On-Demand BI for Insurance Organizations

Lower the Total Cost of Ownership for Information and Accelerate Pervasive BI

Best Practices | White Paper Series

June 2008

COMPACT SOLUTIONS

On-Demand BI for Insurance Organizations: How to Achieve Pervasive BI Faster than Ever Before

Overview

Data. It's the fuel that runs the insurance business. But if you're like most insurance carriers, you're drowning in it. While data surrounds the business, ironically, it's not getting to the right people at the right time to add value—make the organization more nimble, more flexible, and more competitive.

As the volume of data required to run the business expands exponentially every year, so does the diversity of that data, and the frequency with which users need to interact with it. Competitive pressures, diversity of product lines, mergers and consolidations, and changing regulations as well as new technologies create imposing challenges to the insurance data warehousing executive.

Because the insurance business is increasingly complex and evolving, information is seen as a "silver bullet" solution. To gain a competitive advantage in today's markets, the business needs on-demand access to granular information to manage risk, price new products competitively, analyze the state of the business, and bring new products to market faster.

The IT executive and the data warehousing manager are tasked to ensure that all of that information can be made available to users. For them to make the best decisions, they need to be able to analyze historical data—often on millions of transactions using complex statistical models—to understand exposure, loss, and pricing. Unfortunately, from an IT perspective, large amounts of historical data require more refined models. The mainframe, where so much of this data currently exists, makes an inefficient platform. That requires moving large amounts of data onto more efficient and cost-effective UNIX-based platforms.

> ["It has been demonstrated that non-quality data can cause business losses in excess of 20 percent of revenue and can cause business failure."]
>
> -- INFORMATION IMPACT International

The question is: how do you achieve this quickly and cost-effectively?

This white paper outlines a strategy IT executives and data warehouse managers can quickly and effectively deploy to move large volumes of legacy data to UNIX platforms, optimize that data for efficient and timely access, and empower users with pervasive business intelligence.

Integrating Data, Delivering Information: Key Challenges

The insurance business is increasingly complex and evolving. To gain a competitive advantage in today's markets, you need on-demand access to granular information to manage risk, price new

products competitively, analyze the state of the business, and bring new products to market faster. To ensure the best decisions, you need to be able to analyze historical data—often on millions of transactions using complex statistical models—to understand exposure, loss, and pricing. But large amounts of historical data require more refined models. And all of that is compounded with other pressing needs, two of which we'll review in lengthier detail:

- The explosion of transactional data provides a significant challenge for business intelligence and data warehousing for insurers.
- Multiple legacy product/source systems are not integrated so decisions are delayed and productivity is lost.
- Multiple silo'd data separated by function where every unit has its own set of data that may not agree with any other business unit.
- Legacy of legacy data stores—tens of terabytes big—and growing every year.
- More new products being introduced with new innovative pricing and innovative ways to get that product and pricing to customers (via internet portals).

["Explosive data growth demands serious attention . . . digital data volumes are growing by 30% a year and will be approaching 1 zetabyte by 2010. "]

–Forrester Research, 2007

Let's review two of these challenges in more detail: legacy data and distributed data silos.

Rich Granular Legacy

Vast stores of legacy data for policy premium and claims data challenge IT managers. In some cases you're managing 20 terabytes or more and that volume is growing every year. Multiple legacy product and source systems are not integrated, creating data integrity questions and delays in decision making.

In fact, most of the rich history needed for in-depth product pricing and underwriting is lodged in legacy databases. These are time-consuming and expensive to access—especially when it comes to developing new products, pricing them attractively for your customers, and getting to market before your competitors. The solution is to provide key users, Pricing Analysts, Product Operations, Marketing , and others with fast, easy and efficient access to that valuable historic information—but in a format that can be easily and quickly accessed.

Where's the Data?

You're faced with duplicate data in functional silos—every department seems to have its own set of data that may or may not agree. Often the information on which key performance indicators (KPIs) are based may not be as reliable as they need to be. How do you run a business if every unit has a different version of the truth?

The silos of legacy data, of functional data streams, won't fully support a nimble business and the time between asking for information and getting the "intelligence" is getting shorter and shorter. In the age of the internet, it's still common to find actuaries, marketers, claims, and risk managers all working from silo'd data, all waiting hours on a report only to do it all over again when they examine the data and need to make a modification.

Solving the Challenge: High-Performance Solution Maximizes Speed to Value

The optimized data warehouse (ODW) solution will provide the ability to deliver self-service business intelligence across the business—from claims management and policy administration, from underwriters to actuaries, and from marketing to risk analysts. These disparate business units can become more closely linked by leveraging the same historical data. This ensures "one version of the truth" for better more accurate pricing and decision making.

