Is your business risking a USD 100 million loss?

How new technologies can help protect your data and your bottom line

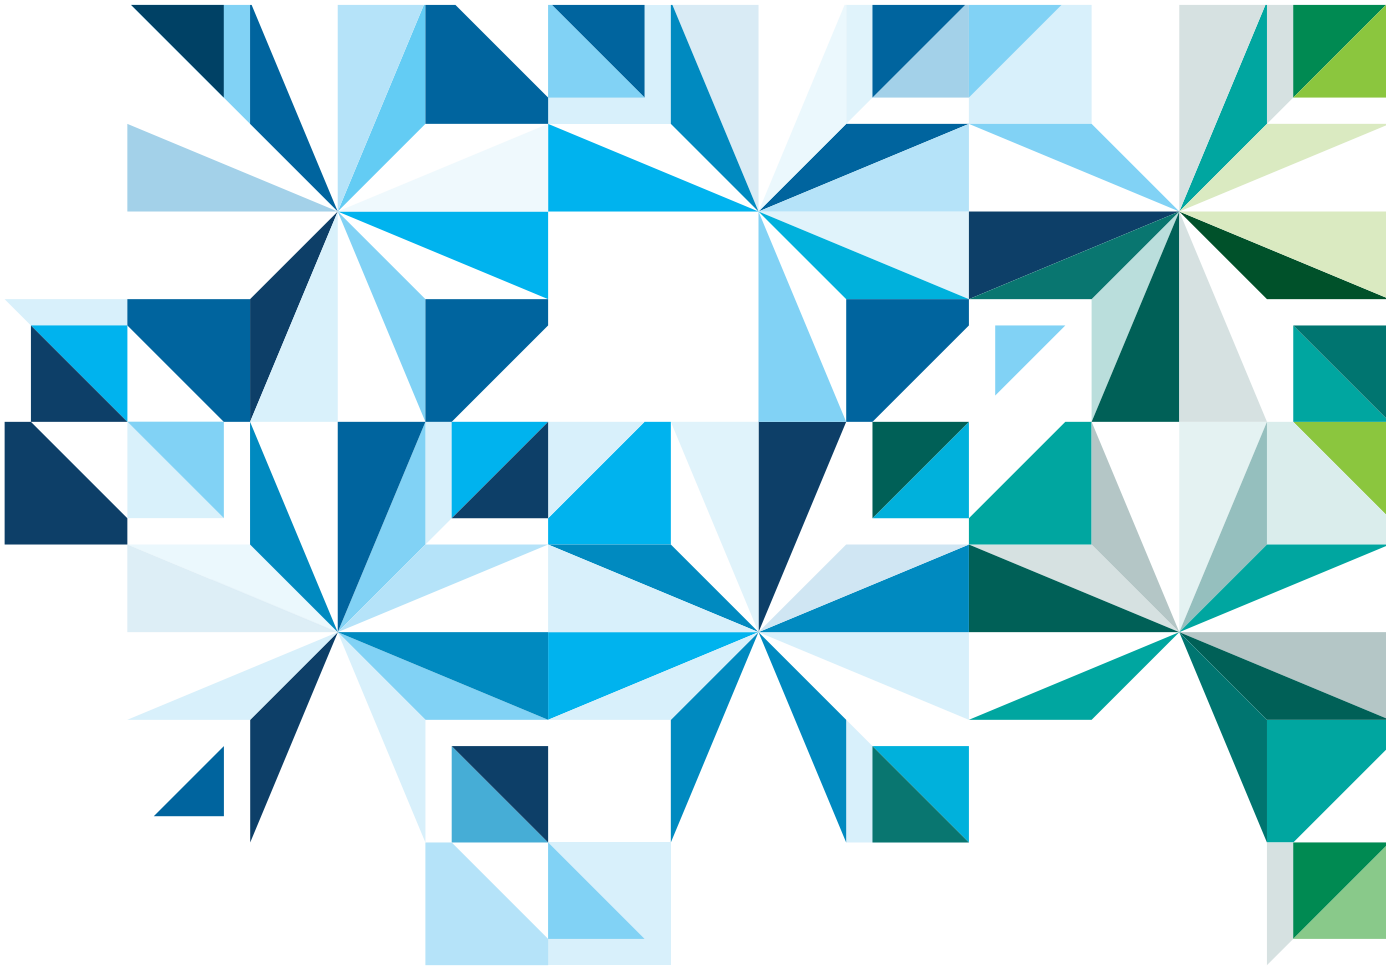

IBM

Contents

Your data is not safe

Technology trends and forces, including cloud, mobile and big data, can create large opportunities for your enterprise to exploit analytic insights. But the same things that enable these opportunities can skyrocket your risks if proper data security and governance controls are not in place.

