   B. **Veracity** is the trust that business leaders place in the results they obtain from big data analytics, which is partly based on the veracity of the original data.

Real-Time Data and Long Data

There are, in fact, two types of big data that have to be managed – *real-time data* and *long data*.

1. Real-Time Data

Real-time data is used for immediate analytics and business decisions. The most immediate data available to a big data analytics engine is data that is streamed (pushed) to it, such as tweets, web clicks, emails, and customer calls. Other real-time data must be pulled from their sources, such as Facebook posts, news stories, and blogs.

Real-time data is characterized primarily by velocity and variety. Real-time data arrives in a variety of formats, and the big data analytics engine must be able to parse and process all of the real-time data that is presented to it with minimal processing delays.

2. Long Data

Long data is a massive data set that extends back in time over an extended period, such as over the life of an organization.

Long data is important because it places real-time data in its proper perspective. If an organization does not look at events from an historical viewpoint, it will analyze current events as the norm and will be blinded to what has happened previously, missing repeated or unusual events and the opportunities (or threats) thereby presented. This perspective is why the addition of long data to an organization's source of information is so important. It provides context for current events. Consider climate change, for example. Real-time data tells us that our ice caps are melting and that sea levels are rising. Is the culprit our increased carbon emissions, or is it a natural cycle that has gone on for eons? Long data can help answer this question.

The Time Value of Data

The value of information is a function of time. Interestingly, the relationship of value to time is opposite for real-time data and long data, as shown in Figure 2.

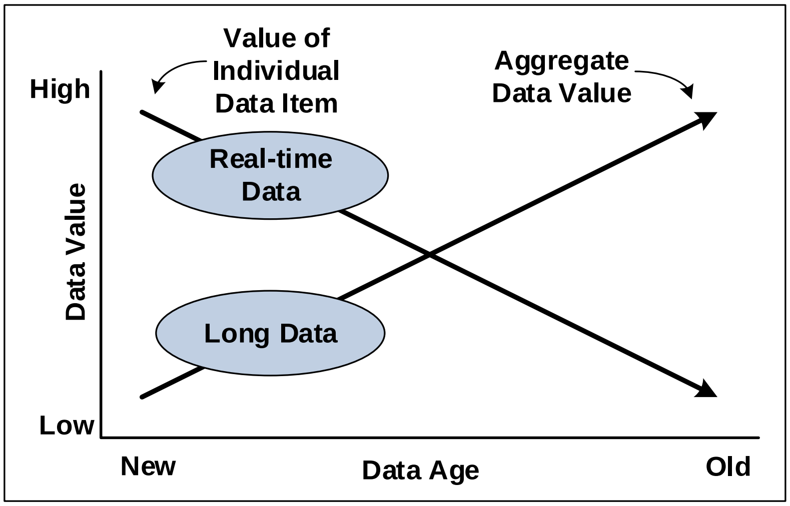

Figure 2 – The Value of Data Changes Over Time and Type

Real-time data is typically used for real-time decision making. The older it gets, the less useful it becomes. Some real-time data items may have half-lives of minutes. Others may have half-lives of microseconds (as is the case for algorithmic, high-frequency stock trading).

Long data, on the other hand, provides historical context for real-time data. The more historical data that is collected, the better is the context. Therefore, the value of long data increases over time as more and more data is accumulated.

The Big Data Analytics Engine

The general structure for a big data analytics engine is shown in Figure 3. As mentioned earlier, three types of data sources make up the information that flows into the analytics engine:

- Streaming data is *pushed* into the analytics engine, including sources such as Twitter and email. *Stream processors* are provided for each source to parse these streams and to deliver pertinent information to the analytics engine.

- Static data is *pulled* from other sources, such as Facebook postings and news stories. *Fetchers* are provided for each static source to fetch new data that has been added to that source and to deliver the fetched data to the analytics engine.

- Transactional data from the organization's transaction processing systems is sent to the analytics engine for its real-time value as well as for its historical value.

