Geeknet

Slashdot sourceforge

free(code):

IBM

Why Your Enterprise Can't Do Without a Proven Information Governance and Optimization Solution

IT Manager's Journal
March 2012

This IT Manager's Journal takes a close look at information governance and optimization as a requisite element for success in almost any information management project – as well as delivering an efficient IT infrastructure, and shows why IBM's System z platform is an ideal hub for information governance solutions. Readers will learn how IBM System z, which is used by 95 percent of all Fortune 1000 companies to store their data, is designed to offer uptime in terms of years, not weeks or months. With the ability to absorb additional workloads without having to add or manage more servers, IBM System z and its many solutions offer proven and advanced security, availability, disaster recovery, and single-platform strengths that are unmatched by competitive offerings.

Information governance is a new(-ish) term for the solution to an age-old human issue: As reasoning and record-making creatures, we work with data all the time. But as useful and essential as data is to us, we're hugely vulnerable to its failings. Wrong information. Incomplete information. Information of uncertain source, or of uncertain meaning. Information taken out of context, or otherwise misinterpreted – these things can subvert our ability to reason and make the right decisions.

In today's business world, the problem is magnified manifold by a variety of factors: increasing the amount of structured and unstructured data we can access, its internal redundancies and deltas, its necessary (and sometimes unnecessary) duplication and spread; increasing our reliance on the metadata we wrap around information to represent its provenance and indicate its meaning; and increasing our real need to use data as inputs for decision-making, and then again, as metrics, to judge if decisions were sound.

And as our enterprises grow, our needs increase. There's no way an executive team, a board, or other top-level stakeholders can manage an enterprise of global scale, except in a fly-by-wire mode, where all they ever really see are reports. This situation is increasingly duplicated all the way through the food chain to the lowest IT worker or consumer of data on the totem pole: the one looking at the server logs all day as well as the office clerk unsure which data to use when processing a claim.

Everyone, in other words, is faced with too much data to check, but with the need to filter and act on that data effectively.

Enter information governance – a concrete methodology for discovering, combing out, labeling and schematizing, optimizing, protecting, permitting, and forbidding access to data – in short, everything you need to be sure that information and its use are being handled in sound ways, under control of policy. The basic toolkit (and of course, it's far from basic) is IBM InfoSphere Guardium and its companion applications – the 'metal' is IBM's System z mainframe product line, DB2 for z/OS and IMS databases, and a host of other data repositories and processing engines that find, store, protect, and apply information in the course of business.

In this IT Managers' Journal, we'll review four informational assets that together offer a comprehensive picture of what information governance itself means, the problems it solves, and how these best-of-breed IBM solutions make good governance achievable.

1 The whitepaper **Information Governance: Audit and Protection on the IBM System z Platform**, by Mike Ferguson of Intelligent Business Strategies, begins by offering an exhaustive outline of the appropriately deep and comprehensive scope for systems to remedy problems, impose order, and reduce risks. The first half of the paper is written in a products-free, objective style that should be perfectly accessible to executives and planners outside the IT chain. It can serve very well as the backbone for any organization's initial local assessment of specific conditions, then to help enumerate business-level and basic technical requirements for an RFQ.

Ferguson then goes on to show how IBM InfoSphere Guardium enterprise audit and protection software provides an adaptable, single-point-of-administration solution for enterprise-wide information governance – one compatible with a huge assortment of databases and with the repositories associated with typical large-scale business process management and business intelligence (BI), applications, such as SAP, PeopleSoft, and IBM Cognos. Finally, the paper details how InfoSphere Guardium can be deployed on System z mainframes, and how it goes about auditing and protecting the VSAM, DB2 for z/OS and IMS, and other data repositories commonly associated with these critical enterprise systems. Ferguson also enumerates how InfoSphere Guardium itself can interoperate with a host of other IBM information management, data development, data archiving, and software develop/test solutions. This latter section, clearly intended for IT specialists, will answer many "So how, exactly…?" questions raised by more general sections within the document.

2 Next up, the IBM Software thought leadership whitepaper **Delivering information you can trust: The benefits of quality data and IBM System z** drills into massively important questions surrounding data at the point of use, e.g. "Is this datum right or wrong?" "Is it complete?" "Do we know its provenance?" and "Do we understand its semantics – what does it really mean in the context of our business?"

Solving these problems is critical for business decision-making, process integrity, regulatory compliance (and proof of compliance) – everything, really – from the heights of five-year-planning to the most mundane ducks-in-a-row transactions like "making sure the actual inventory is actually covered by the actual insurance policy in place." All of these things depend on data that's correct, complete, contextualized, and actionable. And yet, based on surveys enumerated in the paper, many businesses are just muddling along. For example 71% of over 400 respondents from global organizations said that poor data quality was a fundamental concern, and about the same percentage of organizations said that data quality issues had emerged for them over the prior three years (survey completed late 2010).

The fix, the paper argues, is to impose an information governance protocol backed up by software tools to define, systematize and label, store, protect, authorize access to, audit, and deliver data to applications and people. While IBM InfoSphere Information Server for System z is used as the concrete and arguably best-of-breed example, most of the paper's explanations are generic. In this discussion, emphasis is given to the initial steps of data discovery, cross-referencing, schema-building, tagging and metadata provision, quality analysis, and standardization – the IT and process-governance parts of the equation – that end up providing high-trustworthiness information across the organization. Towards the

end of the whitepaper is a brief mention about the role and capabilities of System z in insuring resilient support of this enterprise information governance model.

3 Our third asset, a whitepaper entitled **Information lifecycle management on IBM System z,** from IBM's Systems and Technology Group, picks up the thread of resilience and business continuity by providing a comprehensive outline of System z environment data backup, issue-assessment analytics, and recovery tools for DB2 for z/OS and IMS – typically the enterprise-internal and WebSphere-based data repositories in many Fortune 1000 implementation scenarios. IT experts will appreciate the inside view, and the fact that highly automated, rich tools are offered on both sides for fine-grained analysis and rapid intervention following compute or database failure, making it possible, in many cases, to achieve a known state and restart clean (much faster RTO), rather than recover to an arbitrary point in time.

The paper goes on to discuss database cloning tools and facilities; optimization tools for maintaining fast throughput, verifying service levels, and identifying performance bottlenecks and errors. All this is used as background for a brief but trenchant high-level discussion of organization-wide data architecture and governance planning – though in truth, this article will be of greatest use to decision-makers with strong IT background.

4 Finally, the whitepaper **Building business by lowering costs and increasing revenue,** from IBM's Information Management Division System z Group, returns to executive-level insights with a penetrating, practical analysis of how information governance in the InfoSphere/System z mold can influence business viability. It can accomplish this both on the cost side by reducing exposures, enabling agile management of risks, avoiding bad decisions at every level, and maintaining profitable relationships with customers and with supply chains (while also reducing overall IT costs, both in terms of data administration and in terms of asset utilization efficiency), and on the revenue side by supporting organizational awareness, ready access to KPIs, and continual optimization of the relationship of information to business outcomes.

The stack of enhanced processes works together in what the whitepaper's writers identify as the "IBM Information Governance Council Maturity Model," which can be (from a very high level) summarized as a methodology and set of prescriptive guidance for managing data in project contexts, defining process and data sources, and eventually migrating forward to a state where quantitative management of process and data and ongoing optimization-to-goals can be maintained.

The paper concludes by detailing how this can be achieved, using IBM InfoSphere Guardium and related tools, and – omitted in the works discussed above – goes the extra step to describe how this information governance methodology and operational toolkit can work in the context of overarching IT initiatives, like datacenter consolidation and virtualization, which have their own obvious payoffs in cost reduction, increased efficiency, and improved alignment of IT with business goals.

Taken together, these four assets offer a complete story arc through the information governance problem set and offer concrete solutions you can implement, starting today. Once you've read through them, visit http://www.ibm.com/software/data/db2imstools/solutions/data-governance.html ***for more information.***

WHITE PAPER

INTELLIGENT BUSINESS STRATEGIES

Information Governance: Audit and Protection on the IBM System z Platform

By Mike Ferguson
Intelligent Business Strategies
November 2011

Prepared for:

Table of Contents

INFORMATION GOVERNANCE: THE RISK TO DATA

Many organisations today are introducing additional controls and accountability into management practices to improve, governance and mitigate against risk

In the last few years since the banking catastrophes of 2008, we have seen many organisations across different vertical industries introducing additional controls and accountability into their management practices to improve governance and mitigate risk. Some have done this voluntarily while others have been forced into it through legislation and regulation. A good example of the latter would be the Solvency II risk management regulations imposed on the European Union insurance market by the European parliament. These are designed to ensure that insurers can meet their obligations in a worst case scenario with regards to risks they have insured across all classes of business e.g. motor, property, casualty, professional indemnity, etc. Sarbannes-Oxley (SOX) is another example where legislation is forcing companies to tighten procedures, introduce new process controls and recording business activity to improve business practices.

A key part of improving governance, mitigating risk and remaining compliant is associated with getting control of information within the enterprise. The term given to this practice is information governance. Information governance describes the overall management and control of information throughout the entire organization and can be defined as:

"The people, processes, policies and technology used to formally manage and protect structured and unstructured data assets to guarantee commonly understood, correct, complete, trusted, secure and findable information throughout the enterprise".

Information governance has become important because people need trusted data to help them manage risk and remain compliant

Information governance has risen to the fore over the last few years because organisations have realised that their ability to manage their business, remain compliant and mitigate risk will be compromised without 'rock solid' data. People need to have confidence in the information they are using. Poorly governed information can impact on many areas including operational efficiency and decision making effectiveness. Breaches in data security may also occur or at the very least, the likelihood that these breaches can occur is increased. It is the issue that this paper is concerned with, namely, the risk to data.

Breaches in data security have been increasing steadily over the last few years with governments and private sector organisations falling victim to cyber-attacks and unauthorised user access. These kinds of incidents can seriously damage brand image and customer confidence if news of any breaches become public. This in turn, can impact on share price and bottom line performance. Customers may also fall victim to identity theft as a result. Any worthy information governance strategy needs to address this problem by ensuring that information is properly secured and protected through the implementation of an information protection program aiming to lower business risk. Information protection can be defined as:

Information protection is part of an information governance strategy aimed at preventing security breaches and data leaks

"The people, processes, policies and technology used to formally protect structured and unstructured data assets to guarantee trusted and secure data throughout the enterprise"

Information protection includes the management of information confidentiality, information integrity and information availability. It includes establishing preventative measures as well as monitoring and reporting potential problems and acting before they become major issues. All of this reduces the threat of data breaches and unauthorised changes to sensitive data. This paper looks at the information protection problem, the requirements it imposes and how one vendor, IBM, is addressing these requirements on System z platform using software that is part of its InfoSphere tool suite for information governance and enterprise information management.

THE NEED FOR ENTERPRISE WIDE DATA PROTECTION AND SECURITY

Data is becoming more fractured making it difficult to manage and protect

In many enterprises today, the data landscape is becoming increasingly complex. Data is becoming more fractured as companies continue to implement functional packaged applications (e.g. e-commerce, CRM, Finance, HR), data warehouse appliances, cloud based software-as-a-service applications, and workload optimized systems. Most organisations are also suffering from data redundancy both in terms of their structured data and also their semi-structured and unstructured data. In addition, an increasing number of external data feeds are now coming into the enterprise. The result is that there are many different databases and files spread across multiple database management systems, multiple file systems on multiple operating systems, and across multiple locations across the enterprise.

Data volumes, variety and velocity are also increasing

In addition we are now in an era where we are seeing unprecedented growth both in terms of data volumes and in the variety of data types in use by businesses. The arrival of machine generated data such as sensor data is a good example of this new growth. It has the added characteristic of velocity whereby data is generated at rapid rates. The exploding volume of social network data is also getting attention from marketers and product development personnel interested in understanding who the influencers are in the customer base and what their customers are saying about products and services.

Sensitive data could be widely distributed making it difficult to manage and protect

It is not difficult to deduce from a data landscape like this that sensitive data could be widely distributed across multiple databases and file systems. This increases the risks of security being compromised. Finding, controlling access to and protecting sensitive data content is already a challenge in this kind of environment. However without software assistance, trying to protect it can be very costly. Access privileges need to be controlled across many different applications, many different tools and many different databases on different platforms across the enterprise to prevent unauthorised changes to sensitive data and data breaches.

Data needs to be protected in development, test and production environments

In addition, amongst all of this is the IT department where development, testing and production environments are created. So, we are not only concerned about production data. Live data (some of which could be sensitive) is often used to support development activity and to conduct testing. Therefore not only do we have the problem of a complex production data landscape but we also have to take development and testing environments into account as well before new applications and databases make it into production.

Figure 1 summarises the complexity that companies are facing when it comes to information protection. Many things need to be taken into account including:

- The existing data landscape
- Different types of environments
- Different types of users
- The existing application portfolio
- End user, developer and privileged user tools
- Compliance regulations and legislation
- The need to be able to assess vulnerability to breaches in security
- Information protection prevention measures
- Information protection enforcement

In addition integration with other infrastructure software also needs to be considered (e.g. corporate LDAP directories for role based access enforcement).

Organisations needs to consider many different facets when formulating a holistic approach to information protection

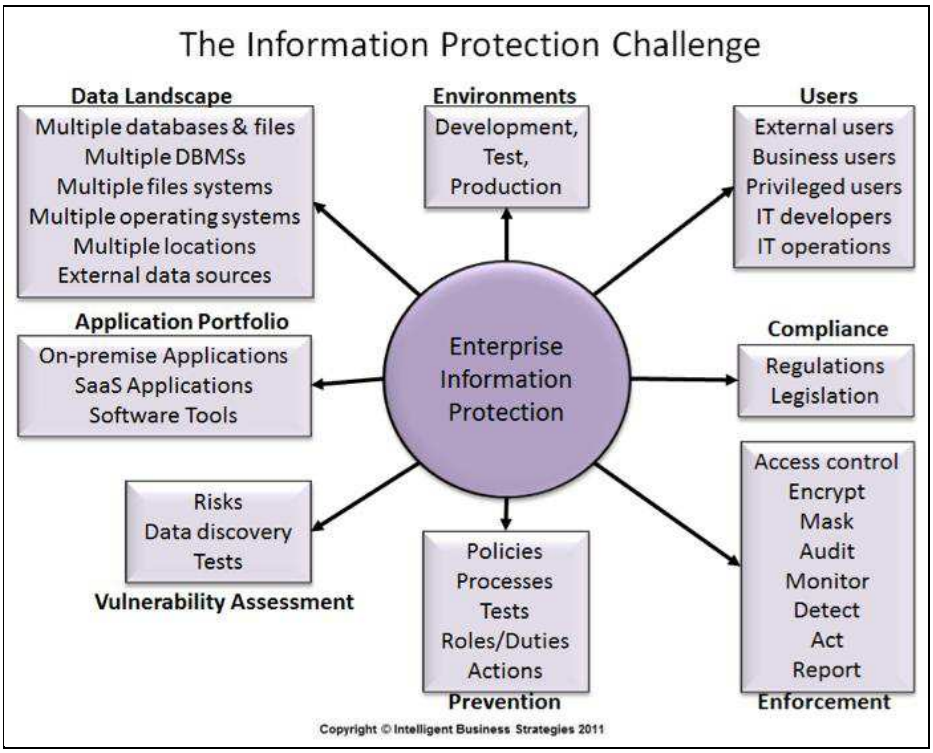

Figure 1

The vulnerability of the organisation to potential breaches in security needs to be assessed

Enterprise information protection therefore has to be holistic. It has to cover all bases. It is a risk management initiative established to avoid information risks that might breach legislation, cause non-compliance with regulations or adversely impact the organization's ability to meet its business objectives. It involves being able to locate sensitive data, assess the vulnerability of the organisation to potential breaches in security, implement prevention measures to avoid putting data at risk, monitor events that may signal a problem and respond in a timely manner to minimise the impact of these events on the business as a whole when they occur. Access control and sensitive data masking is at the heart of this, because without these there is no solid foundation on which to implement enterprise information protection.