As an example, in 2015 one of the largest health benefits companies in the United States reported that its systems were the target of a massive data breach. This breach exposed millions of records containing sensitive consumer information such as social security numbers, medical IDs and income information.

Various sources, including The Insurance Insider, suggest that this company's USD 100 million cyber-insurance policy would be depleted by the costs of notifying consumers of the breach and providing credit-monitoring services. And that policy payout doesn't consider other significant costs associated with a breach such as lost business, regulatory fines and lawsuits.

Cyber criminals now know the same thing every industry analyst knows—data is so important it has a value on the balance sheet. For that reason, every single industry has been attacked by hackers and experienced data breaches, including healthcare, government, banking, insurance, retail and telecommunications. Furthermore, once one company has been breached, hackers focus on other companies in that same industry hoping to exploit similar vulnerabilities.

So while this incident is one of the higher-cost examples, with a worldwide average data breach cost of USD 3.79 million,[1] coupled with the long-term brand damage, loss of faith and customer churn, consider this question: how exposed is your business to a similar type of breach? To answer this question, you must first ask, "Where does the data that feeds our analytics processes originate?"

The data origination challenge

For many enterprise clients, the answer to the data origination question is that the data comes from an IBM® z Systems™ mainframe. That is because organizations often run their mission-critical applications on z Systems to take advantage of its industry-leading service qualities such as availability, reliability and Evaluation Assurance Level (EAL) 5+ security.

Unfortunately, a large percentage of these organizations then weaken that security by replicating or transferring sensitive data off their IBM z Systems mainframe. They do this for a variety of reasons including realizing perceived cost savings, conducting analysis on massively parallel processing (MPP) systems, combining with data from external sources or satisfying end user technology preferences.

Whether it be quick, short online transaction processing (OLTP), analytical operational queries or highly concurrent, lighter-weight workloads, the IBM DB2® for z/OS® database has historically handled mixed workloads well. However, to run

CPU-intensive queries that involve processing large amounts of data in parallel, including analytics that group, sort and aggregate, organizations using DB2 for z/OS would typically employ one of the following options:

- Use data administration (indexing, partitioning, materialized query table) to address each individual query, which can be labor intensive.
- Purchase additional hardware resources, which can be cost prohibitive.
- Use a resource limit facility (RLF) or similar product to prevent these queries from consuming too many resources, which can shunt a potentially valuable or necessary workload.
- Extract data to disparate systems capable of performing the CPU-intensive parallel processing, which is the most common approach because it appears to be less costly.

The reality is that replicating or extracting data off z System mainframes has led to a proliferation of data repositories containing personally identifiable information and silos of disparate people, processes and infrastructure (PPI). Figure 1 illustrates this graphically.

For larger organizations, this structure can occur repeatedly. That is because each line of business (LOB) often has its own PPI that is typically managed and billed through chargeback. Each LOB then uses its own data repository, sourced from a z Systems platform, choosing PPI because it is available and familiar rather than the perceived challenge of employing the data from the source.

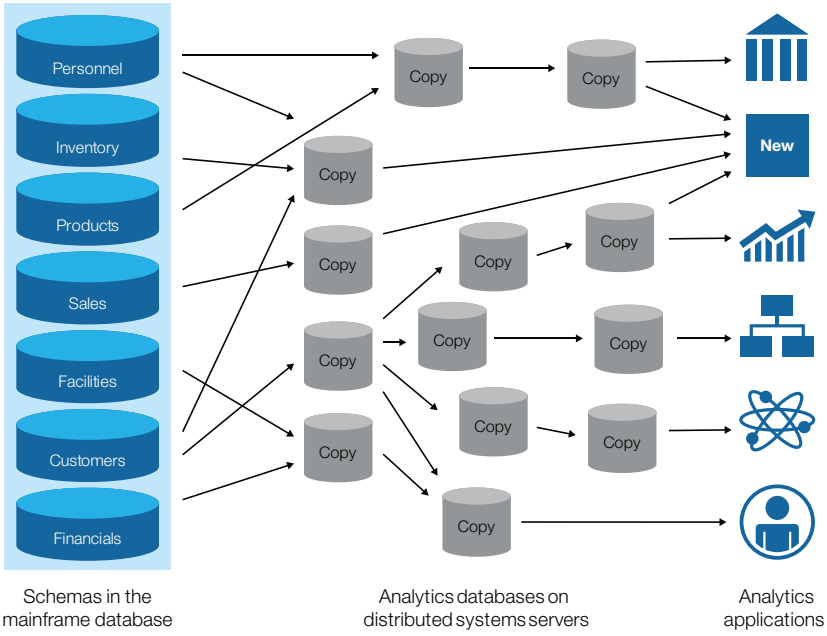

Figure 1. A replicated and extracted data repository structure.