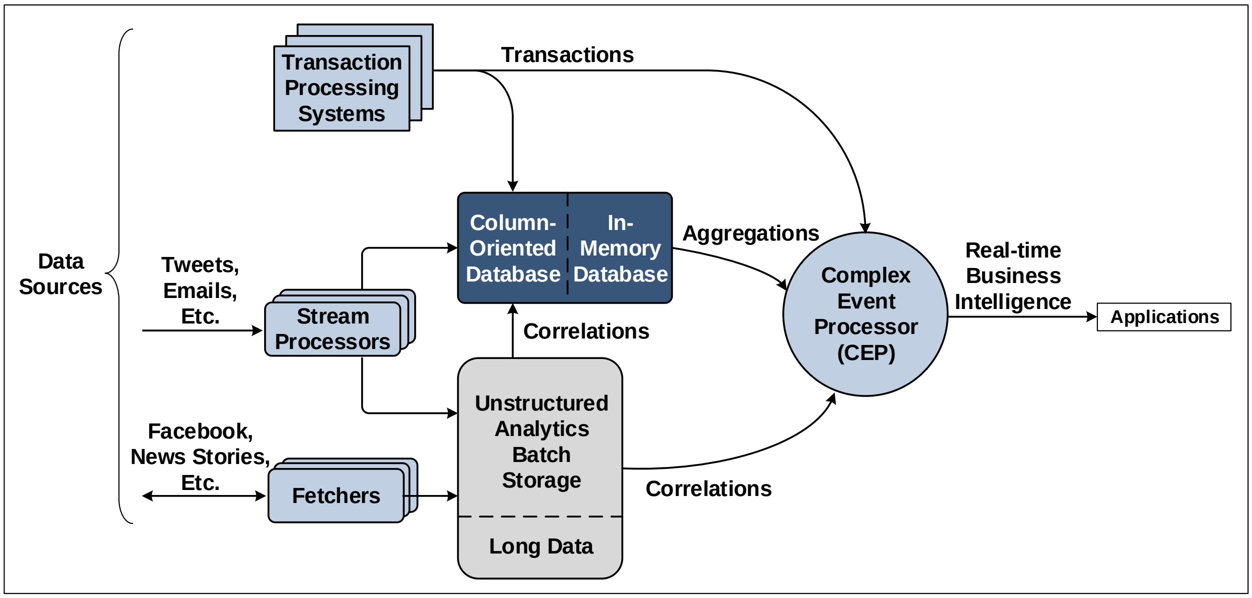

Figure 3 – Big Data Analytics Engine

At the heart of the analytics engine are several (typically massive) components. The implementation of each can require tens or even hundreds of commodity servers:

- A batch-storage analytics engine is capable of storing unstructured data of any kind and can search that data rapidly for correlations. It receives all of the pertinent streaming data and all of the data that is being fetched from static sources. It analyzes this real-time data in context with the long data that it has stored to determine patterns of importance. These patterns, or correlations, are sent to other elements of the analytics engine for analysis processing.

- A column-oriented database stores intermediate results. A column-oriented database stores relational tables as columns rather than as rows. Many types of queries deal only with one or a few columns of a row and can be satisfied much more rapidly and efficiently with this architecture.

- An in-memory database typically holds the contents of the column-oriented database, which improves performance by eliminating disk-seek and transfer times and can further significantly speed up queries. Coupled with the column-oriented database, complex analytic queries can be completed in real-time.

- A Complex Event Processor (CEP) combines data from multiple event streams in real-time to create more encompassing events. These latter events are the real-time business intelligence generated by the big data analytics engine. They provide in-depth insight into what is happening in the business. The goal of the CEP is to identify meaningful events such as business opportunities or threats and to allow immediate responses through the applications that the CEP feeds.

HPE Shadowbase Data Replication is Integration "Glue" for the Big Data Analytics Engine

As described previously, a big data analytics engine comprises many different systems with different missions. Each system is implemented on a "best-fit" platform with a "best-fit" database manager. There may be a myriad of heterogeneous platforms, applications, and databases that make up the analytics engine. A powerful, flexible, fast, and reliable data distribution fabric is required to interconnect these systems. The Shadowbase data replication engine from Gravic, Inc. fulfills this role[3].

Data must be passed from external information sources to the analytics engine systems, between the systems themselves, and from the CEP to a variety of applications, typically when actionable events are identified. For big data real-time applications, also referred to as "operational applications," this data transfer must be very

[3]For more information, please see these other Gravic white papers: *HPE Shadowbase Streams for Data Integration* and *HPE Shadowbase Streams for Application Integration*.

fast, with high capacity and minimum latency. The Shadowbase process-to-process architecture eliminates disk-queuing points that can slow down information delivery. Sub-second replication latency is achieved. The Shadowbase architecture can be multithreaded, including the communication channels, so that any desired data transfer capacity can be attained.