An information protection strategy is needed

It follows therefore that companies need an information protection strategy to protect information as it flows though out the enterprise. This strategy needs to include a vision, statements on policy towards protecting information, statements on risk tolerance, identification of staff responsible for information protection, staff reporting structure for information protection related issues and protection management reports that go to authorised individuals and organisational bodies. It should also include details how the company measures success of its information protection program using key metrics indicators.

Information protection controls help to manage and prevent major risks ocurring

In addition, information protection controls are needed to control information that needs to be protected. Companies need to understand what the information risks are and what controls are in place to protect information to reduce these risks. These controls may be in the form of access approval processes, data masking and encryption processes, auditing, backup policies, retention policies, and other checks and balances. If an information protection violation occurs, then there needs to be a damage limitation process to manage losses and manage changes to procedures to avoid the same thing happening again.

It is important to have tested procedures in place to deal with disasters

Companies also need procedures in place to prepare for information ‘disasters’. Companies need to identify and rank information confidentiality, integrity and

availability disasters in order of importance, stress test each of them and put any necessary contingency plans in place to respond in a robust way if they occur.

IMPACT OF ENTERPRISE COMPLIANCE ON AUDIT AND DATA PROTECTION

Minimising the cost of compliance is of paramount importance when it comes to information protection

Compliance is another factor in Figure 1. Regulations may mandate that selected database activity is recorded for audit purposes. This includes recording updates to sensitive data, schema changes, policy changes, privilege escalation and privileged user behaviour monitoring across all systems in the enterprise. Remaining compliant may also mean real-time monitoring for suspicious behaviour (e.g. unauthorised access to sensitive data) so that this can be blocked. The issue here is the cost when enforcing information protection across the enterprise. The impact of compliance can be expensive and so minimising cost is paramount importance.

KEY PLATFORMS THAT NEED TO BE PROTECTED AND SECURED

Core transaction processing and data warehoouse databases need to be protected

When it comes to implementation, the vast majority of sensitive data resides in 'core platform' databases. This includes sensitive data in:

- Core operational transaction processing databases and files
- Enterprise data warehouses and data marts including DW appliances

IBM System z is an important platform to include in a information protection program

Transaction processing systems is a classic place where sensitive data (e.g. customer financial information) resides. Many of these systems run on mainframes making the IBM System z platform an important platform to include within the scope of an information protection program. In addition, protection of data in office documents (e.g. Microsoft Excel files and Sharepoint lists) is also needed.

NEW TRENDS – VIRTUALIZATION AND PUBLIC CLOUD

Data in virtualised environments also needs protected

In addition to the core systems mentioned above, the emergence of virtualisation software has seen the introduction of both virtualization inside the enterprise and public cloud computing outside it. Virtual servers, virtual storage and virtual networking have all been created with virtual storage bringing all disks into a seamless pool of storage to persist data. Data in the storage pool physically resides somewhere in the enterprise but it still needs to be protected. With respect to public cloud, the use of software-as-a-service applications (e.g. Salesforce.com) and adoption of services to back-up / archive corporate data to the cloud means that some corporate data may now reside outside the enterprise. Ultimately this means that the scope of enterprise information protection may have to be extended to reach beyond the corporate firewall in order to protect that data as well.

INFORMATION PROTECTION SCOPE

Information protection policy scope can initially be restricted to make dealing with the problem more manageable

Scope is also important. Data is typically created as part of an operational business process and flows throughout the enterprise in an information supply chain. As processes execute, copies and subsets of data may end up in multiple operational applications and data stores as well as data warehouses and data marts for reporting and analytical processing. The challenge for information protection is to consistently apply protection policies to data throughout the entire information supply chain no matter where that data flows to. Furthermore, protection policies must be enforced while data is in motion and while it is at rest. Information protection can be implemented incrementally by identifying the data to be protected, defining information protection policies to this data and then deciding on scope. The scope of these information protection policies can start off as being limited to specific systems, processes, organisational units or business entities (e.g. customer, employee, patient etc.). Scope can then be widened as each incremental phase of an information protection strategy is completed until all necessary data is protected.

Scope can then be widened gradually until all necessary data is protected

KEY REQUIREMENTS FOR PROTECTING DATA AND PREVENTING SECURITY BREACHES

Given everything that needs to be considered in Figure 1, the next question is "What are the requirements for protection and securing information to prevent security breaches?" and also "How do information audit and protection technologies stack up when it comes to meeting these requirements?" To answer these questions, we must define a list of information protection requirements that software products should support. These are detailed below in no particular order of preference and are grouped into categories shown in Figure 1 for easy reading.

DATA LANDSCAPE REQUIREMENTS

- It should be possible for enterprise information audit and protection software to simultaneously and continually protect and secure data held in multiple heterogeneous relational and non-relational DBMSs throughout the enterprise
- It should be possible for enterprise information audit and protection software to simultaneously and continually monitor and audit database activity across multiple heterogeneous relational and non-relational DBMSs throughout the enterprise

Data needs to be protected and secured across heterogeneous databases and file systems

- It should be possible for enterprise information audit and protection software to protect and secure data held in multiple files systems on different operating systems across the enterprise
- It should be possible for enterprise information audit and protection software to simultaneously and continually monitor and audit activity for data held in multiple file systems throughout the enterprise
- It should be possible for enterprise information audit and protection software to simultaneously and continually protect and secure data held in specific locations and continually monitor and audit activity for data held in specific locations

Sensitive data needs to be protected and secured at the lowest level of granularity

- It should be possible for enterprise information audit and protection software to protect and secure access to sensitive data down to the individual record and field levels in multiple heterogeneous relational and non-relational DBMSs throughout the enterprise

SOFTWARE ACCESS REQUIREMENTS

- It should be possible to register instances of software applications and tools on different systems and in different locations throughout the enterprise as authorized or non-authorized so as to be capable of detecting access to sensitive data from non-authorised application and tool instances
- It should be possible for enterprise information audit and protection software to be able to identify, monitor and audit all activity from all instances of applications and tools used to access and manipulate sensitive data across multiple heterogeneous relational and non-relational DBMSs and file systems throughout the enterprise

ENVIRONMENT REQUIREMENTS

- It should be possible for enterprise information audit and protection software to protect and secure data in all:
  - Development,
  - Test
  - Production

  environments across multiple heterogeneous relational and non-relational DBMSs and file systems throughout the enterprise. Since many customers still don't have information audit and protection controls in production, the production environment is the priority.

Development, test and production environments all need to be monitored, audited and protected

- It should be possible for enterprise information audit and protection software to monitor and audit database and file based activity in all
  - Development
  - Test
  - Production

  environments across multiple heterogeneous relational and non-relational DBMSs and file systems throughout the enterprise

USAGE REQUIREMENTS

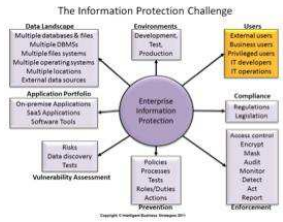

- It should be possible for enterprise information audit and protection software to be able to distinguish between:
  - External users
  - Internal business end users
  - Privileged users e.g. database administrators
  - IT developers
  - IT operations personnel

Privileged user behaviour needs to be monitored across all systems

- It should be possible to monitor and audit privileged user behavior across heterogeneous relational and non-relational DBMSs and file systems throughout the enterprise
- It should be possible for enterprise information audit and protection software to integrate with corporate user directories such as LDAP directories or Microsoft Active Directory to automatically discover active users within the enterprise and to automatically monitor new user creation
- It should be possible for enterprise information audit and protection software to integrate with relational and non-relational DBMSs to automatically monitor new user creation
- It should be possible for enterprise information audit and protection software to automatically discover privileged users declared in relational and non-relational DBMSs and operating systems throughout the enterprise
- It should be possible for enterprise information audit and protection software to automatically discover user privileges and who granted these privileges and when

VULNERABILITY ASSESSMENT REQUIREMENTS

Vulnerability assessment requirements are associated with the ability to gauge the likelihood of data risk exposure such as unauthorised access to sensitive data and to specifically pinpoint risks that need to be addressed. The requirements are as follows

- It should be possible to mark data items defined in a business glossary as 'data at risk' so as to create a 'data at risk register' within a business glossary that is visible to authorized business users to make people aware of sensitive data

Need the ability to automatically discover sensitive data to find out where it is located

- It should be possible for enterprise information audit and protection software to be capable of automatically discovering sensitive data in heterogeneous relational and non-relational DBMSs and files throughout the enterprise to determine where sensitive data is located and what policies to apply to protect and secure access to it. Automatically discovered sensitive data attributes should be mapped to common definitions in the business glossary so that it becomes possible to use the glossary to highlight sensitive data items in heterogeneous data stores right across the enterprise that qualify for vulnerability assessment testing

- It should be possible to define protection policies and enforcement mechanisms for specific business glossary data items marked as 'at risk' and have these policies enforced enterprise wide

It should be possible to assess an organisations vulnerability to security breaches

- Enterprise information audit and protection software should provide a pre-built set of vulnerability tests available to test if sensitive data is exposed to unauthorized access or if it is not masked

- Enterprise information audit and protection software should provide a pre-built set of vulnerability tests available to test exposures cause by 'loose' privileges allocation in DBMSs e.g. GRANT….WITH GRANT OPTION or GRANT…TO PUBLIC

- Enterprise information audit and protection software should provide a pre-built set of vulnerability tests to test for non-compliance of specific line item requirements defined within regulations and/or legislation

- It should be possible to extend pre-built vulnerability tests and to create additional custom built vulnerability tests that can be run to assess vulnerability of data to other risks

PREVENTION REQUIREMENTS

These requirements define the capabilities enterprise information audit and protection software need to provide to minimise exposure to data risks across DBMSs and file systems in the enterprise. The requirements are as follows:

- Enterprise information audit and protection software should provide a set of pre-defined
    - Policies
    - Roles
    - Tests
    - Templates

  to speed up implementation and enforcement of information protection across the enterprise and to comply with security and privacy regulations

Out-of-the-box pre-built templates help organisations get started quickly

- It should also be possible to define custom policies to protect and secure access to sensitive data residing in one or more heterogeneous DBMSs and file systems across the enterprise
- It should be possible to define custom policies to mask and encrypt sensitive data on one or more instances of a database, a file, or structure in one or more heterogeneous DBMSs and files systems across the enterprise
- It should be possible to group policies in any way to make it easier to administer the protection and security of at risk data. For example:
  - Common policies that can be enforced enterprise wide
  - System specific policies
  - Policies associated with a specific master data entity e.g. customer or employee data
  - Policies associated with a specific transaction and its data
  - Policies associated with specific business intelligence
  - Policies associated with a specific database or file and its data
  - Policies associated with a specific data structure
  - Policies associated with a specific location and data located there
  - Policies associated with a specific compliance regulation

Policy based access to sensitive data is paramount

- It should be possible to define policies that control access to sensitive data. These policies should be capable of being applied at multiple different levels including:
  - Access to specific data in all systems across the enterprise
  - Access to specific data across all instances of a database
  - Access to specific data in specific database or file instances
  - Access to specific data in a specific database or file structure
  - Access to specific data by specific users or user groups
  - Access to specific data by applications and software tools
  - Access to specific data at a specific location
  - Access to specific data only at specific times e.g. only within working hours
- It should be possible to define policies that govern the masking and encryption of data in one or more heterogeneous relational and non-relational DBMSs and files throughout the enterprise

It should be possible to restrict sensitive data manipulation, changes to schema and priviliges escalations

- It should be possible to define policies that restrict the ability to change the schema of any database in any DBMS instance or instances
- It should be possible to define policies that prevent manipulation of data by unauthorized transactions, software tools and users
- It should be possible to restrict the ability to change information protection policies to only authorized users
- It should be possible to restrict the powers of privileged users across one or more databases, files, DBMSs, file systems and locations
- It should be possible to separate the duties of privileged users from approvers so that privileged user activity can be formally controlled across development, test and production environments
- It should be possible to lock down production databases to prevent privileged users from creating changes to information protection policies,

and schemas and also to unlock a database and once unlocked to monitor it to make sure all changes by privileged users audited so that they can be reported

Approval workflows are needed to control privileged user behaviour

- It should be possible to define processes that prevent privileged users from escalating their own privileges or the privileges of others
- It should be possible to flag policies as needing to be enforced in real-time or on a scheduled basis
- It should be possible to define expiry dates for information protection policies and to implement policy version control

Policies need to be enforced in real-time on a continuous basis

- It should be possible to define information audit and protection 'agents' that can be deployed to monitor specific data access behavior and database activities in real-time. This includes monitoring privileged user activity, end user access to sensitive data, access to sensitive data from application transactions and tools, file open and close, etc. These software information audit and protection agents should be capable of being deployed in specific heterogeneous relational and non-relational DBMSs and file systems across all defined locations and platforms in the enterprise
- It should be possible to define policy-based actions or action sequences that can be invoked by information audit and protection software agents in real-time to uphold protection policies and neutralize threats when unauthorized or suspicious behavior occurs
- It should be possible to monitor all units of work associated with sensitive data in a database
- It should be possible to monitor access to database image copies containing sensitive data
- It should be possible to monitor changes to database metadata

A centralized auditing function is needed

- It should be possible for information protection agents to generate and emit audit entries that can be sent back to a centralized auditing function for the enterprise that automatically records all audited activity in a tamper-proof repository
- It should be possible to define policies to prevent circumvention of a DBMS by monitoring and denying access to underlying files used by a DBMS to store sensitive data

ENFORCEMENT REQUIREMENTS

Real-time monitoring of privileged and application users is needed

- It should be possible to audit all of the following in real-time
  - Schema changes including what the changes were, who made the changes, when they were made, who approved them and when
  - Policy changes including what the changes were, who made the changes, when they were made, who approved them and when
  - Access and changes to sensitive data by any user, application object, transaction unit of work, query plan or software tool
  - Access to image copies
  - All SQL statements accessing sensitive data
  - Time of access to sensitive data (e.g. outside working hours)
  - Outbound transactions from databases e.g. caused by triggers
  - Privilege escalations
  - Login failures
  - New user IDs created
  - Sharing of user IDs
- It should be possible to detect, audit and block the following in real-time
  - Access to sensitive data by an unauthorized end user

Unauthorised access and changes should be blocked

- Access to sensitive data by an unauthorized application object, transaction or query plan
- Unauthorized database activity e.g. privilege escalation
- Unauthorized policy changes
- Unauthorized schema changes
- Unauthorized opens and closes of a file
- Unauthorized outbound transactions coming from a database
- Access to sensitive data by an unauthorized software tool e.g. TOAD
- User account creation by an unauthorized user, application or tool

- It should be possible to notify nominated user(s) about security exceptions, login failures and privilege escalations in real-time via a user-defined alerting mechanism (e.g. email, SMS, dashboard alert, etc.) and escalate the alert if necessary if receipt of alert is not acknowledged

Pre-built reports and dashboards help speed up time to deployment

- It should be possible for enterprise information protection software to provide pre-built out-of-the-box reports and dashboards for monitoring and auditing purposes on unauthorized access, suspicious behaviour, privilege escalations, login-failures, schema changes, policy changes etc.
- It should be possible for enterprise information audit and protection software to provide pre-built out-of-the-box reports on regulatory compliance violations
- It should be possible to create custom built reports from the audit repository to satisfy specific questions

PERFORMANCE REQUIREMENTS

A solution needs to scale

- It should be possible for enterprise information audit and protection software to scale across all databases and file systems in the enterprise
- It should be possible for enterprise information audit and protection software to impose minimal overhead on application databases and the daily running of operational and analytical workloads by using agent software to emit data back to a centralised audit repository

PROTECTING AND SECURING DATA ON IBM SYSTEM Z

Having defined the requirements that enterprise information audit and protection software should support, this section of the paper looks at how one vendor – IBM - steps up to meeting these requirements to protect information across the enterprise. We will then focus on the IBM System z mainframe platform which typically runs some of the most mission critical transaction processing systems within the enterprise. This is an important platform because many of the databases and files housed on IBM System z may hold sensitive data.

