

REAL-TIME DATA REPLICATION FOR PRESCRIPTION DRUG FRAUD PREVENTION AT A LARGE GOVERNMENT HEALTHCARE AGENCY

Paul J. Holenstein, Executive Vice President, Gravic, Inc.

Introduction

In most developed countries, managing the prescription drug process is one of the largest components of healthcare, regardless if it is run by a socialized or private agency. Every aspect of the process, from pharmaceutical drug manufacturing, to doctors writing prescriptions, to pharmacies dispensing prescriptions, to patient, doctor, pharmacy, and manufacturer reimbursement, requires constant oversight and control. This distributed process quickly becomes very complex, with many interactions between varying parties, as each prescription ultimately makes its way from the prescribing doctor to the patient receiving the medication.

Managing this process crosses many functional areas, including:

- **healthcare information management** (record, track, and secure private patient medical information)
- **finance** (match the payment for the drug and various reimbursements)
- **verification** (confirm that the right medication is given to the correct patient)
- **accounting** (check that proper statistics are generated for analysis and reporting purposes)
- **fraud prevention** (ensure that the correct patient uses the medication properly and not for illicit activity)

It is critical to put checks and balances in place and to integrate each functional area, or a shortcoming in one process could lead to a host of societal impacts, such as unwarranted side effects, disease, or even death.

Prescription fraud is a growing worldwide problem. In the U.S., one fraudulent scheme involved 243 people and 46 health professionals and defrauded Medicare of \$21.2 million (USD)¹. In Canada, one pharmacist cheated taxpayers \$471,000 (USD) in an elaborate methadone-billing scheme². In England, patient prescription fraud costs the country about \$6.7 billion (USD) per year³. According to the World Health Organization, \$415 billion (USD) is lost to healthcare fraud worldwide every year.⁴

In large healthcare agencies, the problem is often exacerbated by the absence of standardized case reporting and information sharing between various jurisdictions. In many instances, the fraud goes undetected until a later date, feeding the addiction problem, costing the government large sums of money in improper reimbursements, and creating a difficult, involved investigation of who did what, and when. In many cases, the perpetrators often are long-gone and cannot be found by the time the fraud is detected.

As the sheer amount of prescriptions proliferate, the complexity in tracking and protecting this sensitive information increases exponentially, opening the door to fraud and even waste that needs to be uncovered and eliminated⁵. For the largest agencies, billions of dollars (USD) change hands daily.

Two primary issues are avoiding prescription processing fraudulent prescriptions and reimbursements, and eliminating narcotic addiction and abuse. As the current opioid crises overwhelms the healthcare sector, agency administrators know that immediate action must be taken.

Prescription Claims Adjudication System - Overview

A publicly funded healthcare agency is handling all medical insurance claims for a large and diverse population. The systems track key components as prescriptions are written, filled, and money changes hands. There are numerous IT applications involved in the management of this agency. More specifically, an application running on a pair of HPE NonStop servers processes prescription drug claims. The pair are located in geographically separated datacenters (DC1 and DC2). These servers run in an active/passive disaster recovery configuration, and run HPE NonStop RDF for transactional data replication between them, keeping both active and passive databases synchronized. If the active system fails, a failover is performed and online processing is switched to the backup (passive) system. In normal operation, the backup system is used for read-only query and reporting type activities.

In order to attack the prescription fraud problem, administrators and systems analysts implemented a centralized claims adjudication application to consolidate and analyze all claims throughout the country's regional jurisdictions to flag suspicious activity. Typical types of suspicious activity the system is designed to detect include:

- Multiple claims for the same drug and patient over a short period of time [either from the same or different doctors/pharmacies]
- Multiple claims for similar drugs and patient over a short period of time [either from the same or different doctors/pharmacies]
- Multiple claims for the same or similar drug for the same patient address over a short period of time [either from the same or different doctors/pharmacies]
- A doctor or pharmacy receiving healthcare reimbursements for prescriptions that do not match those that they prescribed
- A doctor or pharmacy receiving manufacturer reimbursements for prescriptions that do not match those that they prescribed
- A doctor or pharmacy submitting for reimbursements for the same prescriptions through multiple (regional) jurisdictions
- A doctor or pharmacy prescribing/dispensing an atypical amount of controlled substances or receiving an atypical amount of reimbursements for their size

This new claims adjudication application runs on a remote system while it collects all of the pertinent data from the existing NonStop-based prescription drug processing system and accesses a new decision support database containing all current and historical claims.