In some cases, the data is generated by an application; in other cases, it is taken from a database or application log file. The data must be delivered to a target application or to a target database. Shadowbase replication supports heterogeneity. It can receive data as it is generated from any supported application or database and can deliver it to any supported application or database, including support for filtering, redefining, and enriching the information in-flight to satisfy any target environment formatting needs.

When replication takes place between diverse applications and databases, complex reformatting and restructuring of the data is required. Shadowbase technology includes provisions for many types of reformatting. If Shadowbase replication does not inherently support a specific type of reformatting, it provides *user exits* so that the user can define customized reformatting algorithms.

Shadowbase technology can replicate data synchronously[4] or asynchronously. Asynchronous replication takes place after the fact and has no impact on the source application. Synchronous replication guarantees that an update is either made to both the source and target systems or that it is made to neither. Both synchronous and asynchronous replication preserve the transactional consistency of the target database, and ensure that the target database is available for application processing during replication. Whether synchronous or asynchronous replication is required depends on the particular requirements of the application.

Shadowbase software is architected to provide continuous availability. If one of its components fails, it is automatically restarted and replication continues uninterrupted. If the target system fails, then the Shadowbase engine queues all events until the target system is restored to service. It will then drain its queue of saved events to bring the target system back into synchronization with the data source, and automatically resume replication of real-time events.

The data distribution fabric between the many components in a big data analytics engine must be low-latency and provide high capacity. It must be fundamentally heterogeneous and be able to deal with any application or database as a source or as a target. It must be able to reformat and restructure data on the fly as it moves data from one source to a totally different target. It must be highly reliable.

All of these are attributes of Shadowbase solutions: low-latency, high capacity, heterogeneous, powerful message processing, flexible end points, high availability. Integrating heterogeneous data resources is a formidable challenge, one that has been solved by Shadowbase data replication, which provides the ideal foundation as the data distribution fabric for a big data analytics engine.

[4] Contact Gravic for the availability of synchronous replication.

Big Data Protection

Businesses today are driven by data, and the quality of the business depends upon the quality of that data. Consequently, data has become one of a company's most valuable assets, and other people want it. Stealing or corruption of this data can result in significant business losses, pose serious security threats, and result in regulatory violations. As hackers become increasingly sophisticated, protection of data from unauthorized access is a number one priority for any IT department.

Protection of data from unauthorized access within a big data environment becomes much more complicated because the data is being consumed from many sources (trusted and otherwise), and moved between systems for analysis. For example, a data source may be restricted to only a certain set of users; if this data is replicated to a big data repository, measures must be taken to ensure that access to the data remains restricted to only this set of users.

Fortunately, using a data replication engine (such as Shadowbase) for the data distribution fabric addresses many of these data protection issues. When the data is in motion (being copied between systems) techniques such as IPSec and/or proxy servers (SSL/TLS) can be used to authenticate and encrypt each data packet. For data at rest, as the replication engine applies the data to the big data repository, Shadowbase user exits can be customized to encrypt or obfuscate the data as it is written. Encrypted target file systems can also be used when available. Via these means, the data replication engine ensures that data replicated to a big data repository remains protected, regardless of the source of the data.

Case Studies

In the following section, we provide examples of the ways in which customers are taking advantage of Shadowbase capabilities for data integration to benefit their business and customers.