While often the default choice, deploying replicated and extracted data repositories can increase your organization's risk exposure across the spectrum. For this discussion, we'll focus on the four most crucial risk areas: system costs, data security, system performance and data archiving.

System costs

Some organizations view reducing data breach liability as cost avoidance. Since IT architecture decisions are often made based on the best hard dollar cost solution, they may not account for cost avoidance best practices that include limiting the data breach threat. This mindset, along with the pervasive view that DB2 for z/OS is primarily useful for OLTP workloads, has contributed to the growth of disparate data repositories for analytics workloads.

New DB2 for z/OS technologies targeted at analytics applications are challenging this traditional view and transforming the hard dollar cost comparisons. As figure 2 illustrates, comparisons now often reveal the hard dollar costs to extract, transfer and load (ETL) or replicate data to a disparate environment for analysis exceeds the hard dollar cost to keep the data on z Systems and apply new technologies such as IBM DB2 Analytics Accelerator or IBM DB2 Value Unit Edition.

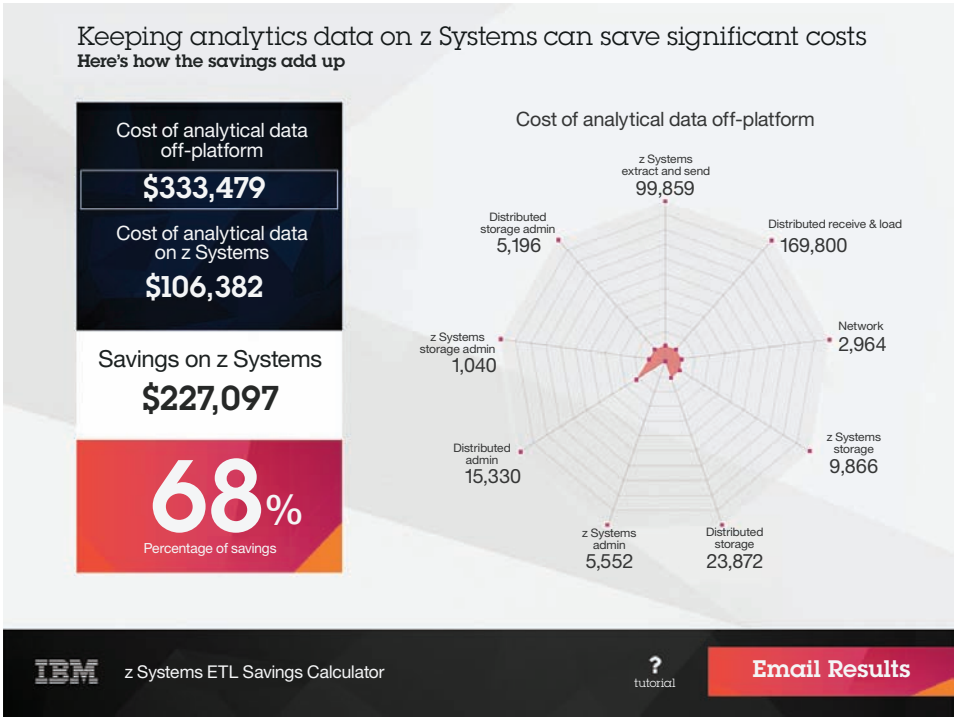

Figure 2. Potential cost savings compared to replicating 200 GB of data per day per database environment.[2]

While the USD 333,479 expense is substantial, it is just one instance. A typical LOB organization will have multiple instances of each data repository for unit test, system test, user acceptance test and production. Consequently, the costs—and the savings—are generally multiplied by a factor of four.

In this example, that means the traditional ETL approach costs almost USD 1 million more than executing analytics on the z Systems mainframe. For an enterprise with multiple LOBs, you can see how the extra annual costs for data extraction or replication can quickly add up to millions of dollars.

If you closely examine the cost components, you can see how this happens. In addition to costs usually accounted for in a hard dollar cost comparison, the true costs should account for all actual expenditures including additional managed infrastructure costs as well as additional IT administration costs.

For example, once data is replicated to a disparate environment, there are real managed infrastructure costs associated with using, administering and maintaining that environment. Many organizations rely on various internal or outsourced infrastructure teams to manage these environments and then bill the client through internal chargeback.