One of the project requirements for this new application is

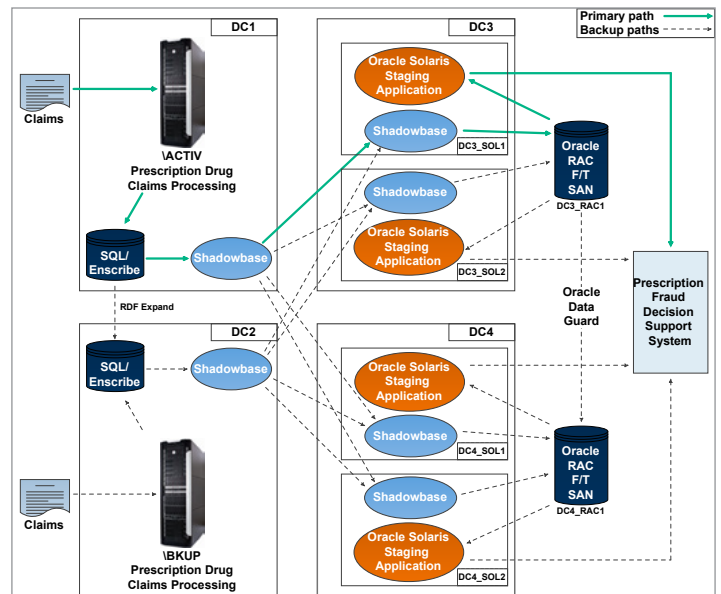


Figure 1 - Prescription Fraud System Architecture

to keep the decision support database current with all active claims being processed, while continuously maintaining high availability for the application and data. This requirement is further complicated by the fact that the NonStop system pair is managed by a service provider, and the new claims adjudication application is managed by the government healthcare agency, with neither party having direct control of the other. Additionally, the new claims adjudication application must have minimal impact on the existing claims processing environment, because the government does not want to upgrade its existing infrastructure or systems to handle any additional material load.⁶ Consequently, the fault tolerant characteristics of the whole application required automation with the least minimal impact. HPE Shadowbase software was key in meeting the project's goals.

Prescription Claims Adjudication System - Detailed Architecture

Figure 1 illustrates the application configuration where claims are sent to the active NonStop server running in datacenter 1, DC1 (\VACTIV system), with HPE RDF replicating over Expand to a backup NonStop server running in DC2 (\BKUP system). The government agency installed intermediate staging systems [comprised of Oracle Solaris servers with an Oracle RAC database], which feed the new decision support system and run the claims adjudication application.

- For local fault tolerance, i.e., to provide localized "high availability", the government agency installed two Solaris servers, each accessing the same Oracle RAC database via a shared fault tolerant SAN. At any one time, only one of the Solaris servers is actively processing requests, with the other acting in a passive backup role.
- For geographic fault tolerance, the Solaris staging systems are duplicated, running in two different datacenters (DC3 and DC4).

Oracle Data Guard maintains synchronization on the two Oracle RAC databases running on duplicated staging systems (DC3_RAC1 and DC4_RAC1). At any one time, only one of the Solaris systems is active, with the other three basically acting as local and remote backups. The government agency also developed a private

application to feed data from the staging system into the support system for adjudication decision processing.

Prescription drug claims transactions are processed by applications running on the active NonStop server, which writes the necessary claims information into TMF audited NonStop SQL and Enscribe databases. The NonStop maintains the “database of record.” HPE Shadowbase software is installed and operates on both NonStop servers and all four Solaris systems. The software is configured for uni-directional replication from each NonStop system to each Solaris system and its Oracle RAC database (with only one path being active at a time). Shadowbase architecture replicates change data as it is generated by the NonStop claims processing application to the primary active Solaris/Oracle RAC system (DC3_SOL1), thereby keeping the database of record and the staging database synchronized.