IBM INFOSPHERE GUARDIUM

IBM InfoSphere Guardium is enterprise infprmation audit and protection software

IBM provides a suite of tools for enterprise information governance and information management under the IBM InfoSphere brand. One of the products within that product line is IBM InfoSphere Guardium. IBM InfoSphere Guardium is enterprise information audit and protection software that runs on a range of operating systems including:

- IBM AIX
- HP-UX
- Red Hat Linux (including on System z)
- SUSE Enterprise Linux (including on System z)
- Solaris — SPARC and Intel/AMD
- Tru64
- Windows 2000, 2003, 2008
- iSeries IBM i5/OS ®
- IBM z/OS

InfoSphere Guardium Data Sources

IBM InfoSphere Guardium can runs agents on multiple platforms and can audit and protect information in multiple databases and files

With respect to data sources, InfoSphere Guardium can audit and protect information on the following DBMSs and file systems across the enterprise.

DBMS Platform Supported	Versions
Oracle Database including ASO/SSL	8i, 9i, 10g (r1, r2), 11g, 11gr2
Microsoft SQL Server	2000, 2005, 2008
Microsoft SharePoint	2007, 2010
IBM DB2 ® (Linux, UNIX, Linux for System z)	9.1, 9.5, 9.7
IBM DB2 (Windows)	9.1, 9.5, 9.7
IBM DB2 pureScale ®	9.8
IBM DB2 for z/OS	8, 9, 10
IBM IMS™	10, 11, 12
IBM VSAM	1.10 (5694-A01) or later
IBM DB2 for IBM iSeries ®	V5R2, V5R3, V5R4, V6R1
IBM Informix ®	7, 9, 10, 11,11.50, 11.7
Oracle Sun MySQL and MySQL Cluster	4.1, 5.0, 5.1
SAP Sybase ASE	12, 15, 15.5
SAP Sybase IQ	12.6, 12.7, 15
IBM Netezza ® NPS	NPS 4.5, 4.6, 5.0, 6.0, 6.02
PostgreSQL	8,9
Teradata	6.X, 12, 13, 13.10
FTP	
Windows File Share	Windows 2003, 2008

It can also monitor the following packaged enterprise applications

- Oracle E-Business Suite, PeopleSoft, Siebel
- SAP
- IBM Cognos
- SAP Business Objects Web Intelligence
- JDA Supply Chain

Custom applications and 3rd party software can also integrate with InfoSphere Guardium

For DBMSs and applications not directly supported out-of-the-box, InfoSphere Guardium also supports a new Universal Feed capability whereby customers and 3rd party software vendors can integrate niche DBMSs and applications with InfoSphere Guardium via an API to allow InfoSphere Guardium to monitor, audit and protect the data in these systems in real time. This allows organisations to move towards full enterprise wide security and compliance.

InfoSphere Guardium Deployment Options

InfoSphere Guardium ships as a hardware or software appliance

To minimise deployment costs, InfoSphere Guardium can be delivered either as a pre-configured hardware appliance or as a preconfigured software appliance. The hardware appliance can optionally be deployed into a IBM System zEnterprise BladeCenter® Extension (zBX) frame. This is discussed in more detail in the "Guardium Integration with Other System z Infrastructure" section later in this paper. The software appliance can run either on user supplied hardware or in a user supplied virtualized environment running under the control or a hypervisor[1].

InfoSphere Guardium Architecture

InfoSphere Guardium uses software agents to monitor and audit multiple data sources

In addition InfoSphere Guardium is architected to scale across multiple DBMSs and data centres throughout the enterprise. This is achieved via the use of software agent technology known as software taps (S-TAPs). InfoSphere Guardium S-TAPs are lightweight software agents (or probes) that are installed on various databases and can be deployed to run on multiple operating systems in a heterogeneous environment. Each S-TAP captures all local activities by privileged users and in addition also monitors all access, local and remote, to databases and supported files by users, applications and tools. These S-TAP agents send the captured data back to a central audit repository via so-called data collectors to avoid impacting on the performance of the databases and files they are monitoring, as well as to ensure the security of the audit trail.

InfoSphere Guardium Grid allows the information protection service to scale smoothly as more sources are added

Most enterprises typically start using InfoSphere Guardium by monitoring a few major data sources and then broaden the use of audit and protection software to gradually bring more data under management until all core information is under control. As more data sources are added, more InfoSphere Guardium S-TAPS can be gradually deployed which means that more data needs to be collected. In order to cater for smooth growth in data collection, InfoSphere Guardium can assign S-TAPs to virtual IP addresses. This allows it to decouple S-TAPs from data collectors and to use its InfoSphere Guardium Grid technology to load balance data collection across a grid of InfoSphere Guardium data collectors. By allowing the number of data collectors and S-TAPs toggled separately it means that InfoSphere Guardium can scale easily without any need to re-plan configurations as the use of information audit and protection grows. This introduces elasticity into the configuration and also helps to minimise the cost of administration.

InfoSphere Guardium Components

InfoSphere Guardium comes with a number of components including

- Database activity monitoring on the aforementioned data sources
- A central manager and aggregator
- A tamper-proof audit trail repository
- Vulnerability assessment capability

[1] EMC VMware ESX is currently supported

- A database and sensitive data finder
- Data level access control
- Entitlement reports
- Workflow automation
- Agents for monitoring databases and other repositories
- Software that automates integration with LDAP, databases and other IT infrastructure
- Accelerators for regulatory compliance e.g. PCI, SOX

Workflow automation is supported to help automate audit report generation, distribution to key stakeholders, electronic sign-off and escalations. Workflow processes are completely user customizable allowing specific audit items to be individually routed and tracked through sign-off.

Vulnerability assessment is supported because configurations, privileges, etc. are constantly changing, and can introduce easy paths to gain unauthorized access to sensitive data, Organizations can't afford highly skilled DBAs to constantly check database configurations to identify configuration errors. Vulnerability assessment makes it possible to identify these errors without the need for DBAs thereby enabling configurations to be constantly hardened to eliminate security gaps.

INTRODUCING IBM INFOSPHERE GUARDIUM ON SYSTEM Z

InfoSphere Guardium on System z can protect IMS, DB2 and VSAM data

In many large enterprises the mainframe is a platform that runs mission critical transaction processing applications. The databases and files used by these applications typically hold considerable amounts of detailed and often sensitive data including customer financial information. In order to protect these data, IBM has extended the reach of InfoSphere Guardium to the IBM System z platform. The following System z data stores can be protected:

- DB2
- IMS
- VSAM
- 3rd party databases via the Universal Feed

InfoSphere Guardium uses a variety of methods to collect events in order to gather the necessary information to protect these IBM System z data stores. This includes memory inspection, use of base product instrumentation, and use of information that is sent to the IBM System z systems management facility (SMF). Each InfoSphere Guardium S-TAP may use a combination of techniques depending on the variety of event information that needs to be collected. The selection of the technique is made based on ensuring there is a comprehensive record of events, balanced against minimizing overhead, while providing separation of duties as necessary.

Auditing and Protecting DB2 Data on System z

InfoSphere Guardium for IBM System z can also protect information in DB2 for z/OS environments.

Organisations can use InfoSphere Guardium auto-discovery and information classification capability to identify where confidential data is stored in DB2 z/OS databases. Having done this, they can then use customizable classification groups to automate enforcement of security policies that apply to particular classes of sensitive objects. This ensures that sensitive data can only be utilised by authorized users.

DB2 on z/OS can be monitored to look out for changes to data structures, privileges, and sensitive data

In addition, InfoSphere Guardium can continuously monitor DB2 for z/OS database operations in real time alerting when unauthorised activity is detected. This includes monitoring of:

- DB2 for z/OS security exceptions such as SQL errors

- Changes to DB2 structures using CREATE, DROP and ALTER SQL statements
- SELECT SQL statements and cursor operations such as OPEN and READ
- Data manipulation via SQL INSERT, UPDATE and DELETE
- DB2 for z/OS GRANT and REVOKE statements that control accounts, roles and permissions

It is also possible to conduct vulnerability assessment tests to highlight risk exposures such as privileges that have been GRANTed TO PUBLIC or GRANTed WITH GRANT OPTION. This allows organisations to reduce risks and to separate duties. For example a DB2 system administrator may have SYSADM authority allowing him or her access to sensitive data. Identifying this during vulnerability testing allows organisations the opportunity to change this to SYSCNTL authority without data access.

By continuously tracking all DB2 for z/OS database actions, InfoSphere Guardium can be used to proactively identify unauthorized or suspicious database activity. In addition, malicious or unapproved activity by DBAs, developers and outsourced personnel can be detected without the need to rely on native logs, triggers or other DB2 DBMS-resident mechanisms.

Pre-configured reports and are also available to view database activities in detail. This includes login failures, escalation of privileges, schema changes, access during off-hours or from unauthorized applications and access to sensitive tables.

Auditing and Protecting IMS Data on System z

IMS is a hierarchical database management system that is optimized for performance and mission critical applications. It has been available on System z for over forty years and is still widely used in a large percentage of Fortune 1000 companies around the world. IMS is mainly used to support applications with very high volume transaction workloads. Just like any other database management system, many organisations are now seeking to improve the protection of the information they hold in IMS databases. In some cases this information protection requirement has been brought about because of new regulations or legislation (e.g., SOX, Solvency II), however in the vast majority of cases it is the implementation of data governance that often brings information protection under the spotlight. Some geographies around the world, also put a higher priority on information protection than others. This is particularly true in Europe.

With respect to information protection requirements in an IMS environment, there is a need to support all variants of IMS including batch IMS. In addition protection of IMS data and metadata is needed. This includes the need to:

- Monitor privilege user activities
- Automate identification of policy violations
- Create a granular audit trail of all database activity
- Assure separation of duties

Information protection policies can be applied at the database and segment level in IMS

InfoSphere Guardium for IBM System z steps up to this challenge by providing a new S-TAP software agent for IMS. One S-TAP agent is needed for each IMS instance with a single audit server being shared across S-TAP for IMS agents.

Each agent can protect information by implementing policies and capturing IMS events in real time at both the IMS database level and the segment level. IMS events captured include all

- DB READs,
- INSERTs,

- UPDATEs
- and DELETEs

associated with both IMS online regions and IMS batch jobs.

All events are fed back to an InfoSphere Guardium for IMS collector in order to detect policy violations, provide support for a secure audit trail and to reporting on policy violations. Compliance workflow is also supported.

Event auditing and monitoring can be controlled at multiple levels of granularity by applying filters. This includes being able to filter by

- All databases
- One IMS database
- All segments
- Some segments

InfoSphere Guardium agents stream captured events off System z to avoid interfering with high volume workloads

S-TAP for IMS can then stream the captured events off the z/OS server to avoid interfering with high volume transaction processing workload performance.

In addition InfoSphere Guardium for IMS can also protect IMS metadata held outside the IMS runtime. It does this by auditing and watching out for changes to IMS objects. This includes monitoring and auditing who is touching image copies, logs or recon datasets.

Auditing and Protecting VSAM Data on System z

In addition to the IBM IMS and DB2 database management systems on System z, it is also possible to protect information in VSAM files using the IBM InfoSphere Guardium S-TAP for VSAM.

Monitoring VSAM on z/OS prevents privileged users from circumventing the DB2 DBMS

VSAM files are an important data source to protect in a System z because they are often used by mission critical mainframe applications such as ATM and core banking applications for example. They are also used as the underlying file system of relational DBMSs such as IBM DB2. Therefore even though IBM DB2 for z/OS supports a SYSADM privilege without data access, it is important to protect against circumvention of the DBMS by privileged users who may chose to go behind the back of the DBMS in a 'back door' attempt to access and manipulate information in VSAM datasets.

It is not surprising therefore that many organisations have highlighted the need to protect VSAM files on System z. To cater for this requirement, InfoSphere Guardium provides an S-TAP agent for VSAM. There is one S-TAP agent for VSAM per system

All types of VSAM datasets can be protected by InfoSphere Guardium including:

All types of VSAM dataset can be protected

- ESDS
- KSDS
- RRDS
- VRRDS
- LDS file types

Everything and everyone that touches these datasets can be monitored in real time with all events being captured at the dataset level. This includes real-time monitoring and auditing of all VSAM data set OPENs, OPEN for UPDATEs, DELETEs, RENAMEs, CREATEs, ALTERs, RACF ALTERs, CONTROLs, UPDATEs and READs.

As with the other InfoSphere Guardium S-TAP agents on IBM System z, they filter events away from the VSAM datasets through an InfoSphere Guardium for z/OS

collector to a centralized tamper proof audit trail to provide organisations with a secure audit trail of VSAM activity. VSAM activity and policy violations can then be reported from this audit trail. In addition it is also possible through the S-TAP for VSAM, to detect policy violations in real-time and take responsive action if necessary.

AUDITING AND PROTECTING 3RD PARTY DATABASE DATA ON SYSTEM Z

With respect to protecting information in third party databases on IBM System z, InfoSphere Guardium's provides the aforementioned Universal Feed capability. This allows customers and third party software vendors to integrate third party DBMSs and applications with InfoSphere Guardium via an API. This allows organisations to extend the reach of InfoSphere Guardium to other database management systems (DBMSs) running on the System z platform so that database activity and access to sensitive data in these third party DBMSs can be monitored, audited and protected in real time.

GUARDIUM INTEGRATION WITH OTHER INFOSPHERE TOOLS ON SYSTEM Z

InfoSphere Guardium is only one of the tools in IBM's InfoSphere product line. The entire product line is listed below together with a description of each tool's capability. InfoSphere Guardium products are highlighted.