The IT administration costs come from development efforts to make the DB2 for z/OS source data usable and consumable. While typically not included in chargeback or actual data movement costs, these jobs have real costs associated with them. Examples of these costs include building ETL flows to extract the data from DB2 for z/OS, denormalizing the data, which includes combining data from different platforms and performing data transformations, adding surrogate keys and table lookups as well as loading the data into the remote data store.

Often, the cost of building and maintaining these flows, along with adjusting for detected changes, represents the largest ongoing and growing financial liability for the IT organization. As more end users are supported and data volumes grow, the cost to extract, transform, load and render that data explodes.

In addition, the change detection, change management, and change test and deployment activities become very challenging and therefore costly. Fortunately, the IBM Eagle Team has developed total cost of ownership (TCO) analysis tools that incorporate these expenditures into TCO analyses; additional details can obtained by emailing eagletco@us.ibm.com.

System cost solutions

Much of the data sprawl off z System mainframes was triggered by the traditional belief that z Systems were not the most cost-effective option for handling CPU-intensive analytic workloads, and that PPIs allowed for more efficient integration of data from external sources. New technologies are challenging this belief in a variety of ways.

For example, software costs can be drastically reduced using specialty engines such as z System Integrated Information Processors (zIIPs), z System Application Assist Processors (zAAPs) and Integrated Facilities for Linux (IFLs). In addition, z System CPUs typically exhibit 100 percent utilization in a mixed-workload environment. Most importantly, three IBM offerings can dramatically impact the economics of keeping source data in DB2 for z/OS:

- DB2 Analytics Accelerator
- DB2 Value Unit Edition
- z System Collocated Application Pricing (zCAP)

The IBM DB2 Analytics Accelerator for z/OS is a high-performance appliance for DB2 for z/OS. It deeply integrates the IBM z Systems infrastructure with IBM PureData® Systems for Analytics, which is powered by IBM Netezza® technology.

With the DB2 Analytics Accelerator, DB2 for z/OS can offload data-intensive and complex static and dynamic DB2 for z/OS queries, such as data warehousing, business intelligence and analytic workloads, transparently to the application. The DB2 Analytics Accelerator then executes these queries significantly faster than previously possible—all while avoiding CPU utilization by DB2 for z/OS.

The DB2 Analytics Accelerator price-performance curve also opens unprecedented opportunities for organizations to use data on their z Systems. It allows users to run workloads that historically were offloaded from z Systems, or run queries that were governed or shunted in DB2 for z/OS such as ad hoc queries whose performance characteristics are typically unknown at runtime. And IT administrators can allow DB2 for z/OS to choose where to run these queries, or they can force these queries to the DB2 Analytics Accelerator to prevent additional DB2 for z/OS consumption.

While the DB2 Analytics Accelerator is a separate piece of hardware, it is a logical extension of DB2 for z/OS because DB2 for z/OS manages and regulates all access to the DB2 Analytics Accelerator. This means DB2 for z/OS will continue to handle those workloads that make sense to run in DB2 for z/OS like OLTP queries and operational analytics, and it will reroute those queries that best fit an MPP technology to the DB2 Analytics Accelerator.

Beyond the pure analysis capabilities, the DB2 Analytics Accelerator also includes the ability to perform in-database transformations and multi-step processing using Accelerator-only tables. And with technologies like the IBM DB2 Analytics Accelerator Loader for z/OS, users can quickly load data from DB2 for z/OS along with non-DB2 and non-mainframe data directly into the DB2 Analytics Accelerator without interrupting access to production objects. This can reduce the need to denormalize the data, historically a major reason for offloading z System data.

While the DB2 Analytics Accelerator adds a cost-effective, secure solution for CPU-intensive workloads, IBM also offers a DB2 for z/OS licensing option called DB2 Value Unit Edition (DB2 VUE). With DB2 VUE, clients can deploy DB2 for z/OS using a one-time charge (OTC) model as opposed to the traditional monthly licensing charge (MLC) model.

Taking advantage of the OTC approach helps shift the cost of qualifying new workloads from an operating expense (OPEX) to a capital expense (CAPEX). This allows you to consolidate new workloads back to DB2 for z/OS, and manage those workloads as they grow to meet the demands of end users—all without increasing MLC costs.

The third offering that influences the costs of executing data analytics is z System Collocated Application Pricing (zCAP). zCAP can improve the cost of deploying new z/OS applications by allowing new applications to be deployed in existing logical partitions (LPARs) but priced as if they were running in dedicated LPARs. More information is available at https://ibm.biz/BdXQe9.