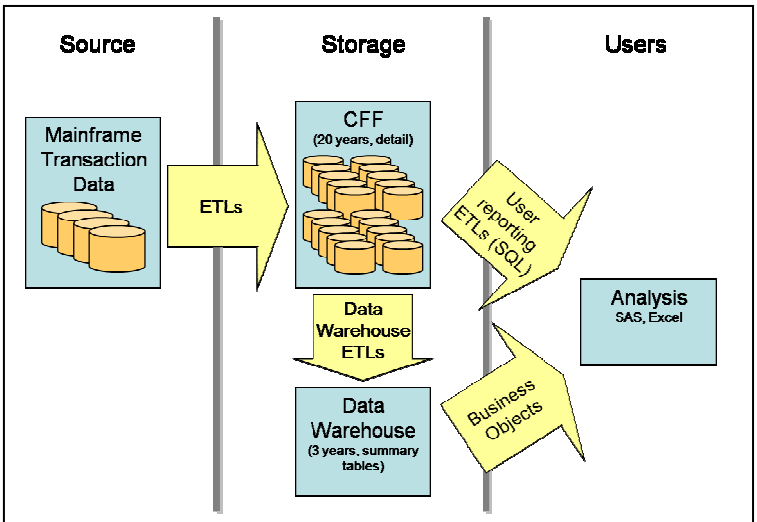

Using compressed flat files delivers an innovative, cost-effective high-performance solution that speeds up access to legacy data stores and the granular data they contain and efficiently provides pervasive BI.

The ODW creates a rich, centralized data repository that integrates core claims and policy and premium transaction data and makes it ready for high performance analysis on demand. This solution transforms tens of terabytes of legacy data into real-time intelligence that can be quickly and easily accessible with BI tools. The business can expect decision making to accelerate and the quality of those decisions to deliver higher profitability.

Specifically, the solution delivers an information architecture that dramatically improves the overall business value of your information. This solution:

- Integrates data from silo'd operations and applications, mainframe legacy databases, third-party sources, to enable a complete picture of the market/customer landscape.
- Re-platforms expensive legacy data residing on mainframes to less expensive distributed open systems.
- Dynamically integrates disparate data sources into one repository to deliver one version of the truth.
- Optimizes the solution for multi-dimensional analysis and modeling tools so analysts can easily drill down and go deeper.
- Scales easily to handle data volume growth.
- Uses the best of breed technologies—open systems platform bringing the data closer to users, no more waiting, no more expensive database storage costs.

Data Warehouse Optimized for Insurance Organizations

The Optimized Data Warehouse design innovates outside of the traditional data warehouse box. For example, to put 20 years' worth of legacy transaction data into a traditional data warehouse would result in an inefficient, 40-80 terabyte monster that would sap resources and processing time. But what's the alternative?

The design behind an Optimized Data Warehouse for Insurance Organizations re-thinks the traditional DW approach to legacy data. It uses compressed flat files to store historical data instead of a conventional relational database. In addition, it provides for the ability, with a SQL query, to retrieve any amount of data based on any criteria the end users wanted, something that can be made available to other systems within the enterprise to facility data integration.

These developments can deliver order of magnitude faster access to data as well as reduction in system maintenance costs. For example, one implementation used compressed flat files in a highly parallel applications designed in an Ab Intio ETL tool, resulting in extremely fast processing of 12 Terabytes of historical data.

Series of Optimized Data Repositories

The ODW solution creates a series of data repositories without using a relational database system. These time series repositories are optimized for fast and complex queries to feed millions of transitions to modeling tools, like SAP and SPSS, etc. This approach speeds up query response time by positioning pre-calculated complex measures within the data repository at frequently used aggregation levels. By understanding what the most frequent calculations are (such as loss, exposure, earnings) and performing them within the repository, after merging policy and claims transactions, significantly improves query response time.

Improved Data Quality

The integrity of the ODW design solution is largely dependent upon the quality of the stored data. As part of balancing and verification process, the records are compared with the legacy system's data. In case of differences the solution re-examines both systems with the common source data system. It is easier to think of this as a triangle: source system, legacy target, and ODW solution target.

Obviously a bi-product of finding discrepancies between the source system and the legacy target system and fixing them means that the overall quality of historic data is significantly increased. This in turn will help in more accurate business analysis, reporting, and modeling efforts.

Enhanced Analytics

Typically, access to a large volume of historic data on a legacy system requires many steps. This in turn creates layers of users that specialize in data retrieval, reporting, analysis, etc. The ODW solution, however, will enable end users to create necessary data sets, in desired format and volume, in a single step. This significantly lowers the processing time, and puts their focus back on analysis rather than queries. For example, one ODW solution used a web based .NET interface to provide this capability and was able to produce CSV and SAS data sets in a matter of minutes. This leads to the ability

to do same day modeling analysis because the ODW solution removes the complexity of the access layer from the picture entirely.

New Technology at Lower Cost

Once the ODW is implemented, all of that historical data once mainframe-based is now stored as compressed files in an open environment. A query processing engine, such as the Ab Initio ETL tool, provides high performance with data and process-level parallelism. This overcomes a significant disadvantage of typical relational databases that are better tuned to process current normalized data at an aggregated level rather than transactional data with 20 years of history. From a tech perspective, compressed files have an open API and can easily be processed by other tools if needed. This eliminates the risk of tight coupling between the data format and data storage/retrieval layer.