IBM has a suite of tools for Information Governance

IBM Information Management Products	Description
IBM InfoSphere Foundation Tools	
• IBM InfoSphere Business Glossary	Define, manage and control common data names and data definitions for all master data and transaction data that needs to be governed
• IBM InfoSphere Data Architect	Data modelling
• IBM InfoSphere Discovery	Discovery of disparate data within and across source systems that needs to be governed, cleaned, integrated and protected so that it is made fit for business use
• IBM InfoSphere Information Analyzer	Data quality profiling to determine the state of data that needs to be governed and to monitor and make people accountable for data quality to maintain business confidence in it
• IBM InfoSphere FastTrack	Capture design specification mappings and generate data integration services to integrate and clean data
• IBM InfoSphere Metadata Workbench	Monitor data flows - metadata lineage and audit
IBM InfoSphere Information Server	
• IBM InfoSphere Blueprint Director	Used to build templates for data warehousing, MDM, data migration, data synchronisation etc. from data quality and data integration services created in underlying IBM InfoSphere Information Server tools as part of an information governance program
• IBM InfoSphere Quality Stage	Data cleansing and matching
• IBM InfoSphere DataStage	Data integration for consolidation
• IBM InfoSphere Federation Server	On-demand data federation to integrate data from multiple underlying data sources
• IBM InfoSphere Services Director	Information service publication for use in managing and governing data
IBM InfoSphere Guardium	Real-time database activity monitoring and database vulnerability assessment
IBM InfoSphere Optim Data Masking Solution	Data masking for privacy in non-production environments
IBM InfoSphere Optim Test Data Management Solution	Subset data to right-size and speed deployment of testing environments. When combined with masking making them secure
IBM InfoSphere Optim Data Growth	Database archiving

InfoSphere Guardium is part of a suite of tools

Solution	
InfoSphere Guardium	Real-time database activity monitoring Monitor privileged users e.g. DBAs Monitor enterprise application users for fraud Enforce database change control Prevent database leaks
InfoSphere Guardium Data Redaction	Remove sensitive data from unstructured environments (documents, graphics…)

A key question to ask with respect to the tools listed, is how does InfoSphere Guardium integrate with other tools in the InfoSphere product line? One of the tools in the above table that is of particularly interest is InfoSphere Discovery. This product is used to discover data and data relationships both within and across disparate systems irrespective of data location or schema.

IBM InfoSphere Discovery can be used in a number of areas including in the identification or data for archiving, test data management and application consolidation projects. It enables users to gain an understanding of data content, relationships and transformations across multiple heterogeneous sources and also look for patterns as part of understanding the data landscape.

Organisations can use InfoSphere Discovery to locate and identify sensitive data that InfoSphere Guardium can then protect

In the context of information protection, having access to a data discovery tool can be extremely useful for obvious reasons. If we don't know where data is located, we cannot fully protect it. InfoSphere Guardium's built-in discovery capabilities address this problem by automatically identifying where sensitive information is located so that InfoSphere Guardium can:

- Monitor and audit application and user access to discovered sensitive data in real time
- Monitor and audit application and user maintenance of discovered sensitive data in real time
- Monitor and audit privileged user activity in real time to check for suspicious behaviour around discovered sensitive data
- Apply and enforce policies to help protect discovered sensitive data in order to prevent data leaks and compliance violations.

Note that for information protection to cover all bases, data discovery needs to include the identification of all redundant copies of sensitive data irrespective of whether these data are in development, test or production systems.

In addition it is also possible to make use of the bidirectional interface between InfoSphere Guardium and InfoSphere Discovery[2] to provide additional capabilities. This allows customers to select the discovery approach that meets their need. Integration with InfoSphere Business Glossary would also be useful to determine what data items in the glossary have been flagged as sensitive and in need of protection. IBM has implemented integration between InfoSphere Guardium and InfoSphere Discovery via a new bidirectional interface that enables the exchange of metadata regarding sensitive data. This integration makes it possible for organisations to make use of InfoSphere Discovery to automatically analyze complex data landscapes to identify sensitive data. Metadata can then be exchanged with InfoSphere Guardium so that discovered sensitive data can be protected. It addition, because the interface between the two products is bi-directional, it also allows information about sensitive objects in InfoSphere Guardium to be leveraged by InfoSphere Discovery so that all relationships to that sensitive data can be identified. In this way, information protection becomes more

[2] InfoSphere Discovery can be used across multiple types of initiatives including archiving, security, etc.

comprehensive because information protection policies can be applied to sensitive data as well as all data related to that sensitive data.

GUARDIUM INTEGRATION WITH OTHER SYSTEM Z INFRASTRUCTURE

InfoSphere Guardium can run across System z LPARs, a SysPlex configuration and plug into the zEnterprise zBX facility

InfoSphere Guardium can be used to protect information across System z logical partitions (LPARs) and SysPlex environments. In the case of LPARs an S-TAP agent can be deployed on each LPAR. In addition, InfoSphere Guardium supports authentication via LDAP.

IBM is further integrating InfoSphere Guardium with other System z infrastructure by making use of the IBM System zEnterprise BladeCenter® Extension (zBX) facility. Given that InfoSphere Guardium can ship on a hardware appliance, this appliance can fit into a zEnterprise zBX frame as a blade extender and communicate with the zEnterprise platform over the high performance private network. This makes it possible to manage InfoSphere Guardium using the zEnterprise Unified Resource Manager just like any other heterogeneous resource and allows InfoSphere Guardium workloads to run on optimized hardware.

In contrast to logging and the systems management facility used by custom developed auditing applications, the use of S-TAPS in InfoSphere Guardium reduces IBM System z resource consumption by offloading events to collectors while also allowing separation of duties. In addition by supporting a wide variety of repositories on IBM System z as well as on Linux, Unix and Windows (LUW) platforms a single enterprise-wide view is provided. InfoSphere Guardium also supports a broad security and compliance capabilities (e.g. vulnerability assessment, sensitive data finder, compliance workflow automation, entitlement reports) and can be operated by non-superusers such as security personnel or a database security officer. Overall therefore, in this broader enterprise wide context, InfoSphere Guardium is a more comprehensive security and compliance solution than custom developed auditing applications.

OTHER SYSTEM Z SUPPORT FOR INFORMATION PROTECTION

System z also provides further support for information protection

In addition to the protection offered by InfoSphere Guardium itself, the IBM System zEnterprise also offers further support for information protection both in terms of hardware and software. zEnterprise hardware support is shown in the table below:

System z Information Protection Hardware	Description
Central processor assist for cryptographic functions (CPACF)	Accelerates the encrypting and decrypting of SSL transactions and VPN-encrypted data transfers
Configurable Crypto Express3 (CEX3)	Optional feature suited to applications requiring high-speed, security-sensitive, RSA acceleration, cryptographic operations for data encryption and digital signing
Trusted Key Entry (TKE) workstation and smart card reader	Optional workstation that provides security-rich local and remote key management

In addition to InfoSphere Guardium real-time information audit and protection of IMS, DB2 and VSAM, IBM also offers the following software to help on System z.

- IBM RACF - to managing role based access to data and services
- IBM Optim Data Privacy and Data Retention
- DB2 for z/OS SYSCNTL authority without data access
- DB2 for z/OS EXPLAIN query plan facility without execute privilege or ability to access data
- Data encryption for IMS and DB2 databases using System z hardware

CONCLUSION

Organisations need to understand their data landscape and locate sensitive data so they can apply protection policies

Information protection is currently implemented in a very fractured and inconsistent way across systems in most enterprises. It has not been helped by the fact that the technology components available in the market to implement end-to-end information protection have also been somewhat stand-alone and lacking in end-to-end integration. Organisations therefore want to find a way to eliminate siloed approaches in order to improve security and reduce costs. This together with the increasing threat of internet fraud plus mounting pressures caused by stricter compliance regulations and risk management has led many companies to start looking for integrated end-to-end solutions that automate security and compliance activities. The challenge is to improve the protection of information no matter where it resides or flows to in the enterprise and to do so at an affordable cost.

Key platforms like IBM System z should be given priority

To get started with information protection, organizations need to understand their data sources and application portfolio as well as any development, test and production environments that exist. In addition they need to understand the different types of users within the enterprise. Having done this, sensitive data needs to be located so that organisations can work out what policies need to be put in place to prevent security breaches and what needs to be monitored and audited in real-time to guard against suspicious behaviour. Protecting key platforms like IBM System z mainframes which typically runs mission critical transaction processing applications should be given priority since these systems see significant changes to data and can be exposed to internet access which potentially makes them a target for cyber-attacks and fraud.

InfoSphere Guardium is capable of auditing and protecting information on System z and across the entire enterprise

In terms of implementation, organisations looking for software capable of auditing and protecting information on IBM System z as well as other heterogeneous data sources would have to seriously consider IBM InfoSphere Guardium as a candidate for the task at hand. IBM InfoSphere Guardium supports automated discovery of sensitive data, vulnerability testing, policy management and real-time continuous monitoring and auditing of activity on System z IMS, DB2 and VSAM data. It also provides these same capabilities on non-System z platforms for data residing in IBM DB2, Oracle, Microsoft, Teradata and SAP Sybase, SQL Server, Sharepoint, IBM Netezza, MySQL, Postrges and Informix databases across a broad range of operating systems. In addition it can control, audit and monitor access to sensitive data from applications and tools as well as privileged user behaviour including privilege, policy and schema changes. It also supports policy based actions to block[3] changes and unauthorised access. Also its agent based architecture is designed to scale with minimal interference to the data sources it monitors. Add to this a centralised manager and aggregator, a tamper-proof repository and entitlement reports and it is not surprising that the product is a serious contender to minimise data risks and compliance violations on System z and across the enterprise.

[3] Available on certain platforms – refer to vendor support page

About Intelligent Business Strategies

Intelligent Business Strategies is a research and consulting company whose goal is to help companies understand and exploit new developments in business intelligence, analytical processing and enterprise business integration. Together, these technologies help an organisation become an intelligent business.

Author

Mike Ferguson is Managing Director of Intelligent Business Strategies Limited. As an analyst and consultant he specialises in business intelligence and enterprise business integration. With over 30 years of IT experience, Mike has consulted for dozens of companies on business intelligence strategy, data governance, master data management, technology selection, enterprise architecture, and SOA. He has spoken at events all over the world and written numerous articles. Mike is a resident expert on the Business Intelligence Network, providing articles, blogs and his insights on the industry. Formerly he was a principal and co-founder of Codd and Date Europe Limited – the inventors of the Relational Model, a Chief Architect at Teradata on the Teradata DBMS and European Managing Director of Database Associates, an independent analyst organisation. He teaches popular master classes in New Technologies for Business Intelligence and Data Warehousing, Enterprise Data Governance, Master Data Management, and Enterprise Business Integration.

Water Lane, Wilmslow
Cheshire, SK9 5BG
England
Telephone: (+44)1625 520700
Internet URL: www.intelligentbusiness.biz
E-mail: info@intelligentbusiness.biz

Information Governance: Audit and Protection on the IBM System z Platform

IBM Software
Thought Leadership White Paper

September 2010

Delivering information you can trust

The benefits of quality data and IBM System z

IBM

Executive summary

Would you decide to have an operation or life-saving medical procedure based on impartial or incomplete data? Imagine if you did—and later found out that you were incorrectly diagnosed and the treatment was not necessary. You would have put your life at risk because of bad data.

In the business world, data quality problems may not be a matter of life and death, but they still pose a serious threat. More and more companies and government organizations are learning that poor-quality data is damaging their ability to support business processes, comply with regulations and make accurate decisions. In a recent survey of 407 various-sized organizations from around the world, the top two concerns for business and IT executives were increasing confidence in information for decision making and improving data quality. Almost 70 percent of surveyed organizations had experienced data quality issues during the past three years (see Figure 1).[1]

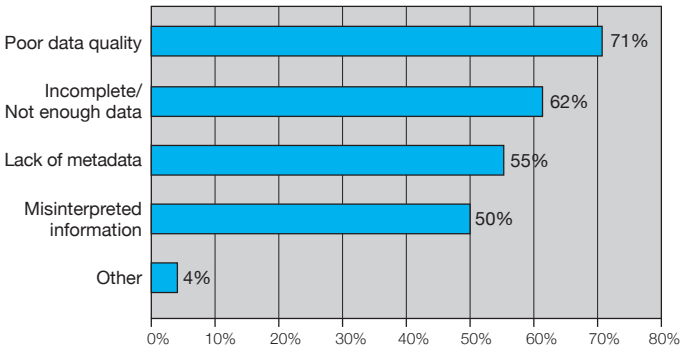

Note: Based on 282 respondents.

Source: "Information Governance as a Holistic Approach to Managing and Leveraging Information." BeyeNetwork Custom Research Report prepared for IBM Corporation by Judith R. Davis.

Figure 1: Survey respondents experienced a variety of data quality problems during the past three years (Note: Respondents could check more than one option).

Even though most business leaders understand the need for high-quality data, they are often not sure how to achieve it. And yet, it has become clear that an investment in the infrastructure must be made to ensure measurably acceptable data quality.

This is where IBM® InfoSphere® Information Server for IBM System z® can make a big difference. InfoSphere Information Server is the foundation of many successful data quality initiatives, helping organizations derive more value from the complex, heterogeneous information spread across their systems. Meanwhile, System z provides a resilient, reliable, high-performance platform for mission-critical data; it's estimated that 95 percent of Fortune 1000 companies store business data on System z.[2]

Data quality can determine business success or failure

Having a clear understanding of customers, partners and suppliers can mean the difference between growing a business and failing to compete. Critical initiatives for information governance, compliance and master data integration simply will not succeed unless the quality of the data in systems is clearly understood and actively managed.

Put another way, bad data is like a virus. A virus is small, yet if left undetected or misdiagnosed, it can spread and become more aggressive, eventually even crippling its host. Imagine that one of your core information systems contained some unchecked or inaccurate data. The data might start out in a part of your strategic master data management (MDM) system used by departmental systems, data warehouses, business intelligence (BI) systems, subsidiaries, trading partners or regulatory reporting staff to make essential business decisions. That bad data will also be used by workers across the information supply chain—and they will consume, process and then spread the "infected" information to others. Along the way, it will skew metrics, reduce report accuracy and ultimately affect business decisions.

The effects of poor data quality include failed business processes, lower productivity and wasted materials. Lost, inaccurate or incomplete information also generates higher costs and extra work, such as hunting down information or additional reconciliation.

The IBM System z platform is designed for co-locating business-critical applications, processes, transactions, BI systems and data warehouses, and can deliver performance, security and operational advantages. However, if left unaddressed, poor quality data will undermine these benefits.

Information governance

Most organizations have not yet evolved their processes, policies and infrastructure to be able to help ensure high data quality levels. As a result, organizations are beginning to adopt information governance, a quality-control discipline for adding new rigor to the process of defining common terminology and managing, using, improving and protecting information.

Effective information governance can enhance the quality, availability and integrity of a company's data by fostering cross-organizational collaboration and structured policy-making. It balances factional silos with organizational interest, directly impacting the four factors that an organization cares about most: increasing revenue, lowering costs, reducing risks and increasing data confidence.

Information governance enables an organization to monitor its information supply chain as an end-to-end system, helping to ensure that information is consistently defined and well understood, of high quality, managed throughout its life cycle and protected and secured wherever it lies. With information governance, organizations achieve many goals, from improving decision making to simplifying and strengthening regulatory compliance.

A forum for information governance

Now more than ever, the challenge to protect and manage data has become a universal concern for organizations. To help better understand the emerging space, IBM created a leadership forum in November 2004 for chief data officers and security, risk, compliance and privacy officers concerned with information governance issues.