This enhancement to sub-capacity reporting removes up to 100 percent of the new application's general-purpose processor time from the machine utilization values reported for other middleware, and up to 50 percent for z/OS.[3] The result is you only pay for the new application's direct usage of the IBM middleware that manages the processing of the new application—such as CICS, IBM WebSphere® Application Server, IMS and DB2 for z/OS—while mitigating the impact on the million service units (MSU) reported for other programs running in the same partition.

Data security

As the health benefits company data breach illustrates, the practice of extracting or replicating data to disparate PPI systems poses serious security concerns. And each PPI instance substantially increases your potential liability from a data breach.

By definition, creating more copies of sensitive information increases the risks of a data breach. This is because each LOB may have its own data encryption policies, its own policies on credential governance, its own policies on disseminating analyzed data and its own level of policy enforcement of these policies. Perhaps more important, each copy of that data has its own set of interfaces for accessing that data.

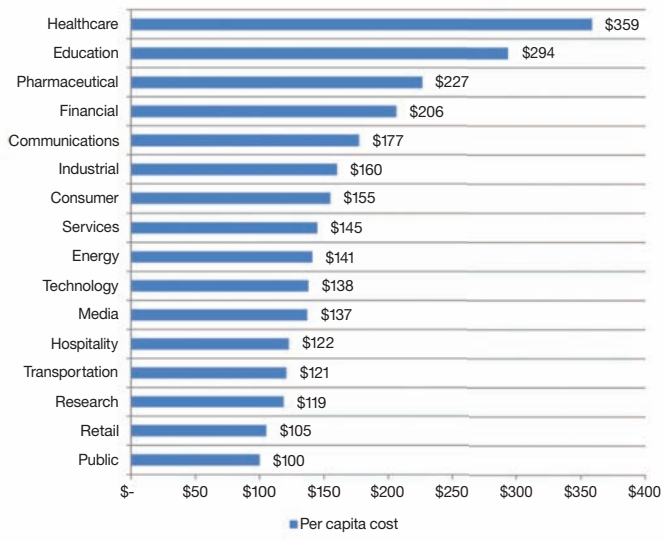

Figure 3. The per capita data breach cost by industry.[4]

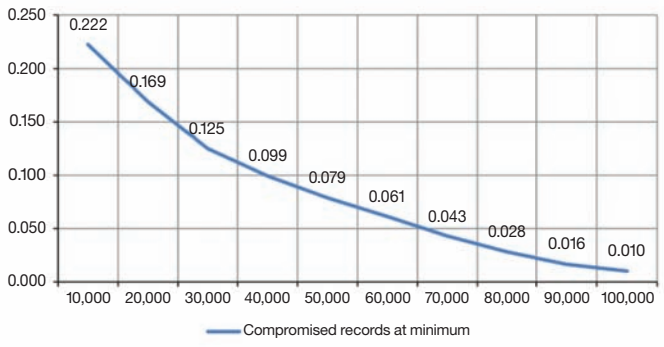

Figure 4. The probability of a data breach over a 2-year period.[5]

So, typically there will be a separate set of authentication mechanisms for each data store. For example, DB2 for z/OS data is protected via IBM RACF® security. But a copy of that data in a distributed database management system (DBMS) will typically be protected using a separate lightweight directory access protocol (LDAP) server that may be combined with DBMS credentials. Each copy then presents an additional potential entry point for a hacker.

Example healthcare company analysis

As an example, let's use findings from the Ponemon Institute's report, "2014 Cost of Data Breach Study: Global Analysis," to illustrate the cost of a 10,000-record data breach risk to a healthcare company.

The per capita data breach cost = USD 359 per record, from Fig. 3.

The probability of a data breach per year: 22.2 ÷ 2 = 11.1 percent per year, from Fig. 4.

At 2 KB per record, data breach costs are USD 0.18 per byte, or USD 180,000 per MB.

For each 10,000 record, or 20 MB, the average data breach cost is 20 MB * USD 180,000 per MB * 11.1 percent probability per year = USD 400,000 per year.

For a company supporting 10 distributed databases that have been extracted or replicated from a DB2 for z/OS system, this can equate to an estimated increased annual cost of USD 4.0 million due to security risks.

Data security solutions

You can see how minimizing the PPIs can benefit your bottom line. Because the DB2 Analytics Accelerator is an integrated component of DB2 for z/OS, users do not need separate credentials to access it. In other words, users don't get separate credentials to a separate data repository; rather, all access is controlled through DB2 for z/OS and through RACF security.