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

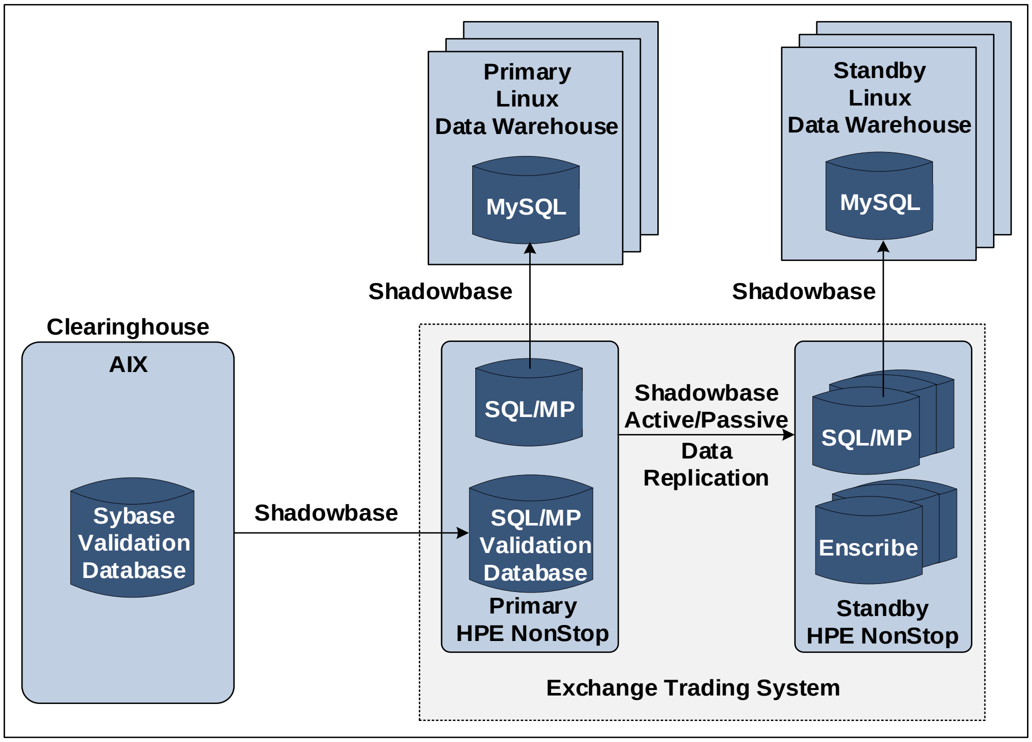

Figure 4 – The Exchange's Trading and Data Warehouse Systems

In addition to its trading system, the exchange provides a series of query nodes that act as data warehouses for trading history, as shown in
Figure 4. These data warehouses are useful for the brokers/dealers to follow a security over a period of time to help them make purchase and sales decisions. The data warehouses use MySQL databases running on Linux servers.

Since the previous replication product did not directly support heterogeneous replication between HPE NonStop SQL/MP databases and MySQL databases, mini-batches of trading activity were instead created periodically and shipped as flat files via the existing replication product from the exchange servers to the data warehouses. These files were then periodically loaded into the target database. This approach meant that the data warehouses were often lacking the latest data for supporting broker queries.

The exchange deployed Shadowbase data replication to take over this function as Shadowbase technology provides real-time, transactional replication of updates from the source exchange HPE NonStop databases into the target Linux/MySQL databases (data warehouses). Trading activity is now replicated to the data warehouses in real-time by Shadowbase replication, and the information in the data warehouses is kept current with the source exchange database. This improves the quality and utility of trade data in the data warehouses, allowing brokers to make better informed decisions and recommendations for their clients.

Large Telco Deploys HPE Shadowbase for Subscriber Data Mining and Real-time Fraud Detection

A major international telephone company based in the United States provides cellular and smart phone services. When a mobile call is completed, the originating cellular switch (the switch that is serving the calling party) generates a Call Detail Record (CDR) that contains information about the call. The CDR is then sent by the cellular switch to the calling party's Home Location Register (HLR). The HLR is the heart of a digital cellular network. There are hundreds of HLRs distributed around the world. Every subscriber is assigned to an HLR, known as the subscriber's home HLR. The primary purpose of the home HLR is to know where its subscribers are at any point-in-time, to manage authorized features, and to record billing information. The HLRs send their CDRs to a central Billing System, which stores them in its billing database. The CDRs in the billing database are used to generate monthly bills to the Telco's subscribers.