Since then, the IBM Information Governance Council has steadily grown to comprise nearly 55 leading companies, universities and IBM Business Partners, including large financial institutions, telecommunications organizations, retailers and even government agencies. The Council designed an information governance framework to help businesses understand the supporting and core disciplines and enablers of information governance. It also produced a maturity model to help assess information governance within an organization. To broaden involvement in the Council, IBM launched an online community to encourage organizations to participate, using crowd-sourcing technology to further enhance the maturity model and information governance as a whole.

For more information on the IBM Information Governance Council, please visit: www.infogovcommunity.com

IBM InfoSphere Information Server for System z

The success of an information governance program and supporting data quality initiatives hinges upon a robust data integration technology infrastructure. Enter InfoSphere Information Server for System z, an IBM software platform that provides breakthrough productivity and performance for understanding, cleansing, transforming and moving information consistently and securely throughout the enterprise.

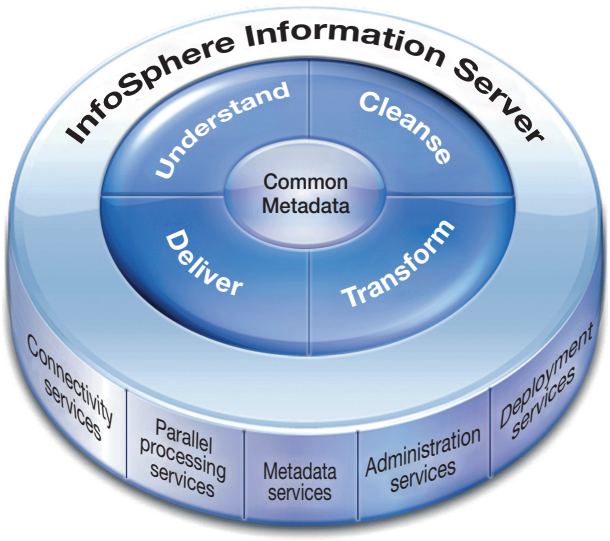

Figure 2: InfoSphere Information Server is built on a foundation of parallel processing and other services.

InfoSphere Information Server for System z helps you leverage information across all of its sources. The software delivers the functions required to integrate, validate, enrich and deliver trusted information for your key business initiatives. It enables you to:

- *Define* a common business language for your information
- *Understand* all sources of information within the business, analyzing its usage, quality and relationships
- *Cleanse and standardize* information to assure its quality and consistency
- *Transform* information to provide enriched and tailored information
- *Federate and deliver* information to make it transparently accessible to people, processes and applications

Those functions are based on a parallel processing infrastructure that provides leverage and automation across the platform (see Figure 2). InfoSphere Information Server for System z also provides connectivity to almost any data or content source, and the ability to deliver information through a variety of mechanisms.

A project lifecycle approach

InfoSphere Information Server for System z employs a unified metadata management foundation that enables the seamless sharing of knowledge throughout a project life cycle—and a detailed understanding of what information means, where it comes from and how it relates to information in other systems.

The common metadata infrastructure facilitates a shared understanding across business and technical domains, helping to reduce the time between specification and build. The infrastructure also provides a persistent record of understanding that can dramatically reduce downstream project delivery times and help improve overall insight and confidence in information.

All functions of InfoSphere Information Server for System z share a metadata model that spans design and operational metadata, making it easy for different roles and functions to collaborate seamlessly. The platform provides comprehensive reporting on data movement, data lineage, business meaning and the impact of changes and dependencies across InfoSphere Information Server for System z modules and third-party tools.

InfoSphere Information Server for System z provides:

- Access to a broad range of information sources
- Extensive integration functionality, including federation; extract, transform, load (ETL); in-line transformation; replication and event publishing
- Flexible integration capabilities, including support for Service Oriented Architectures (SOAs), event-driven processing, scheduled batch processing and standard APIs like SQL and Java™

Proven technologies help ensure data quality

InfoSphere Information Server for System z achieves new levels of information integration, speed and flexibility by providing data quality and transformation capabilities to boost trust in information, automate partitioning and process pipelining for scalability and rapidly deploy services to enhance value. The following components provide additional functionality.

IBM InfoSphere Discovery

Before you can implement an information governance program or information-centric project, you must know what data you have, where it is located and how it relates between systems. For most organizations, the data discovery process is manual, requiring months of human involvement to discover business objects, sensitive data, cross-source data relationships and transformation logic. The result is a time-consuming, error-prone process that slows time to value, establishes doubt about the accuracy of the data within the new system and creates the possibility that the new system will never become operational.

IBM InfoSphere Discovery provides a full range of capabilities to automate the data discovery process. It addresses single-source profiling, cross-source data overlap analysis, matching key discovery, prototyping and testing for data consolidation and automated transformation discovery. InfoSphere Discovery also uses heuristics and sophisticated algorithms that automate analysis to help companies realize 10 times more time and cost savings compared to performing the same tasks manually using a profiling solution.[3]

InfoSphere Discovery capabilities include:

Data profiling: InfoSphere Discovery provides advanced data profiling with results that are "fit for purpose." This includes column analysis, automated primary-foreign key discovery and simultaneous cross-source column overlap analysis of multiple data sources. These sources can be as simple as text files on a PC or as complex as virtual storage access method (VSAM) on System z—or both at the same time.

Unified Schema Builder: Unified Schema Builder takes the output of overlap analysis and uses it as input into a process for helping a data analyst determine the rules by which data will be consolidated for data migration, MDM or a data warehouse, to name a few examples. The Unified Schema Builder component delivers automation software with an embedded workflow to help you complete your consolidation project on time and within budget.

Transformation Analyzer: This component is designed to automate discovery of complex cross-source transformations and business rules by analyzing data values and patterns across two data sources. Transformation Analyzer is used when you know that two data sources are related, but you also know that the relationship can't be described by simple overlaps in data values and requires figuring out how data is transformed between the two sources. Data migration, application retirement, data warehousing and MDM almost always require the mapping and discovery of complex transformation logic between two or more data sources. Transformation Analyzer helps accelerate this process by automating much of the analysis involved and replacing tedious manual work.

The InfoSphere Discovery analysis process establishes an understanding of your data sources and how they relate to each other, generating actionable output that can be immediately consumed by a wide range of information projects, including archiving, test data management, data privacy, data integration, MDM and data consolidation.

IBM InfoSphere Information Analyzer

IBM InfoSphere Information Analyzer, a product module of InfoSphere Information Server, delivers data profiling and rules analysis functions within the context of a complete information integration platform, enabling more accessibility and

consistency throughout the enterprise. Active metadata across InfoSphere Information Server simplifies the collection and management of metadata across the entire integration spectrum. Within the InfoSphere Information Analyzer module, profiling results are stored in the common metadata repository. By using the data profiling capabilities of InfoSphere Information Analyzer in the early phases of your data integration projects, you can:

- Expedite delivery of data-driven projects
- Utilize business-driven Rules Analysis with reusable construction and application across multiple data sources to offer quick time to value
- Help minimize costs and resources of critical data integration projects
- Eliminate the risk and impact of proliferating incorrect and inaccurate data
- Help ensure the timely delivery of trusted information

IBM InfoSphere Business Glossary

Difficulties in understanding and interpreting data, determining what data is important and then managing that information creates roadblocks as business and technical users attempt to collaborate for effective information integration. The problem of business definition inconsistency across enterprise environments is often attributed to the absence of an enterprise-wide data dictionary and stewardship program.

InfoSphere Business Glossary for System z helps you create, manage and share an enterprise-wide controlled vocabulary that acts as the common language between business and IT. This is a critical step in better aligning technology with business goals. In addition to a controlled vocabulary, the InfoSphere Business Glossary hierarchy and classification systems provide additional business context.

Actively connected to InfoSphere Information Server metadata services, InfoSphere Business Glossary enables data stewards to link business terms to technical artifacts shared between IBM InfoSphere Data Architect, InfoSphere Information Server or a third-party data integration solution. The result is a common set of semantic tags used by data modelers, data analysts, business analysts, governance stewards, data architects, developers and end users. To help ensure high quality and tight security, only authorized data stewards can use the administrative functions within InfoSphere Business Glossary to create and manage the glossary.

The solution also serves as a history of records to help ensure compliance with regulatory rules, such as the Sarbanes-Oxley Act and Basel II. Business terminology is always subject to change: What defines a “high-value customer” today may be different tomorrow as business requirements evolve. Being able to see the history of what changed, why it changed and who changed it is as important as the change itself. Such a history is critical to data governance protocols, as it increases the trust and understanding of your information.

InfoSphere Business Glossary has a web-based interface designed to enable data stewards to administer the contents of the common glossary. Plus, InfoSphere Business Glossary Packs for different industry verticals help you accelerate the implementation and deployment of your business glossary. Based on knowledge gained during work with more than 400 IBM clients and 10 years of experience in key industries such as banking, financial markets, retail, telecommunications, healthcare and insurance, those packs allow you to quickly deploy, promote and adopt InfoSphere Business Glossary—and therefore get a fast return on your investment.

InfoSphere QualityStage

IBM InfoSphere QualityStage™, part of the InfoSphere Information Server suite, enables enterprises to create and maintain an accurate view of master data entities, such as customers, vendors, locations and products. InfoSphere QualityStage may be deployed in transactional, operational, or analytic applications, and in batch and real-time environments.

InfoSphere QualityStage enables a comprehensive process to manage and maintain data quality. Its core capabilities include:

- **Investigation:** Understand the nature and extent of data anomalies and enable more effective cleansing and matching
- **Standardization:** Create a standardized view of customer, partner or product data. This capability also enables global address cleansing, validation and certification (for significant postal discounts in select localities) and geolocation
- **Probabilistic matching:** Provides an industry-leading matching engine to help ensure the best match results possible; built on a platform enabled for high connectivity and scalability
- **Survivorship:** Helps ensure the optimum consolidation, householding or linked view of record information; enables consolidated and accurate view of customers, partners, products and more

InfoSphere Metadata Workbench

InfoSphere Information Server is designed to be a complete platform for integrating and enriching information across disparate source systems. By leveraging an active and shared metadata repository layer, InfoSphere Information Server can support a full range of integration activities and user roles with collaboration and reuse principles. These artifacts include technical metadata about the various sources of information, business metadata that describes the business meaning and usage of information and operational metadata that describes what happens within the integration process.

IBM InfoSphere Metadata Workbench provides a powerful metadata management interface that supports not only InfoSphere Information Server metadata but also other key metadata that play critical roles in data integration processes. A centralized and holistic view across the entire landscape of data integration processes, with visibility into data transformations that operate inside and outside of InfoSphere Information Server, arms businesses with critical information that can lead to better decisions.

InfoSphere Metadata Workbench highlights include:

- Web-based navigation of key information assets through an interactive and powerful interface provides an easy way for users to access critical information.
- Visual cross-tool and cross-platform data lineage enables an understanding of the complete information lineage, including where data came from and what happened to it as it moved across data integration processes, with extended visibility into enterprise data flows outside of InfoSphere Information Server.
- Visual cross-tool impact analysis allows complete understanding of the impact of a change before the change is made, even when the impact extends beyond a single tool.
- Reporting on information assets, through simple and advanced search with save, repeat and publish capabilities, helps business and IT users to quickly understand complex environments.
- Automated linkages to InfoSphere Information Server metadata services help organizations reduce their overall IT costs and accelerate productivity.
- Collaboration and shared metadata with InfoSphere Business Glossary promotes data stewardship, business and IT alignment and better understanding of information assets.
- Various access levels for different user types provide the flexibility to customize requirement management and enforce information security.

System z: Business-critical resilience for information governance

Given the far-reaching effects of an information governance initiative, it's critical that the host platform be utterly reliable and always available. Consider the effect of data not being available for a period of time. A report, "Business resilience: Ensuring continuity in a volatile environment," by the Economist Intelligence Unit in 2007, stated that according to the U.S. National Archives and Records Administration, 25 percent of the companies that experienced an IT outage of

two to six days went bankrupt immediately.[4] A disruption in business continuity equates to significant risk with direct costs (loss of immediate business) and indirect costs (long-term damage and credibility to the brand). The damage can extend well beyond the financial realm into key areas, such as customer loyalty, market competitiveness and regulatory compliance.

While some solutions offer weeks or months of uptime, the System z platform offers uptime in terms of years. It provides disaster recovery configurations and is designed to deliver 99.999 percent application availability when implementing IBM Parallel Sysplex® technology to create dynamically balanced clusters with near-linear scalability. This configuration helps avoid the downside of planned downtime, equipment failure or the complete loss of a data center. It also employs some of the most advanced security technologies in the industry—helping organizations meet rigid regulatory requirements that include encryption solutions, access control management, and extensive auditing features.

As a result, System z is the hub for many organizations' information governance solutions; the world's top banks, insurers and retailers to rely on IBM to help deliver continuity and to secure sensitive business transactions.

Make your information work harder for you

InfoSphere Information Server for System z is a fully integrated software platform that profiles, cleanses, transforms and delivers information from mainframe and distributed data sources alike to drive greater insight for the business without added IBM z/OS® operational costs. It can help you derive more value from the complex, heterogeneous information spread across your systems—and support your information governance initiative.

With breakthrough productivity and performance for cleansing, transforming and moving this information consistently and securely throughout your enterprise, InfoSphere Information Server for System z lets you access and use information in new ways to drive innovation, help increase operational efficiency and help lower risk.

For more information

To learn more about information quality and the IBM System z platform as part of your information governance strategy, please visit:

- **ibm.com**/software/data/integration/info_server_system_z
- **ibm.com**/software/data/db2imstools/solutions/data-governance.html

IBM Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
September 2010

[1] Davis, Judith R. "Information Governance as a Holistic Approach to Managing and Leveraging Information." BeyeNetwork Custom Research Report prepared for IBM Corporation. www.beyeresearch.com/study/14243

[2] Moutsos, Kim. "IMS at 40: Stronger than Ever," IBM Database, October 2008. www.dbmag.intelligententerprise.com/story/showArticle.jhtml?articleID=211300235

[3] Time and cost savings based on reports from IBM client engagements.

[4] The Economist Intelligence Unit. "Business resilience: Ensuring continuity in a volatile environment," February 2007. http://graphics.eiu.com/files/ad_pdfs/eiu_Bus_Resilience_wp.pdf

Please Recycle

IMW14326-USEN-00

IBM Systems and Technology
Thought Leadership White Paper

January 2011

Information lifecycle management on IBM System z

Contents

Executive summary

We all know the importance of good maintenance. Whether it's your car, house, electronics, sports equipment—if you manage and maintain them correctly, you will extend their life and their value. By the same token, if you let them fall into disrepair, you will lose money through their degradation and devaluation. Worse still, they could fail unexpectedly, putting you and others at great risk. So it is with an organization's data.

Data is the lifeblood of organizations, where decisions are only as good as the information on which they are based. Data must therefore be protected, and its quality optimized, throughout its life cycle. This paper specifically examines requirements for the following activities during the information life cycle in mainframe computing environments:

- Data backup and recovery
- Data replication
- Data optimization
- Data maintenance
- Schema management
- Performance optimization

In addition, this paper describes IBM solutions for information lifecycle management on the IBM System z® platform for mainframe computing.

Basics of data storage and access

Data flows through the information supply chain, where it is constantly being created, shared, read, updated, archived and deleted. It touches many users, applications and business processes across a wide variety of devices.