Conceptually, with the DB2 Analytics Accelerator, the z System operates analogous to a hybrid automobile. A hybrid automobile has a standard user interface comprised of a steering wheel with accelerator and brake pedals. But at any given moment, the hybrid automobile may be using its gasoline or electrical power source to meet the driver's current needs; the switching is automatic. In much the same way, the addition of the DB2 for z/OS with the DB2 Analytics Accelerator automatically routes data to the optimal processing engine to meet the users' needs—using a standardized interface without user intervention.

Beyond the automated optimization, the N300x hardware used with the DB2 Analytics Accelerator includes self-encrypting drives (SEDs). Also known as hardware-based full disk encryption, these drives implement an encryption algorithm in hardware within the disk drive. Using symmetric cypher AES256, all data written to the drive is encrypted at full interface speed before writing it to media and then decrypted before it is returned.

The symmetric disk encryption key (DEK) is unique to each drive and is stored in a secure, tamper-proof location on the drive, so it never leaves the drive. By default, the SED is in its unlocked state. During installation, an IBM field engineer will activate the encryption. Once this is done, data on the drive is always encrypted. Using the DEK and powering the disk on requires unlocking the DEK with the authentication encryption key (AEK), which means the drives cannot be unencrypted. For additional security, your service provider can manage or change the AEK encryption keys.

Since no system is immune, the inevitable question arises, "What if our DB2 for z/OS credentials are stolen?" To start, most z Systems customers already have sophisticated change management and access control policies in place through RACF and system authorization facility (SAF). However, you can deploy an additional layer of fraud detection by using a technology like IBM Guardium®.

With Guardium, you can perform outlier analysis to detect suspicious behavior against a norm for each user and block the transaction or thread, thus taking a more preventative and proactive approach to protecting data. In other words, Guardium can be used as a form of predictive whitelisting to identify anomalies in a user's behavior relating to sensitive data, and determine whether those anomalies may be related to a credential or data breach.

In addition to these new technologies, IBM also provides comprehensive data tools including discovery, classification and vulnerability assessments as well as real-time monitoring, alerting and blocking of suspicious behavior. These tools can be applied across many databases, data warehouses, off-the-shelf commercial applications, Hadoop clusters and more. The reason is simple—a proactive approach to data security is far less expensive than the millions it can cost to resolve a data breach.

System performance

Another reason for ETL proliferation is perceived performance improvements that could be realized by segmenting workloads based on data access characteristics. For example, OLTP workloads could remain on z Systems while extremely CPU-intensive workloads could be run on an MPP system. This leads to multiple copies of the data for different access method patterns, multiple instances of PPIs, increases in the number of data interfaces, application changes to point at different environments for different queries and more.

While many believe the ETL structure is efficient, our clients tell us differently. Many organizations extract, transfer and load data to a downstream data warehouse so they can execute long-running analytical queries to identify fraudulent and improperly rejected transactions.

One client reported it required over 24 hours to extract, transfer and load content into the data warehouse, then 11 hours to execute their fraud detection analytic suite. The 35 hours of latency not only increased their risk of financial losses due to fraud, it diminished customer satisfaction.

In another example, banks must continually assess and manage their market and asset liquidity risk to maintain minimum asset ratios for regulatory compliance and customer obligations. It is impossible to run analytics on real-time or near real-time data using legacy technologies alone. This forces banks to run nightly batch jobs to assess market and asset liquidity risk. These nightly runs are often too late to address liquidity events.

System performance solutions

Multiple banks have implemented the DB2 Analytics Accelerator and can now assess their trading risk positions in near real time rather than waiting for the overnight batch to run. This dramatically reduces each bank's liquidity risk and improves return on equity.

Another client, a large North American managed healthcare company, turned to the DB2 Analytics Accelerator because of its built-in MPP capabilities. As a result, they saw queries run up to 1908 times faster. Yet another client, an American investment management company, accelerated the delivery of canned reports using real-time data—realizing up to a 700 percent performance improvement in queries on intra-day activity and transaction volume.

In the past, it wasn't cost effective to offer this capability or it required hard coding in applications as part of a development effort. Giving DB2 for z/OS the flexibility to run queries in DB2 for z/OS or on the DB2 Analytics Accelerator eliminates the need to manipulate data specifically for analytic workload performance or to manually segregate workloads across multiple interfaces to force specific access methods.

This can increase your system performance and simultaneously reduce your technology stack, helping improve your customer satisfaction by giving you visibility into what is happening in near real time and allowing more flexibility around ad-hoc inquiries. The result is more analytics possibilities and bottom line benefits without increasing your risk of data breach.

Data archiving

With data quantities exploding, archiving has become an even more important component of your data management strategy. While active data that is still changing remains available in source systems, unchanging historical data is generally archived for future use.