Highly Available, Secure System Architecture

Shadowbase transmission paths are configured between both of the NonStop systems and all four Solaris systems. The staging database is continuously synchronized with the database of record for the following failure scenarios, all of which were rigorously tested and passed the agency’s acceptance criteria.

- 1. Failure of the active NonStop system – Shadowbase software on the backup NonStop system detects this condition and takes over, continuing replication to the primary Solaris system (DC3_SOL1).** During normal operation, Shadowbase software on the backup NonStop system recognizes when RDF is replicating to the configured files and tables, and therefore does not replicate such changes to the Solaris staging database until it detects a failover. As a result, it is possible to have Shadowbase software up and running on both NonStop systems simultaneously. As soon as an RDF takeover performs and the backup NonStop system becomes the active system, Shadowbase software recognizes that RDF replication has stopped, and automatically begins replicating data changes from the backup to the target Solaris system.
- 2. Failure of active Solaris system [local] – Shadowbase software detects this condition and switches to the standby local Solaris system (DC3_SOL2 or DC4_SOL2, as appropriate).** At all times, Shadowbase software is running in standby mode on the alternate/backup Solaris systems. When the primary Solaris system fails, an automatic switchover activates via the use of an intelligent network aliasing method, and the IP address of the target system switches to the alternate/backup Solaris system. Consequently, Shadowbase source replication automatically switches to the alternate/target Solaris system. Whenever replicated data appears on the alternate Solaris system, Shadowbase software will automatically begin applying it to the Oracle RAC database.
- 3. Failure of staging system pair (e.g., whole target datacenter outage) – Shadowbase replication switches to the remote staging system (DC4_SOL1), which takes place the same way as a local system failover – leveraging the intelligent network router.** In this approach, the actual failure is transparent to the Shadowbase software running on the source environment.

The claims data being processed contains sensitive information (patient names, addresses, medical history, etc.), and must be protected against unauthorized disclosure. For this reason, all of the data at rest and in motion is encrypted as it moves across this entire application chain. This process includes the data:

- **from entry into the active NonStop system,**
- **to the data replicated by RDF to the backup NonStop system,**
- **to Shadowbase replication to the staging systems,**
- **to the forwarding of the data to the decision support system.**

For additional protection, once the data is replicated to the decision support system, it is deleted from the staging database. This requirement greatly increases the project’s complexity as multiple layers of encryption and key management are in effect.

Initial Database Load with No Application Outage

Before the new claims adjudication application is deployed, all of the historical and current data residing in the NonStop database of record needs to be made available on the staging system, to load the empty Oracle RAC database instances. This initial load has to be done while the claims processing application is active (i.e., the initial load and online replication has to be performed in parallel, with no outage of the claims processing application running on the NonStop systems).

HPE Shadowbase Online Loading (SOLV) software is designed for this very task. SOLV supports online loading of a source database into a target database, meaning that the source and/or target databases can be open for reading and updating while the load occurs. There is no need to take either the source or target databases offline and make them inaccessible to applications while the load proceeds. SOLV automatically handles the merging of the data being loaded with the data being replicated. Thus, by using SOLV, the staging database is fully loaded while the claims processing application on the active NonStop system remains available and processing transactions.

SOLV achieves zero application outage while the load synchronizes by tight integration with the Shadowbase data replication engine (Figure 2). SOLV is reading and replicating data from the source application database itself (the database of record on the active NonStop system). Meanwhile, the Shadowbase data replication engine is replicating change data from the audit trail (the data actively being updated by the claims processing application). The key transactions happen when the SOLV data and the change data are merged together and applied to the target database (one of the Oracle RAC staging databases). Any changes made to the SOLV data while it was in transit are identified and applied. Hence, the data in the target database is a current and consistent copy of the source database, even as it is being updated by the online claims processing application.

This approach is powerful as there is no ‘outage’ of the target environment required, and the target data is fully consistent with the source data as it loads. There is no large queue of change data that builds up while the load takes place (that must later be replayed) to bring the target consistent, because the queue of change data is actually consumed while the load occurs.