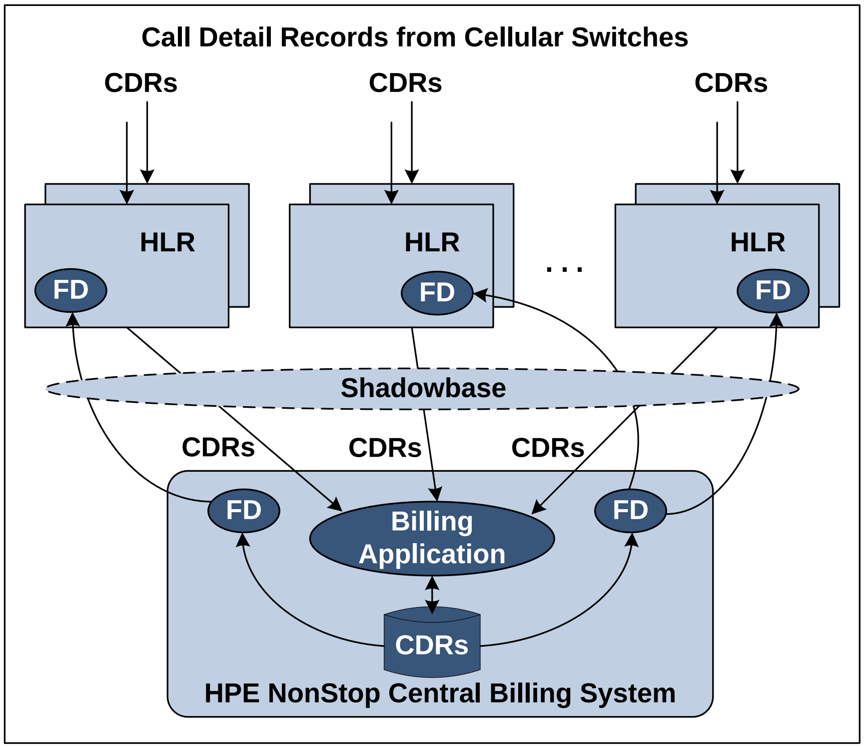

Figure 5 – Real-time Fraud Detection

But the CDR activity in the billing database is valuable to the Telco for more than just billing. It represents a data warehouse that can be used to mine a great deal of data about its subscriber base. What is the geographic distribution of calls? What is the time distribution of calls? Where should additional cell towers be located? What ancillary services are most used?

Another important use of this data warehouse is fraud detection. For instance, the CDR identifies the location of the mobile phone. If that phone makes a call from the U.S. and a short time later makes a call from India, then it is clear that the phone has been cloned. Cloned phones can make apparently legitimate calls, and the cost of the calls will be billed to the unsuspecting subscriber. The Telco has to absorb the cost of these illegitimate calls.

The Telco therefore uses the CDR information to detect a range of suspected fraudulent actions. It will notify the mobile phone's home HLR of the potential fraud. The HLR will block further calls until the situation has been resolved.

All of these capabilities are made possible by the use of Shadowbase replication of the CDRs between the many HLR systems and the central billing system, in real-time (Figure 5). Shadowbase technology is also used in the other direction to send directives related to fraudulent use of mobile phones from the Billing System to the HLRs. By keeping the central billing database current with all the CDRs generated by every HLR, Shadowbase software enables the Telco to collect a vast amount of timely information about its subscribers, and to implement enhanced features such as real-time fraud detection.