Driven by cost, historical data is typically archived to less expensive, tape-based solutions. This includes archives that remain *online* due to auditing requirements and archives stored *offline* for more permanent storage.

While not being actively updated, the reality is that most archived data must still be analyzed by end users. For *online* archives, data access experiences a latency that eliminates real-time analysis options. For *offline* archives, analysis typically requires nightly batch processes to restore the data; this limits using this archived data to trend-type analyses. Figure 5 illustrates these flows.

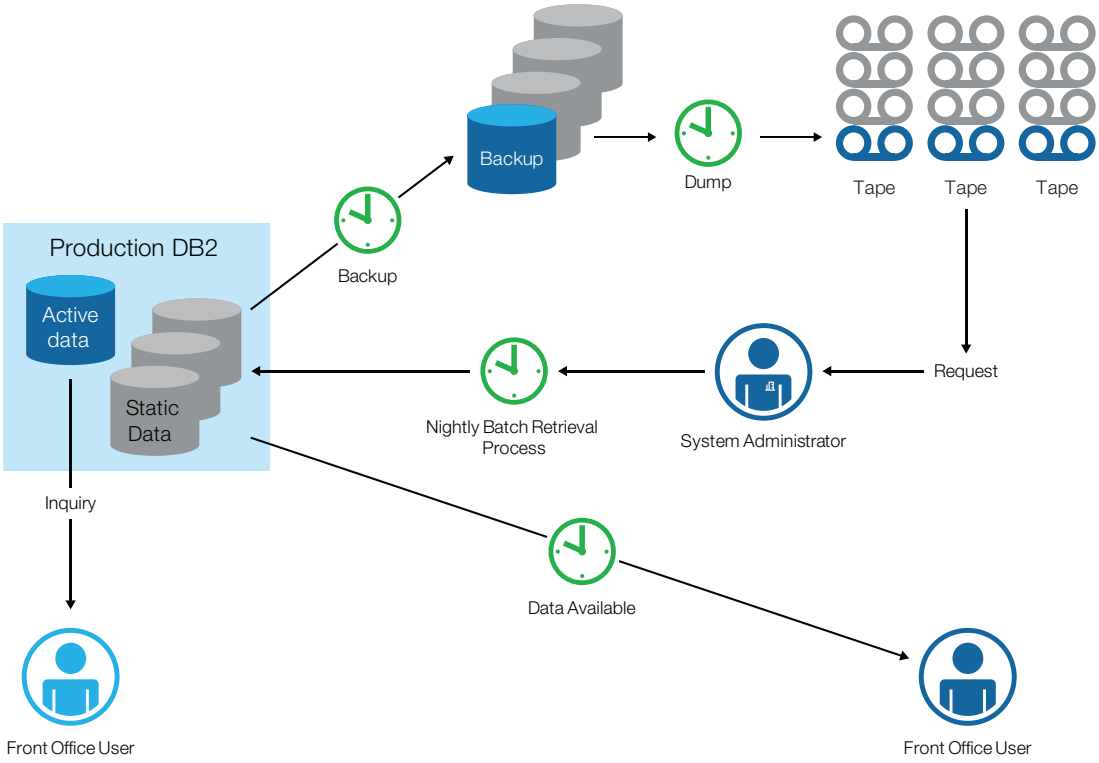

Figure 5. A typical process for analyzing archived data.

Data archiving solutions

To address these challenges, the DB2 Analytics Accelerator includes an online archiving functionality called a high performance storage saver. Using this functionality to move historical data into the DB2 Analytics Accelerator reduces the data volume in the DB2 for z/OS table; this means smaller indexes and smaller materialized query tables.

Just as important, moving this historical data into the DB2 Analytics Accelerator enables active analysis without any restore process. Access to both the historical and active data is managed through the same DB2 for z/OS interface and is protected through RACF security as opposed to interacting with the tape archive or copying the data to a separate DBMS and maintaining a separate data interface. And archive data durability is ensured through existing backup/recovery strategies.

By employing an HPSS, both active and online historical data are always available to the user for analysis and auditing purposes—with near zero latency for online historical data queries. The DB2 for z/OS optimizer manages traffic and directs any queries that touch archived data to the DB2 Analytics Accelerator, while queries against strictly non-archived data follow the DB2 Analytics Accelerator's usual query routing criteria. This means there will be near zero CPU consumption associated with the analysis of the historical data.