Additionally, the SOLV approach can run in parallel with full source (and target) application processing for a continuous (and extended) timeframe, meaning that the load/merge operation can be set up and run in parallel until the load completes. Massive quantities of data can be loaded, and there is never an outage of the application environment required during this process.




Figure 2 – Shadowbase SOLV Architecture

Summary

Prescription drug claims fraud is big business, costing the country billions of dollars. To help stem the losses, a centralized claims fraud adjudication and prevention decision support system is implemented. This system utilizes various analytical techniques to identify, prevent, and recover payments for fraudulent claims, as well as provide evidence to support criminal proceedings against fraudulent claimants.

The decision support system relies upon data gathered from a primary claims processing application running on an HPE NonStop system. This data is required for real-time replication from the NonStop system to an intermediate Solaris/Oracle RAC staging database, in preparation for being loaded into the decision support system. The country’s government healthcare agency chose HPE Shadowbase data replication software to provide this functionality.

The system architecture also requires data protection via multiple levels of encryption at every point in processing the data, and high availability so that any one component’s failure would not prevent claims processing from occurring. In order to support this requirement, HPE Shadowbase SOLV is used to perform the initial load of the staging system. SOLV allows loading of the Oracle RAC staging database without any outage of the claims processing application. In addition, multiple Shadowbase replication paths are established between the various redundant systems so that if an active system or path fails, processing and replication can continue using alternate/standby systems and paths, keeping the claims data current on the fraud adjudication and prevention decision support system at all times.

Now that the new fraud adjudication and prevention decision support system is in place, the agency is looking at improving its overall NonStop and Oracle database availability by leveraging the HPE Shadowbase Sizzling-Hot-Takeover (SZT) architecture. This architecture improves overall system availability more than the RDF and Oracle Data Guard active/passive model to eliminate failover fault risk and dramatically improve recovery time if a failure occurs.

The updated system now assures the government agency that its important healthcare data is fault tolerant, highly available, decreases costs, and allows law enforcement to monitor and investigate fraudulent claims activity. [CS](#)

¹ Ornstein, C. (2015, June 23). [Fraud Still Plagues Medicare's Prescription Drug Program](#), npr.org
² Howlett, K. and Andreatta, D. (2013, December 23). [Rampant Canadian Pharmacy Fraud Sign of a Broken System](#), theglobeandmail.com
³ Agency (2015, September 24). [NHS Losing Billions to 'Fraud by Doctors and Dentists'](#), telegraph.co.uk
⁴ World Health Organization (December, 2011). [Prevention Not Cure in Tackling Health-Care Fraud](#), who.int
⁵ The entire agency also realized unforeseen project benefits of reduced waste and optimized efficiency.
⁶ We appreciate management’s requirement that the new claims adjudication function cannot cost (much) to implement or require any significant investment in additional hardware or software. This request is fairly standard and is not new to cost-constrained IT departments. Fortunately, the solution is efficient, and there was sufficient headroom available to accommodate the new processing and data flows.

Paul J. Holenstein is Executive Vice President of Gravic, Inc. He is responsible for the HPE Shadowbase suite of products. The HPE Shadowbase replication engine is a high-speed, unidirectional and bidirectional, homogeneous and heterogeneous data replication engine that provides advanced business continuity solutions as well as moves data updates between enterprise systems in fractions of a second. It also provides capabilities to integrate disparate operational application information into real-time business intelligence systems. Shadowbase Total Replication Solutions® provides products to leverage this technology with proven implementations. HPE Shadowbase is built by Gravic, and globally sold and supported by HPE. Please contact your local HPE account team to answer any questions about the product. For further information regarding Shadowbase data integration and application integration capabilities that can assist in solving big data integration problems, please refer to the companion documents [Shadowbase Streams for Data Integration](#) and [Shadowbase Streams for Application Integration](#), or visit [www.ShadowbaseSoftware.com](#) for more information. To contact the author, please email: SBProductManagement@gravic.com.

FUTURE LEADERS OF IT
 Interested in investing in next year's future leaders?
 Make an investment that is certain to show a return for years to come!
DONATE TODAY!