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

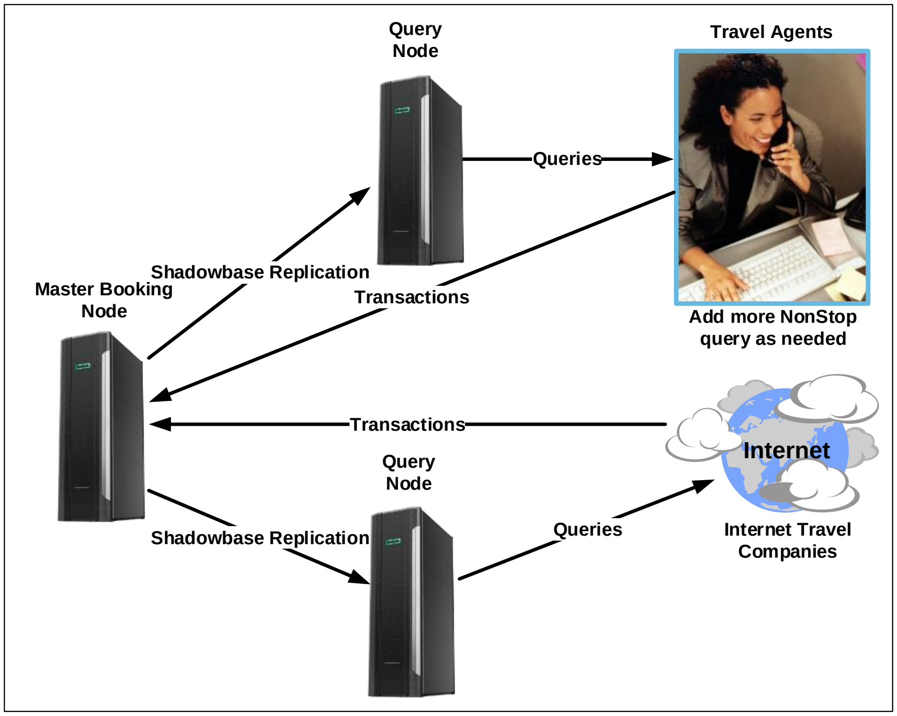

Figure 6 – Tour Operator Multi-node Distributed Query, Booking, and ODS Systems

A large, international packaged tour operator depends upon its IT 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.

To cope with this increased query load, the company decided to move to a multi-node, distributed system and place the query processing and the actual booking transaction processing onto different systems. Its plan was to have one master booking system that handled all booking transactions. This system maintains the master booking database, which is replicated to multiple read-only query systems that handle the heavy query load. The tour operator chose HPE NonStop for its master system. The query systems are Windows systems using Microsoft SQL Server (MS SQL). Shadowbase provides the heterogeneous replication between the master booking NonStop SQL/MP database and the Windows MS SQL query databases (Figure 6).

This multi-node, heterogeneous system architecture has several 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 master booking node in real-time 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

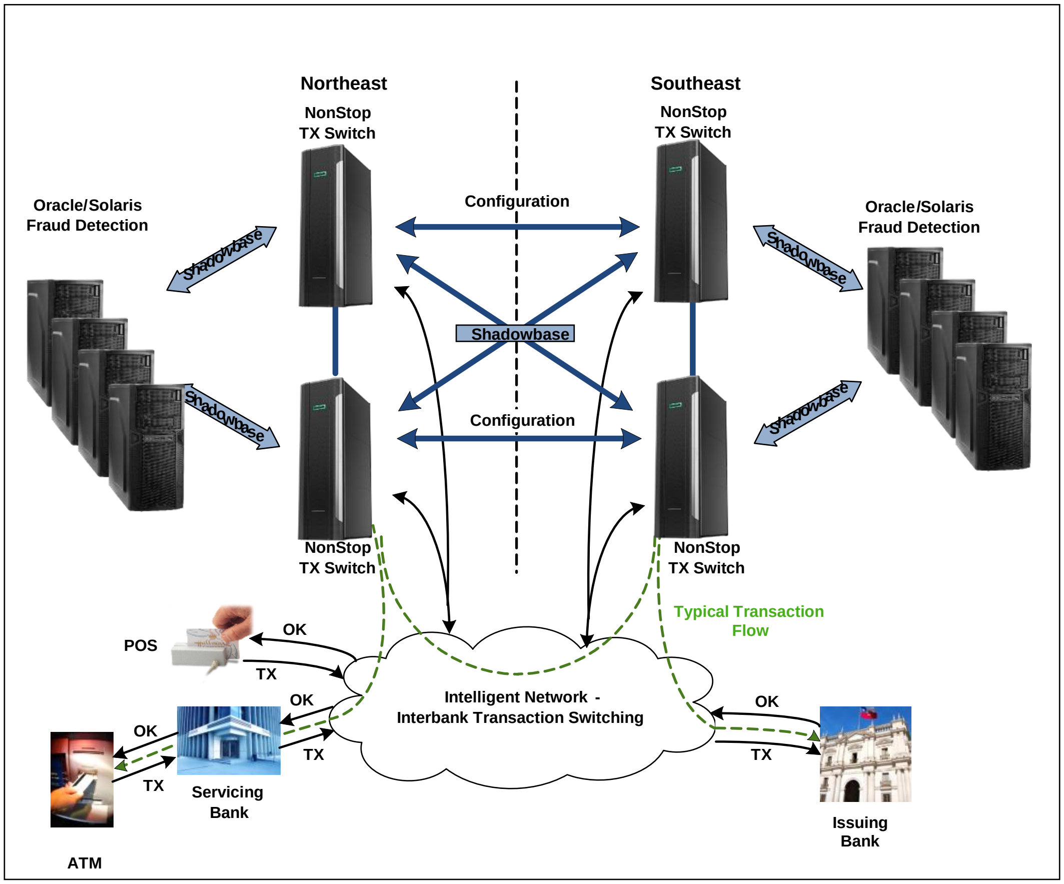

Figure 7 – Fraud Detection Services

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 HPE 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 a Shadowbase replication engine. Shadowbase engine routes the transaction to the particular fraud detection server that is monitoring that card or account (Figure 7). 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 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 Shadowbase engine is a high-performance, low-latency replication engine that can typically replicate between platforms in tens to hundreds of microseconds. 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.