In enterprise environments, IBM DB2® for z/OS® and IBM Information Management System (IMS™) are premier data servers and transaction managers. It is estimated that 95 percent of Fortune 1000 companies store some of their business data on the IBM System z platform because of its high availability, performance characteristics, scalability, transaction integrity and unsurpassed security capabilities.[1]

These two data servers often form the underlying foundation for many other services. On one side is DB2 for z/OS, the repository for IBM InfoSphere™ Warehouse and InfoSphere Master Data Management Server. It is the database for IBM WebSphere® Application Server, WebSphere Process Server, WebSphere Enterprise Service Bus, WebSphere Message Broker and many more applications that run on the IBM System z platform. On the other side, IMS is the repository for mission-critical operational data and the transaction manager for critical online applications. As the cornerstone for so many other products, System z's performance, availability, integrity and ability to scale to meet unforeseen workloads have led many customers to store more data—and to collocate processes, business applications, business intelligence and business analytics—on the System z platform.

The importance of business continuity

Losing data to an IT systems outage, or being unable to recover it after a loss, can have dire consequences. According to the U.S. National Archives and Records Administration, 25 percent of companies that experienced an IT outage of two to six days went bankrupt immediately.[2]

Even small amounts of downtime can be inconvenient and costly. If applications aren't consistently available, business suffers. The damage can extend well beyond the financial realm into key areas such as customer loyalty, market competitiveness and regulatory compliance. High on the list of critical business requirements today is the ability to keep applications up and running in the event of planned or unplanned system disruptions.

While some servers may offer weeks or even months of uptime, the System z platform offers years of uptime. It also provides disaster recovery configurations and is designed to deliver 99.999 percent application availability where Parallel Sysplex® is implemented, minimizing planned downtime, equipment failure or the complete loss of a data center. Additionally, the strength of System z security is one of the many reasons why the world's top banks and retailers rely on the IBM mainframe to help secure sensitive business transactions.

Data backup and recovery: Preventing loss

Backup and recovery is one of the most complex areas of database management. Having the right resources to perform a recovery is critical. Unfortunately, in many cases, this requirement is not addressed until after data is already lost.

Database backup and recovery comes in many forms, ranging from recovering from a dropped object to bouncing back from a major disaster, and everything in between. Recoveries done manually can be error-prone, time-consuming and resource-intensive, and database administrators face many recovery-related challenges. For example:

- Can a transaction be reversed or does the entire database have to be recovered?
- Is it possible to determine which objects have been impacted?
- Are the necessary resources available to recover to a point in time?
- What preparations are in place for a disaster? Can the subsystem be recovered? How much data can the organization tolerate losing?

IBM data backup and recovery tools

Some organizations have used database system-level backup (SLB) methodologies for many years as an efficient and effective way to backup database systems. Created by storage organizations using full volume dumps and storage-based fast replication facilities, these backups are shipped offsite, where they are used to provide the foundation for traditional DB2 or IBM IMS disaster recovery procedures. They tend not to be used by DBAs to perform local site recovery or disaster recovery operations because without supporting automation, it can be difficult to coordinate the data restoration from volume backups with database recovery processes. This complexity, coupled with the familiarity of using DBMS and host-based copy methods, has guided DB2 and IMS DBAs to instead use traditional image copy backups for local site recovery and disaster recovery purposes.

DB2 data backup and recovery tools

DB2 Recovery Expert for z/OS is a storage-aware backup and recovery tool that automates and simplifies backup and recovery processes that are based on SLB methods. It integrates FlashCopy® facilities for copying data and allows FlashCopy facilities to be exposed to DBAs in a transparent manner. DB2 system-level backups performed using FlashCopy offer the advantages of saving time, host-CPU resources and I/O resources.

DB2 Recovery Expert enables the recovery of database objects without the need for a full database recovery. It's a simple, self-managing tool that maintains high availability with minimal system disruption. DB2 Recovery Expert also recommends options for selecting the fastest, least costly method of recovery when time is of the essence. It helps with disaster recovery by building assets that can be moved to a disaster site for rebuilding DB2.

DB2 Log Analysis Tool for z/OS helps to minimize DB2 recoveries by determining what data has changed and identifying points of recovery. It makes it easy to distinguish between changes made to data by end users and those resulting from referential integrity constraints. It also lets organizations identify database updates that have been rolled back instead of committed—and even report on data changes to objects that have been dropped and subsequently recovered. This versatility removes the guesswork and manual labor associated with reversing improper changes to enterprise data.

IMS data backup and recovery tools

IMS DEDB Fast Recovery for z/OS is a fast alternative to emergency restart failure recovery that reduces the time needed to recover data entry databases (DEDBs) after an IMS failure. It corrects online log data sets (OLDS) by invalidating logging for transactions that did not reach the synch point, and recovers DEDBs that were active when IMS failed.

IMS Recovery Expert for z/OS determines the recovery assets to use and the recovery steps to run, and then manages the recovery process for an entire IMS system or application. This transforms disaster recovery into a disaster restart process, meeting or exceeding recovery time objectives (RTOs) and simplifying and automating backup, recovery and disaster recovery processes. It also provides the ability to automate the process of backing up entire IMS subsystems by volume.

IMS Recovery Solution Pack for z/OS provides database backup capability through image copy, provides database recovery processing for all types of recoveries, and integrates all processes into a single step. It also rebuilds index data sets for recovery or builds new indexes.

Data replication: Protecting resources

While it might be ideal to have just one copy of a piece of data to keep track of in an organization, multiple copies are in fact often necessary for production, test and development purposes. There are a number of business drivers and technology issues that necessitate high-volume, robust and secure information replication. Some of the most often cited include:

- Populating a data warehouse or data mart while maintaining data consistency among disparate systems.
- Populating a central repository from remote sites or vice-versa.
- Scheduling and maintaining consistency of data in multiple locations by providing a single version from a central site.
- Replicating information from a central production environment to a mirror image or subset used for analytical or query processes.
- Applying filters to replicate pertinent data to provide an accurate subset for business processes.
- Applying real-time or event-driven rules to ensure synchronization across data stores and locations.

IBM cloning tools

Cloning database management systems allows production data to be used for testing, reporting, data warehouse loading, database utility processing, or other production offload tasks. Offloading these types of activities to a cloned database copy reduces production I/O contention and allows processing activities on a static copy of the data. Storage-aware database cloning utilities such as IBM DB2 Cloning Tool for z/OS and IBM IMS Cloning Tool for z/OS simplify and automate the cloning process to reduce cloning time and administration costs.

IBM DB2 Cloning Tool for z/OS makes it easy to quickly clone a DB2 subsystem or a DB2 table space. The tool can automate the cloning process to provide usable DB2 clones within minutes, thus boosting efficiency and freeing DBAs from performing time-consuming multistep tasks. It has the ability to clone DB2 table spaces and automatically translate the object IDs, simplifying and automating the process of refreshing data. It also includes fast copy technology, which quickly copies DB2 data sets within a subsystem or to a different subsystem, and version-enabled journals, which make it possible to use data from older versions of DB2 Cloning Tool. Using clones with current data can help provide faster resolution if users encounter problems with production data. DB2 Cloning Tool can reduce the time required to create a clone from hours or days to just minutes.

IBM IMS Cloning Tool for z/OS, like DB2 Cloning Tool, uses FlashCopy to copy data quickly, speed up the cloning process and reduce administration costs. Data is copied by using volume or data set-based FlashCopy facilities to reduce copy time and save host CPU and I/O resources. Data can be copied while the source database is running or stopped.

Storage-aware data management tools provide facilities to link and coordinate application and data management organizations with business continuity and storage administrators, as illustrated in Figure 1. Storage-aware database management tools on z/OS systems integrate storage-based, fast-replication facilities with DB2 and IMS systems to provide fast and non-disruptive DB2 backup and cloning solutions. Storage-aware database management tools improve DB2 backup, recovery, disaster recovery and cloning solutions by using IBM FlashCopy facilities to copy data, which saves time, host CPU resources and I/O resources.

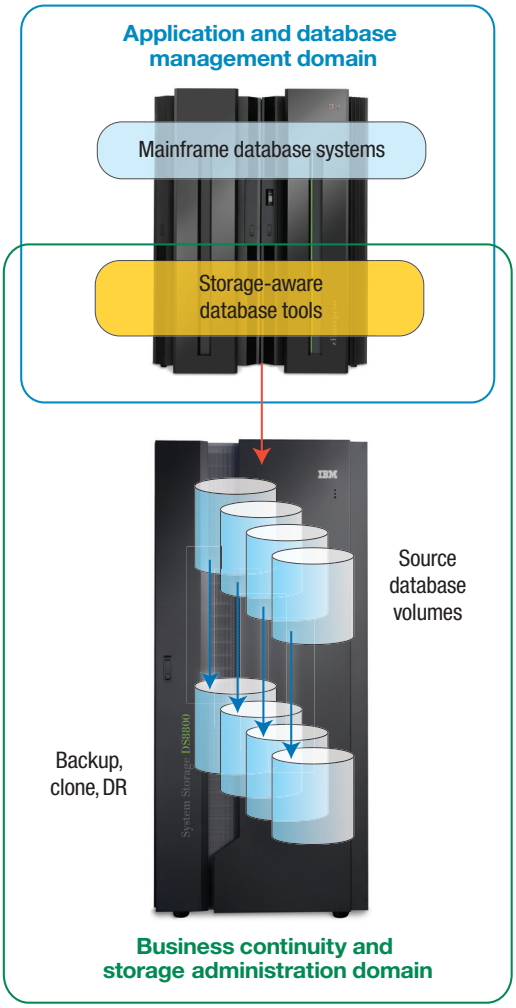

Figure 1: Storage-aware management tools can be used to integrate application and database administration with business continuity and storage administration.

Data optimization: Improving productivity and profitability

Improving productivity and profitability is a top business priority whose impact extends to the data center. DBAs continually work to maximize database performance to support service level agreements (SLAs), drive higher transaction rates and ultimately increase profitability. As they strive for top performance, DBAs are constantly looking for new ways to take a more proactive role in identifying potential performance problems before they adversely affect systems. They must be able to:

- Check the performance health of the DB2 subsystems.
- Determine why online transactions are failing.
- Drive better-performing SQL to allow for faster throughput.
- Identify the most costly SQL statements to reduce costs.
- Set alerts to know when performance thresholds are reached or exceeded.
- Make it easier for users to tune SQL statements.
- Tune IMS systems more efficiently.
- Identify a process for analyzing and isolating IMS transactions and performance bottlenecks.
- Verify service levels and predict trends that will improve the productivity of IMS specialists and/or system programmers.

IBM data growth solutions

One of the greatest impediments to successful data optimization is data overgrowth. In fact, overgrown databases can impair the performance of business-critical ERP, CRM and custom applications. InfoSphere Optim™ data growth solutions, including **Optim Data Growth Solution for z/OS**, solve the data growth problem at the source through intelligent archive management of enterprise application data. They reduce the size of production databases, providing universal access to the data while improving application performance and cutting hardware and software costs. InfoSphere Optim makes it possible to archive

historical transaction records, storing them securely and cost-effectively while maintaining universal access to them. By reducing the amount of data to sift through, InfoSphere Optim helps speed reporting and completion of mission-critical business processes.

InfoSphere Optim data growth solutions help organizations by:

- Reducing hardware, storage and maintenance costs by archiving historical data from applications and systems, freeing valuable resources.
- Maintaining optimal application performance levels by controlling ongoing database growth.
- Archiving, managing and retaining application data according to business and compliance policies.
- Responding quickly to audit and discovery requests with universal access to archived information.
- Managing and controlling archiving centrally across applications, databases, operating systems and platforms.

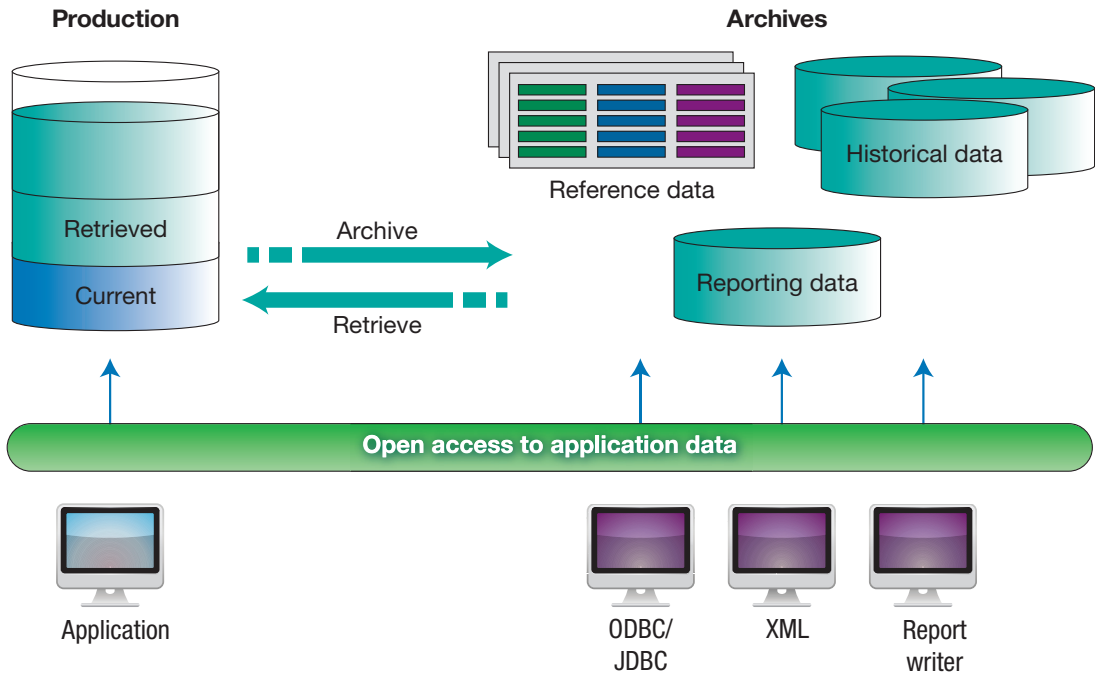

Figure 2: IBM InfoSphere Optim Data Growth for z/OS provides intelligent, application-aware archiving of data.

Optim Data Growth for z/OS can provide significant ROI.[3] In short, Optim Data Growth for z/OS improves the performance of production environments by intelligently archiving infrequently used data to other storage media, thereby enabling organizations to efficiently retain and access data in compliance with regulatory requirements. The solution also facilitates application retirement, while enabling the data to be restored at a later date.

Data maintenance: Managing data complexity

Accommodating rapid information growth and managing the increasing complexity of structured and unstructured data are key challenges for today's IT professionals. Administrators need utilities that can help control data growth and complexity throughout the information management life cycle while sustaining high availability and high performance.

IBM data maintenance tools

IBM works continuously to produce DB2 utilities and tools in response to emerging IT challenges, addressing changing application needs, supporting new DB2 versions and accommodating new data types.

IBM DB2 Utilities Suite for z/OS is a comprehensive tool set for managing DB2 data maintenance tasks. The DB2 Utilities Suite can be used to quickly and easily build utility jobs while helping to ensure the highest degree of availability, performance and data integrity.