In addition, you can realize cost savings by removing any batch jobs that take data from tape and load them in online tables. The size of online tables in DB2 for z/OS can then also be reduced by moving historical partitions to the DB2 Analytics Accelerator. A sample catalog query to uncover archiving opportunities in a DB2 for z/OS environment is available at https://ibm.biz/BdX3gb.

Performance, security and savings

New technologies from IBM are transforming large enterprise analytics. The traditional approach of using z Systems mainframes for mission-critical applications, and then extracting or replicating data from those systems to distributed databases for analysis, is losing its luster.

Now, DB2 for z/OS combined with the DB2 Analytics Accelerator gives you the ability to execute OLTP workloads and perform MPP-type queries on data—while that data remains within the DB2 for z/OS environment. That means you can obtain the competitive advantages that powerful analytics can provide—often with near-zero latency—without incurring the increased costs and elevated security risks associated with data administration, additional hardware, or ETL or replication approaches.

How IBM can help your bottom line

For over a century, IBM has been pioneering technologies and providing services that help companies manage and mine their valuable business data. For 22 consecutive years, IBM has topped the annual list of US patent recipients—receiving a record 7,534 patents in 2014.[6] And by investing USD 24 billion in research and development (R&D) and acquisitions, IBM has established the world's deepest portfolio of big data and analytics technology that spans R&D, solutions and software.[7]

Using this technical expertise, IBM has created a free cost-benefit analysis (CBA) workshop to help you visualize the hard dollar cost savings and efficiencies that can be gained by reducing data sprawl within your IT infrastructure. With the goal of avoiding data breach costs by improving security, while using a hard dollar justification, the CBA workshop considers all functional and non-functional requirements for the existing landing zones, data stores and the feeds they support.

Using those requirements, the CBA workshop experts identify options for consolidating your IT infrastructure using solutions such as DB2 for z/OS sources, the IBM DB2 Analytics Accelerator, the IBM DB2 Analytics Accelerator Loader for z/OS, IBM Cognos® Business Intelligence, IBM SPSS® predictive analytics, Hadoop, additional mainframe hardware and software, as well as suggested deployment and quick start services. The experts then perform a CBA and value analysis using the proposed solution architecture to compare the cost of maintaining distributed databases against working with the data at the source, and finally can recommend a holistic approach to improve your technical architecture and help reduce your costs.

For more information

To learn more about DB2 Analytics Accelerator, or to request your free cost-benefit analysis workshop, please contact your IBM sales representative or visit the following website: https://ibm.biz/BdXTz4

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. For credit-qualified clients we can customize an IT financing solution to suit your business requirements, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: **ibm.com**/financing

About the author

Shantan Kethireddy is an IBM Master Inventor and a technical sales specialist for Information Management solutions on z Systems products. In this role, he is responsible for leading consultative business value validation engagements for big data and analytics. He holds a Masters in Computer Engineering and Electrical Engineering from the University of Iowa and possesses two dozen patents, primarily focused on data-centric technologies.

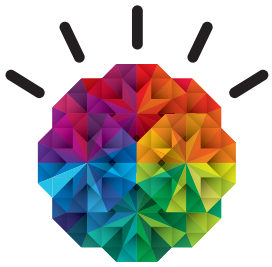

Systems Group
Route 100
Somers, NY 10589

Produced in the United States of America
September 2015

Please Recycle

[1] Poneman Institute. "2015 Cost of Data Breach Study: Global Analysis." May 2014.

[2] This is based on an IBM internal total cost of ownership (TCO) calculation using the z Systems ETL Calculator App, which generates cost comparisons using customer-specific inputs. For this comparison, data usage was assumed to be 200 GB per day over a one-year period, using an ODS with one datamart with the z/OS operational software stack at medium complexity. More details on the ETL Calculator App, this specific example and customer-specific savings can be obtained by contacting the IBM Eagle Team at eagletco@us.ibm.com.

[3] IBM. "IBM z Systems Collocated Application Pricing for z/OS can improve the cost of deploying new z/OS applications." Accessed September 15, 2015. http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=gpateam&supplier=897&letternum=ENUS215-174

[4] Poneman Institute. "2014 Cost of Data Breach Study: Global Analysis." May 2014.

[5] Poneman Institute. "2014 Cost of Data Breach Study: Global Analysis." May 2014.

[6] IBM. "IBM Breaks U.S. Patent Record in 2014." Accessed August 18, 2015. http://www-03.ibm.com/press/us/en/pressrelease/45793.wss

[7] IBM. "IBM Delivers New Big Data Capabilities on IBM Cloud Marketplace." Accessed August 25, 2015. http://www-03.ibm.com/press/us/en/pressrelease/44188.wss

IMW14836-USEN-00