One version of truth

The ODW solution is designed to consolidate multiple historic data sources to provide a single entry point for access to the time-series of transactional and aggregated insurance policy and claims data. For example, by re-platforming 20-30 year old legacy system data, the ODW strategy eliminates the silos and makes these data stores available for analysis and modeling with modern BI tools—something unavailable before.

Iterative approach

The ODW solution is designed to be implemented in multiple iterations. It establishes a working prototype early on, where all stakeholders can become continuously engaged in the project and where the team can try different ideas in action, selecting the “best of” what works. This approach brings credibility early by showcasing the ability to adapt to changing requirements quickly while maintaining high data quality.

Enabling Self Service Architecture

Service Oriented Architecture enables enterprise level integration and simplified access via published services. For large enterprises the challenge is to leverage their existing IT infrastructures and legacy applications prior to making significant investments in new applications. The ODW solution is an excellent example of how to take a legacy application—even a 20-year old one--and make it available via a web based interface. The design of the system allows any other system to have their requests serviced, regardless of their operating environment. This refocuses the issues away from compatibility to common interfaces which are available to every application.

Leveraging Technology Assets

The optimized data warehouse re-thinks the approach to legacy data. Its use of compressed flat files to store historical data, instead of a conventional relational database, had never been attempted before at this scale. For example, a SQL-like query capability can retrieve data based on any criteria the end users wanted and can be made available programmatically to other applications in the enterprise, allowing for easier data integration. In addition, use of ready-to-use compressed flat files in highly

parallel applications resulted in extremely fast processing of 12 Terabytes of historical data.

Agility

When dealing with legacy systems, the enormity of unknown areas, such as exact formats and business rules for data 20 or more years old, can lead to extended analysis periods and rigid but incomplete requirements. This is why an optimized data warehouse strategy should assume, from the start, that new requirements will emerge on an ongoing basis. Maintaining agility and adaptability will be paramount so the solution design should be modular, dynamic, and highly driven by metadata.

Results

Once implemented, this design approach can allow for a dramatic reduction of the time to information—by removing the latency of inherent in relying on mainframe-based legacy data. By re-platforming the data onto UNIX systems, insurers can realize substantial annual cost savings in hardware, software licensing, and maintenance. In addition, an optimized data warehouse solution can deliver top and bottom line results to insurers:

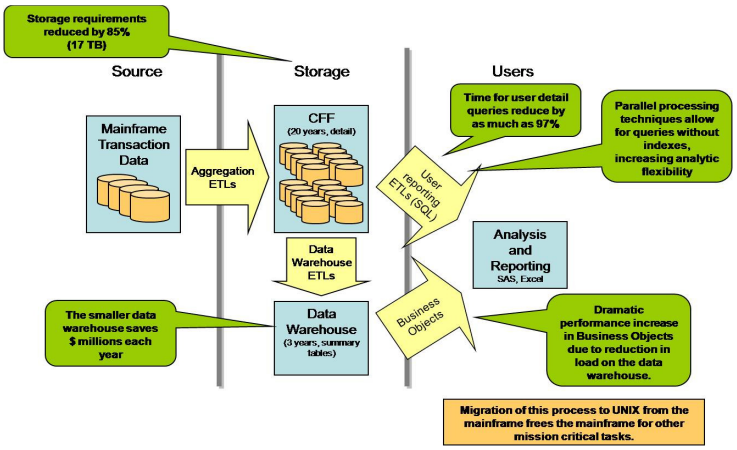

The optimized approach delivers cost-effective innovation and ROI throughout the process.

- Much faster access to integrated data for quicker analysis and modeling on granular data—that means faster time to market for you.
- Ability to include large swaths of historical data for greater accuracy in calculations, especially for risk assessment, earnings, loss, pricing, etc.
- One version of the truth to support improved regulatory compliance and much faster response to regulatory requests.
- Faster time to market because analysis can be completed in minutes or hours instead of days or weeks using mainframe-based legacy data now that data is directly fed to modeling and analytical tools.
- Reduced system maintenance costs for RDBMS licenses, DMA costs—millions per year in software license saves alone.
- Empowered users able to self-serve because they are closer to the data they need.
- Enabled IT governance strategies with standardized BI tools sets—functional groups all use the same tools to access the same information.
- Fully supports SOA initiatives by making fundamental policy and claims data readily available to applications.

About Compact Solutions

Compact Solutions delivers high performance information integration solutions for your industry, your enterprise, and your business processes. We help you realize the speed of real-time availability, the power of information, and the profit from better decisions. Compact Solutions experts offer data warehouse and enterprise data integration assessments that can help you understand how to optimize your data warehousing or data integration environment.

Compact has been helping Global 1,000 organizations solve the toughest data problems to achieve faster time to decision and action for greater productivity, competitive advantage, and profitability.

Visit www.compactsolutionsllc.com or email us at sales@compactsolutionsllc.com.

North America
Two TransAm Plaza Drive, Suite 400
Oakbrook Terrace IL 60181
+1 708 524 9500

Europe
United Kingdom
132 Burnt Ash Rd.
Lee, London SE12 8PU UK
+44 20 3051 1782

Compact Solutions Polska Sp. z.o.o
Ul. Wielopole 18b
Krakow 31-072
+4 812-429-1168