IBM DB2 Sort for z/OS is designed to provide the highest degree of performance with significant reductions in elapsed time and CPU during utility sort processing. DB2 Utilities Enhancement Tool enhances the manageability of the maintenance window and utility processing. When combined with DB2 Sort and the DB2 Utilities Suite, organizations can eliminate the need for duplicate utilities from various vendors, helping to reduce use of system resources and total cost of ownership.

IBM DB2 Automation Tool for z/OS gives DBAs more control over DB2 utility maintenance. By allowing users to set up conditions stored in easy-to-use profiles, they can control when, how or if a utility is run. They are able to set policy rules that help establish standards at various levels, from individual objects to enterprise-wide. Senior-level DBAs can also set utility jobs for lesser experienced personnel. By streamlining utility maintenance, organizations can increase availability because of shorter batch windows and can save on system resources such as CPU and DASD.

Schema management: Responding quickly to change

Today's business environment has increased the complexity and rate of change of DB2 objects throughout the information management life cycle. The ability to respond quickly to a changing environment is constantly challenged by the explosion of data growth combined with a decline in experienced work staff.

IBM schema management tools

The **IBM DB2 Administration Tool**, when combined with **IBM DB2 Object Comparison Tool**, enables users to effectively and safely manage changes and migrations of DB2 objects and schema throughout the application life cycle, while having the least possible impact on availability and data integrity.

Performance optimization: Reducing costs

IT budgets are shrinking, and organizations must find ways to extract more value from existing assets. Understanding where and how it's possible to improve performance to meet ever-stringent service-level demands from both management and customers can save on costly and sometimes unnecessary storage and processor costs.

IBM performance optimization solutions

IBM provides capabilities for the System z platform to analyze and determine where performance bottlenecks exist and where they might be likely to occur.

DB2 performance optimization solutions

IBM Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS evaluates the efficiency of, and optimizes performance from, DB2 databases in a z/OS environment. It provides extended insight into application response time, helps monitor, analyze and tune the performance of IBM DB2 for z/OS and IBM DB2 applications, and provides robust views of performance to help improve productivity. Predefined rules of thumb make it quick and easy to identify performance bottlenecks. It also combines batch-reporting capabilities with real-time monitoring and historical tracking functions, and supports an enterprise-wide integrated systems management strategy activated by the IBM Tivoli OMEGAMON XE family. It stores performance data and analysis tools in a performance warehouse.

IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS helps resolve critical performance issues by monitoring, analyzing and optimizing the performance of DB2 Universal Database™ and DB2 on z/OS applications online in real time and in batch reports. It acts as a single engine for performance data collection logic, monitors individual data-sharing members or entire data-sharing groups, and watches applications running in a parallel query environment, even if the parallel tasks are executed on different processors. It keeps a history of recent DB2 performance metrics available for online analysis and prevention, and it enables online analysis of database objects like disks, tables, table spaces and other elements to tune performance.

DB2 Query Monitor for z/OS enables efficient customization and tuning of SQL workload and DB2 objects, arming staff with the ability to spot trouble before it can cause a significant waste in DB2 resources. It can be used to view and configure monitoring across the enterprise from a single console and to identify users and locations that are issuing problem SQL. It offers extensive choices for determining what monitoring information to gather and when to collect it and allows user-defined settings that trigger warnings when query activity crosses defined thresholds. To make managing data assets more cost-effective, it has

GUI-based reporting, viewing and configuration capabilities for access to consolidated data and events for DB2 subsystems across multiple z/OS images. It also allows users to research DB2 commands, display host variables, display SQL Communications Area and find SQL error patterns. Autonomic functionality makes it possible to execute user-configurable responses to a wide variety of events.

InfoSphere Optim Query Workload Tuner for DB2 for z/OS empowers DBAs to more efficiently manage performance by proactively optimizing the performance of SQL queries and query workloads. Built on an open-source Eclipse-based environment, Optim Query Workload Tuner makes it easy to access candidate queries and define workloads from a number of common sources, including:

- The DB2 catalog (e.g., SQL from packages or stored procedures).
- The dynamic statement cache.
- InfoSphere Optim Development Studio.
- A text file.
- DB2 Query Monitor.
- Tivoli OMEGAMON XE for DB2.

With Optim Query Workload Tuner, DBAs can efficiently capture, format and analyze SQL statements and workloads, creating graphic visualizations of query plans and costs. The solution removes much of the time and effort required for query analysis by formatting queries for readability, annotating them with statistics, visualizing the access plan and more.

DB2 SQL Performance Analyzer for z/OS helps DB2 users improve DB2 application design to achieve maximum productivity, and tests different "what if" scenarios quickly and economically to determine the performance of various database design and production volumes. It can also access any report for any input source (DBRM, plan, package, etc.). Summary reports enable access to a specific report directly from the summary report, and efficient navigation enables quick access to information that's needed most.

IMS performance optimization solutions

IMS Buffer Pool Analyzer for z/OS improves database performance by reviewing the buffer pool environment, projecting the impact of buffer pool changes and recommending changes to the number of buffers. It also provides information on the efficiency of OSAM cached buffers, models user changes to the buffer pool configuration to see their impact before they are implemented, identifies wasted storage and determines whether adding or removing buffers will improve performance, and performs "what if" scenario analysis, visualizing the impact of change in the buffer environment.

IMS Network Compression Facility for z/OS can help reduce both end-user response time and line costs by compressing data streams to end users of 3270 terminals. It is a practical, affordable tool that uses 3270 commands to compress unnecessary or redundant data—without requiring any application changes or alterations to users' screens. It also includes an ISPF interface for monitoring the performance of NCF and IMS message traffic.

IMS Performance Solution Pack for z/OS provides an affordable, complete portfolio of IBM database performance management tools for fast, easy analysis of IMB transactions:

- IMS Connect Extensions to improve availability, reliability and performance of IMS Connect
- IMS Performance Analyzer to provide information on IMS system performance
- IMS Problem Investigator to help determine the cause of problems and trace the flow of events end to end

These components collectively help improve productivity for problem analysts, increase the efficiency of IMS application performance, improve IMS resource utilization and increase system availability.

Transaction Analysis Workbench for z/OS enables analysis of transaction performance and behavioral problems and simplifies problem analysis. It extends the scope of traditional analysis techniques to make it easier to identify problems. The workbench saves a history of each problem session, attaches logs and other historical data to problem sessions, analyzes the data files associated with a problem, and provides performance interactive analysis to track a transaction and identify significant events in its life cycle.

Tivoli OMEGAMON XE for IMS on z/OS helps optimize the performance and availability of IMS systems, reducing the potential for delays or outages by reporting on a number of critical IMS attributes. It provides a single point of control over IMS in Parallel Sysplex environments and offers rich monitoring of IMS Connect TCP/IP transaction requests. Dynamic Workspace Linking (DWL) enables easy navigation between Tivoli Enterprise Portal workspaces and enhanced integration with OMEGAMON XE for DB2.

Summary

Managing the health of an organization's data can prolong the life of the systems and processes that consume and use it. Failing to manage it can, in the worst-case scenario, result in a failed business. IBM System z provides organizations with information lifecycle management capabilities for effective data backup and recovery, replication and optimization, to help minimize business risk, increase productivity and profitability, and reduce costs.

For more information

To learn more about information lifecycle management on System z, please visit: **ibm.com**/software/data/db2imstools/solutions/data-governance.html

Additionally, financing solutions from IBM Global Financing can enable effective cash management, protection from technology obsolescence, improved total cost of ownership and return on investment. Also, our Global Asset Recovery Services help address environmental concerns with new, more energy-efficient solutions. For more information on IBM Global Financing, visit: **ibm.com**/financing

[1] Moutsos, Kim, "IMS at 40 – Stronger then ever," IBM Database Magazine, November 2008. **ibm.com**/developerworks/data/library/dmmag/DBMag_Issue408_IMSat40/index.html

[2] "Business resilience: Ensuring continuity in a volatile environment," Economist Intelligence Unit, 2007. Sponsored by ACE, IBM and KPMG. **ibm.com**/services/us/bcrs/pdf/wp_business-resilience.pdf

[3] "The Total Economic Impact™ of IBM Optim Integrated Data Management Solutions: Multicompany Study," Forrester Consulting, October 2009. Prepared for IBM. public.dhe.ibm.com/software/data/sw-library/data-management/optim/papers/forrester_optim_total_economic_impact.pdf

IBM Systems and Technology Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
January 2011

ZSW03177-USEN-00

IBM Systems and Technology
Thought Leadership White Paper

April 2011

Building business by lowering costs and increasing revenue

Information governance and IBM System z provide key tools for the enterprise

Introduction

Whether in an economic downtown or a market boom, organizations seek to achieve at least two objectives: capture a larger "wallet share" from their customers and reduce their operational costs.

Increasing wallet share makes sense, as studies have consistently shown that it is far simpler and cheaper—by a factor of six to nine times—to increase sales to an existing customer base than recruit new customers.[1] The revenue challenge then becomes understanding who the customers are, determining what products they already own and identifying their propensity to buy additional products. Such insight into how best to up-sell and cross-sell customers can position a company not only to retain customers but also to outsmart competitors.

Reducing operational costs applies to every organization from small businesses, charities and government agencies to large financial institutions, retail stores, telecommunication companies and others. In today's business environment, key techniques for achieving cost savings include:

- Lowering the risk of financial penalties due to non-compliance or fines.
- Reducing duplication of work and processes due to departmental silos.
- Avoiding unnecessary IT upgrades by optimizing existing investments.
- Reducing customer churn.

A closer look at increasing revenue

An organization can address its needs to increase revenue and reduce costs in many ways—including the use of technology. A software solution such as IBM InfoSphere™ Master Data Management Server for System z®, for example, can play a significant role in helping the organization understand its customers' identities and behaviors, as well as the relationship the organization has with those customers. But the success of such a solution—indeed, the success of any initiative intended to gain insight from customer and business information—depends on the quality of the data behind it.

The relationship between customer- and product-focused operational systems and the master data management system is therefore symbiotic. While each operational system requires access to the master data (known as the "source of truth"), this master data also needs to be maintained with an understanding of changes that occur within the operational systems. Every touch point with a customer, regardless of the delivery channel (including face to face, telephone, web or broadcast commercials), becomes an opportunity to influence that customer's next move. And that opportunity potentially can result in increased revenue through up-selling and cross-selling.

The process, however, is dependent on two critical success factors: data quality and semantic consistency.

Data quality

A clear understanding of customers, partners and suppliers can mean the difference between business success or failure. Critical initiatives for master data management and data warehouse projects simply will not succeed unless the quality of the data in systems is clearly understood, managed and kept accurate.

Central to this understanding is the concept that bad data operates like a virus. While small in itself, when left undetected or misdiagnosed it can aggressively spread to the point where it can cripple its host. Unchecked or inaccurate data may begin, for example, in a small departmental system but then quickly spread to a strategic master data management system, data warehouses and business intelligence systems used by other departments.

Throughout the organization, bad data will skew metrics, reduce report accuracy and ultimately affect business decisions. It can lead to lost opportunities, lack of consumer confidence, failed business processes, lower productivity, wasted materials and massive costs.

Consider these data quality issues gathered from IBM customers engagements:

- Inaccurate or incomplete data is a leading cause of failure in business intelligence and customer resource management projects.
- 25 percent of end-user time is spent clarifying bad data.
- 83 percent of data integration projects either overrun or fail because people underestimate the complexity and inconsistency of data sources.

To combat these risks, IBM System z is an excellent platform for collocating business-critical applications, processes, transactions, business intelligence and data warehouses. System z delivers broad performance, security and operational benefits. However, if left unaddressed, poor-quality data will undermine the benefits of the System z platform.

Semantic consistency, classification and data architecture

Another key challenge lies in determining a common understanding of what data means and how it is represented. Architecting data in a consistent manner across and beyond the enterprise in collaboration with trading partners can provide a common understanding of the data for both business and IT users. This understanding can be achieved using a business glossary that enables application end users, architects and developers to understand what the information contained in a field of a form or text of an email really means—as well as who in the organization is responsible for defining and managing that meaning. By architecting data at the enterprise level instead of at the individual project level, a successful business glossary will help reveal the structure and meaning of data, increasing its value as an enterprise asset.

A closer look at reducing operational costs

Good maintenance is important. Whether it's a car, house, electronics or sports equipment—managing and maintaining your assets correctly can extend their life and their value. By the same token, allowing them to fall into disrepair can bring degradation, reduce investment value and increase cost to repair. Worse still, they could fail unexpectedly.

In the same way, maintaining data is important as it flows through the entire information supply chain, where it is constantly being created, shared, read, updated, archived and deleted. The value of data is immense—as can be the cost of an inability to recover from an accidental or deliberate system outage or data loss.

Today, organizations are storing data for longer periods than ever before. This is due mainly to the requirements of government and industry regulatory standards. Both the Sarbanes-Oxley and BASEL II requirements, for example, not only specify the periods that data must be kept, they require that in the early years of storage, data remain readily accessible. As a result, the cost of data storage may increase and the performance of business applications and service level agreements may be impacted.

As another example: while global communications and the Internet have removed geographical boundaries, enabled 24x7 business operations and allowed companies to reach a wider customer set, security breaches and their costs have also increased.

To manage and reduce the cost of business operations, it is therefore critical to carefully manage issues related to the collection, storage, protection and use of business data. Three challenges emerge as particularly important: the management of data growth, the optimization and performance of data, and the establishment of appropriate data security and privacy measures.

The role of governance in managing costs and increasing revenue

Recognizing the need to raise the strategic importance of enterprise-level information, IBM and some 40 other industry leaders formed the IBM Information Governance Council in 2004. The aim was to encourage new thinking about how data could be leveraged beyond the individual project level to reach across the enterprise.

The role of data management and governance on building revenues and reducing costs is underscored in a 2010 report based on a survey designed by IBM, which revealed that poor data quality, data retention and archiving, managing data growth and performance degradation, unauthorized access to data, and privacy issues are key areas of concern.[2]

IBM provides solutions to help meet these concerns. Through the use of vulnerability assessments, ROI calculators, data quality assessments and business value assessments, organizations are able understand that it is far cheaper to put into place proactive preventive measures or automated solutions than it is to attempt to fix problems manually or otherwise tackle them after an incident has occurred.

One such asset has been the IBM Information Governance Council's Maturity Model, which assesses an organization's achievements in information governance in 11 categories across five phases—from organizational behavior and stewardship to how the organization audits, logs and reports information.