Summary

Big data offers the opportunity for businesses to obtain real-time business intelligence (RTBI) that they could never reach in the past from their typical internal systems. A big data analytics engine can mine social media, the press, email, blogs, videos, and a variety of other data sources to determine what customers are thinking, to plan new products, to find the strengths and weaknesses of competitors, to monitor fraud and cyber-attacks, and for many other purposes.

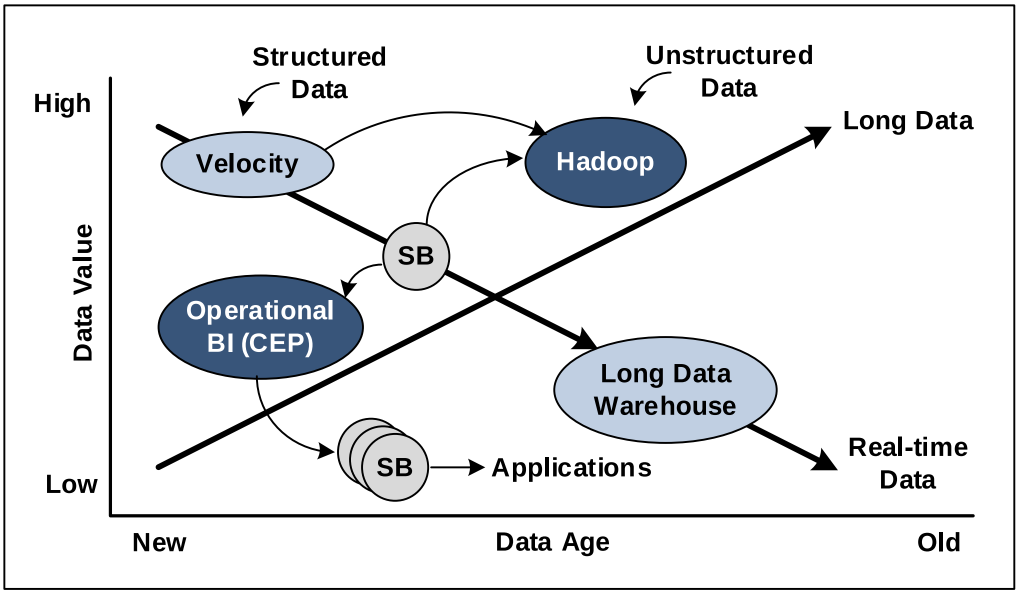

Figure 8 – HPE Shadowbase (SB) Data Distribution Fabric for Big Data Environments

Mining this data is the function of a big data analytics engine. It is a highly complex mixture of many different, built-for-purpose systems that must be integrated with a high-speed, flexible, and reliable data distribution backbone. HPE Shadowbase data replication solutions are ideal for providing this function.

An overview of using Shadowbase replication in a big data analytics engine is shown in Figure 8. This figure indicates the age relationship between real-time data and long data, seen earlier in Figure 2. However, it also shows functional components of the analytics engine. The stream processors and fetchers are represented by the "velocity" component. The CEP is the Operational Business Intelligence (OBI) component. Shadowbase technology can replicate structured data from the stream processor and fetcher databases to the CEP for BI analysis. Additionally, Shadowbase software can replicate the database change results of the operational BI analysis to downstream applications for additional processing.

Figure 8 is only one example of where the many powerful Shadowbase replication capabilities[5] can play a role in delivering inputs and outputs to key processes for analyzing big data. Wherever there is a need to transfer data from a source to another target, regardless of the nature of those devices, Shadowbase software solutions can be placed into service to get the job done efficiently and reliably.

---

[5]For further information regarding HPE Shadowbase data integration and application integration capabilities that can assist in solving big data integration problems, please refer to the companion documents *HPE Shadowbase Streams for Data Integration* and *HPE Shadowbase Streams for Application Integration*, or visit ShadowbaseSoftware.com for more information.

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.

Copyright and Trademark Information