IBM Information Governance Council Maturity Model

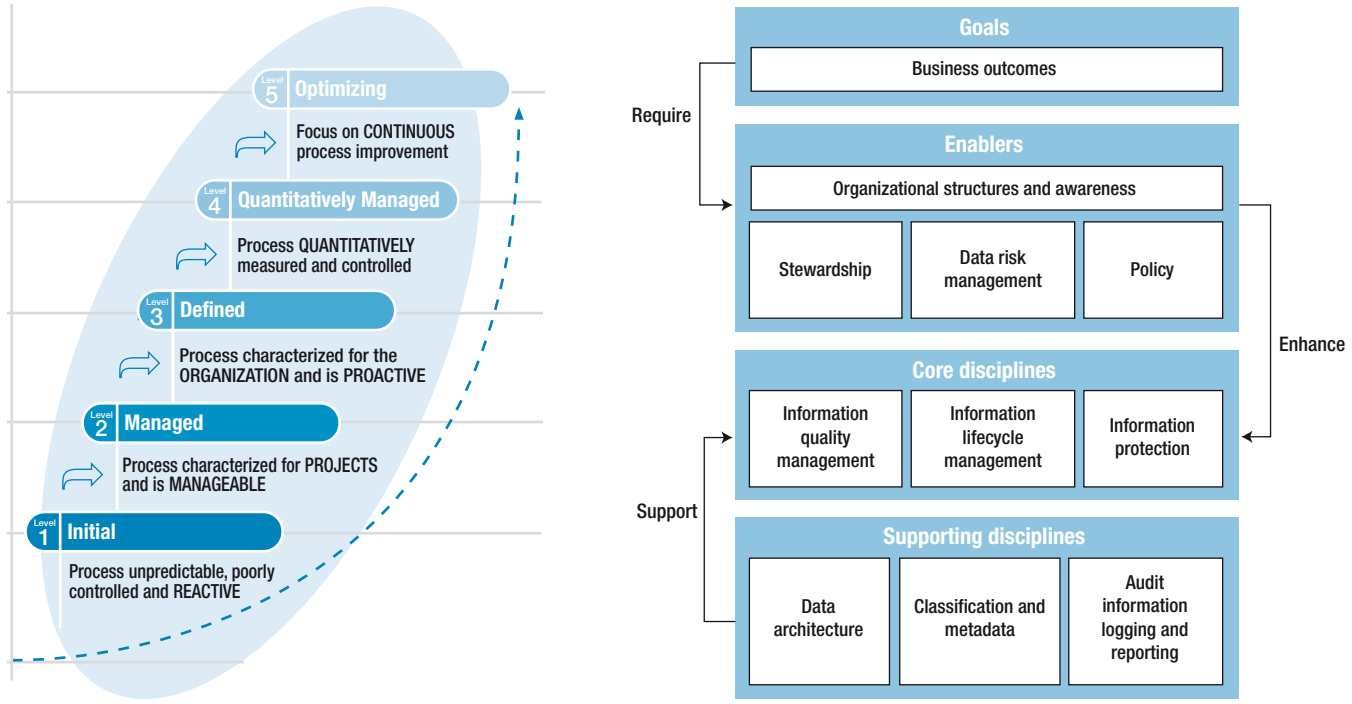

The IBM Information Governance Council Maturity Model enables organizations to identify their current level of maturity, setting the stage for movement to higher levels. The core disciplines in the diagram at right are targeted to an organization's immediate needs.

In addition to user-focused governance initiatives, IBM has structured three "entry points" designed to make information governance solutions more consumable and targeted to an organization's more immediate needs. These entry points are:

- Information quality management
- Information lifecycle management
- Information protection

Increasing revenue with information quality management

Poor data quality costs organizations millions of dollars each year. Two key solutions from IBM can help address this issue.

IBM InfoSphere Information Server for System z is a software platform that provides breakthrough productivity and performance based on a parallel processing infrastructure. The software helps leverage information across all of its sources, and it delivers the functions required to integrate, enrich and deliver trusted information for key business initiatives. It enables users to:

- Understand all sources of information within the business, analyzing its usage, quality and relationships.
- Cleanse information to assure its quality and consistency.
- Transform information to provide enriched and tailored information.
- Federate and deliver information to make it transparently accessible to people, processes and applications.

IBM InfoSphere Business Glossary for System z helps organizations create, manage and share a controlled, enterprise-wide vocabulary that acts as the common language between business and IT. This is a critical step in better aligning technology with business goals. In addition to a controlled vocabulary, the solution's hierarchy and classification systems provide additional business context.

Actively connected to InfoSphere Information Server metadata services, InfoSphere Business Glossary enables data stewards to link business terms to technical artifacts shared with IBM InfoSphere Data Architect, InfoSphere Information Server or a third-party data integration solution. The result is a common set of semantic tags used by data modelers, data analysts, business analysts, governance stewards, data architects, developers and end users. The solution also serves as a history of records to ensure compliance with regulatory requirements.

Reducing costs with information lifecycle management

One of the greatest impediments to data optimization is uncontrolled data growth, as overgrown databases can impair the performance of business-critical applications.

InfoSphere Optim™ Data Growth Solution for z/OS helps solve the data growth problem at the source through intelligent archive management. They reduce the size of production databases, providing universal access to data while improving application performance and cutting hardware and software costs.

InfoSphere Optim data growth solutions help organizations by:

- Reducing hardware, storage and maintenance costs and freeing valuable resources by archiving historical data from applications and systems.
- Maintaining optimal application performance levels by controlling database growth.
- Archiving, managing and retaining application data according to business and compliance policies.
- Responding quickly to audit and discovery requests with universal access to archived information.
- Managing and controlling archiving across applications, databases, operating systems and platforms.

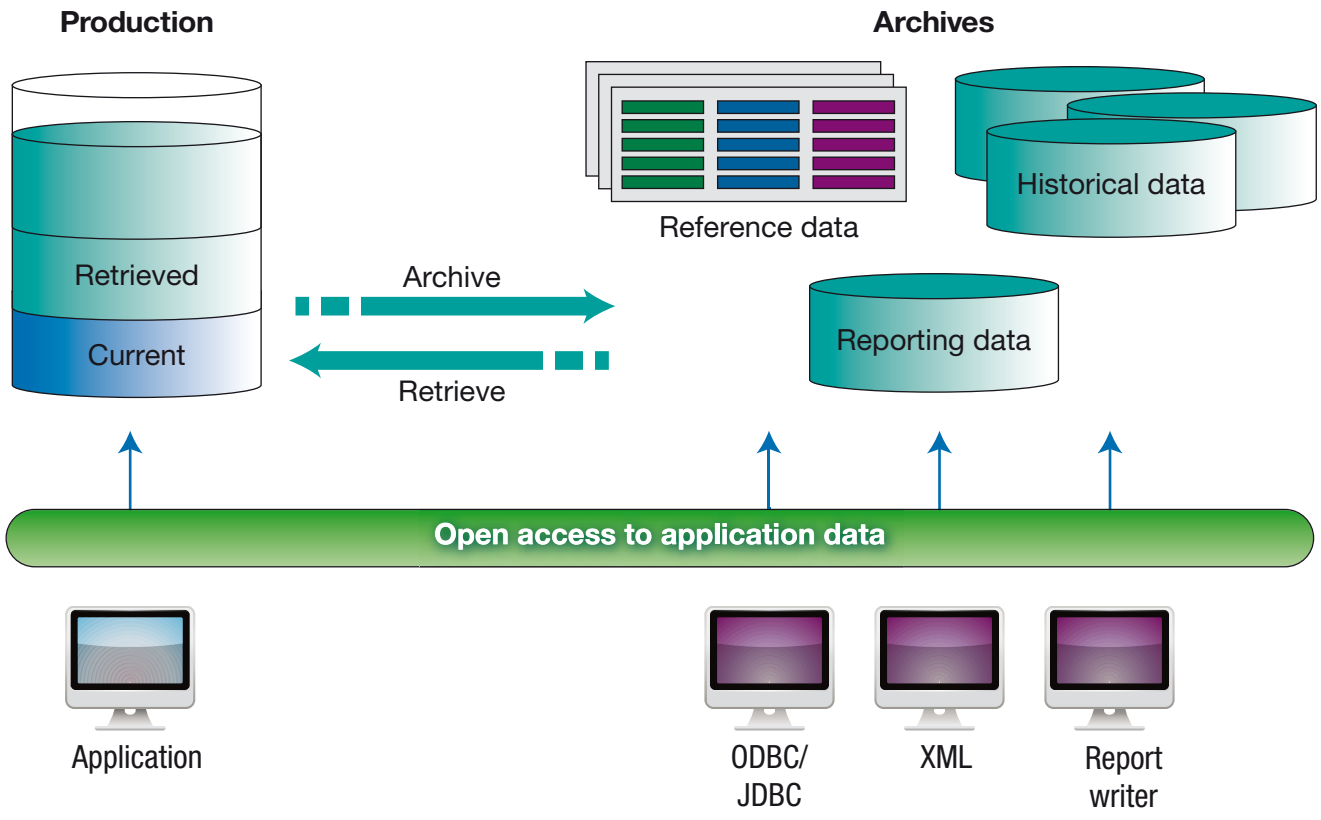

IBM InfoSphere Optim Data Growth Solution for z/OS provides intelligent, application-aware archiving of data.

InfoSphere Optim solutions make it possible to archive infrequently used data and records so that they can be securely and cost-effectively stored elsewhere, while maintaining universal access to them. By reducing the amount of data to sift through, theses solutions help increase the performance of applications as they improve the speed of reporting and mission-critical business processes.

The solution can improve performance of production environments, resulting in significant return on investment.[3] By intelligently archiving infrequently used data to appropriate storage media, it enables the organization to efficiently retain and access data in compliance with regulatory requirements. The solution facilitates application retirement, enabling data to be restored later without having to reinstall the retired application.

Optimizing performance for information lifecycle management

With IT budgets shrinking, organizations must find ways to extract more value from existing assets. Understanding where and how it is possible to improve performance to meet increasingly stringent service-level demands from management and customers can save on expensive and sometimes unnecessary storage and processor costs.

IBM provides capabilities for the System z platform to analyze and determine where performance bottlenecks exist and where they might be likely to occur. Below are some of the key tools that provide capabilities for reducing costs by keeping IBM DB2® for z/OS® and IBM Information Management System (IMS™) data and applications optimized.

DB2 performance optimization solutions

DB2 Buffer Pool Analyzer for z/OS optimizes database performance by helping buffer pools function at their highest efficiency. It supports object placement, including support of inactive objects (table spaces and buffer pools), and provides expert analysis through an easy-to-use wizard. The solution also provides:

- Capabilities for batch trace collection and Interactive System Productivity Facility (ISPF) Collect Report Data functionality.
- Comprehensive reports for browsing or printing, including reporting of buffer pool activity.
- Data collection of virtual buffer pool activity via DB2.
- Simulation of buffer pool usage for varying buffer pool sizes and different object placement.

DB2 Query Monitor for z/OS enables efficient customization and tuning of SQL workload and DB2 objects, arming staff with the ability to spot trouble before it can cause a significant waste in DB2 resources. It can be used to view and configure monitoring across the enterprise from a single console and to identify users and locations that are issuing problem SQL. It offers extensive choices for determining what monitoring information to gather and when to collect it. To make managing data assets more cost-effective, it has GUI-based reporting, viewing and configuration capabilities for access to consolidated data and events for DB2 subsystems across multiple z/OS images.

DB2 SQL Performance Analyzer for z/OS helps DB2 users improve DB2 application design to achieve maximum productivity, and tests different "what if" scenarios quickly and economically to determine the performance of various database design

and production volumes. It can also access any report for any input source (including DBRM, plan or package). Efficient navigation enables quick access to information that's needed most.

IBM Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS evaluates the efficiency of, and optimizes performance from, DB2 databases in a z/OS environment. It provides extended insight into application response time, helps monitor, analyze and tune the performance of IBM DB2 for z/OS and IBM DB2 applications, and provides robust views of performance to help improve productivity. Predefined rules of thumb make it quick and easy to identify performance bottlenecks. It also combines batch-reporting capabilities with real-time monitoring and historical tracking functions.

IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS helps resolve critical performance issues by monitoring, analyzing and optimizing the performance of DB2 Universal Database™ and DB2 on z/OS applications online in real time and in batch reports. It acts as a single engine for performance data collection logic, monitors individual data-sharing members or entire data-sharing groups, and watches applications running in a parallel query environment, even if the parallel tasks are executed on different processors.

IMS performance optimization solutions

IMS Buffer Pool Analyzer for z/OS improves database performance by reviewing the buffer pool environment, projecting the impact of buffer pool changes and recommending changes to the number of buffers. It provides information on the efficiency of OSAM cached buffers, models user changes to the buffer pool configuration to see their impact before they are implemented, identifies wasted storage and determines whether adding or removing buffers will improve performance.

IMS Network Compression Facility for z/OS can help reduce both end-user response time and line costs by compressing data streams to end users of 3270 terminals. It is a practical, affordable tool that uses 3270 commands to compress unnecessary or redundant data—without requiring any application changes or alterations to users' screens.

IMS Performance Solution Pack for z/OS provides an affordable, complete portfolio of IBM database performance management tools for fast, easy analysis of IMS transactions:

- IMS Connect Extensions to improve availability, reliability and performance of IMS Connect.
- IMS Performance Analyzer to provide information on IMS system performance.
- IMS Problem Investigator to help determine the cause of problems and trace the flow of events end to end.

These components collectively help improve productivity for problem analysts, increase the efficiency of IMS application performance, improve IMS resource utilization and increase system availability.

Transaction Analysis Workbench for z/OS enables analysis of transaction performance and behavioral problems and simplifies problem analysis. It extends the scope of traditional analysis techniques to make it easier to identify problems. The workbench saves a history of each problem session, attaches logs and other historical data to problem sessions, analyzes the data files associated with a problem, and provides interactive performance analysis to track a transaction and identify significant events in its life cycle.

Tivoli OMEGAMON XE for IMS on z/OS helps optimize the performance and availability of IMS systems, reducing the potential for delays or outages by reporting on a number of critical IMS attributes. It provides a single point of control over IMS in Parallel Sysplex® environments and offers rich monitoring of IMS Connect TCP/IP transaction requests.

Reducing costs with information protection

A 2009 report by the Ponemon Institute showed that data breaches are increasing yearly, with the average cost of a breach now reaching US$6.65 million.[4] The report put the cost per compromised record at $202, comprising $50 in direct costs and $152 in indirect costs such as customer churn due to lack of faith in the company's ability to protect their data, and reduced ability to attract new customers as a result of damaged confidence in the brand. Interestingly, the cost per loss has risen by 38 percent over just the past four years. A subsequent Ponemon Institute report noted that the most expensive data breach cost a company $31 million to resolve, while the least expensive was $750,000.[5]

IBM provides a set of comprehensive capabilities to enhance information protection and lower costs by helping reduce the numbers of data breaches. These solutions help secure access, provide data encryption, ensure that privacy controls are in place and combine powerful but flexible analysis and reporting tools for:

- Security: preventing, controlling, restricting and monitoring access so only those who are authorized can use the data.
- Privacy: de-identification of data that enables organizations to substitute masked data for sensitive data, as well as encryption of data sources.
- Audit of data: collecting and reporting on who and which applications have access, what levels of privileges they have and what users have been doing with data.

Tivoli Security Management for z/OS includes IBM Security zSecure Admin, IBM Security zSecure Audit for RACF® and IBM Security zSecure Command Verifier. This comprehensive security solution can help organizations automate compliance reporting, enhance their security posture to reduce risks, and improve business agility and reliability. The solution includes the ability to:

- Facilitate compliance with security requirements and policies.
- Achieve effective administration of mainframe security and improve productivity.
- Monitor and audit for threat incidents.
- Audit configurations and resource use to help detect and prevent security exposures and report compliance.
- Integrate IBM Resource Access Control Facility (RACF) capabilities, mainframe security capabilities and other Tivoli security management solutions for comprehensive identity management and control.

Optim Data Privacy for z/OS offers comprehensive, proven capabilities for de-identifying application data for non-production environments, including those in which significant amounts of development and testing are conducted by contractors and offshore companies. Its key capabilities are to:

- Protect sensitive data as it is used for other purposes.
- Remove, mask or transform elements that could identify an individual.
- Ensure that masked or transformed data remains appropriate to the context.

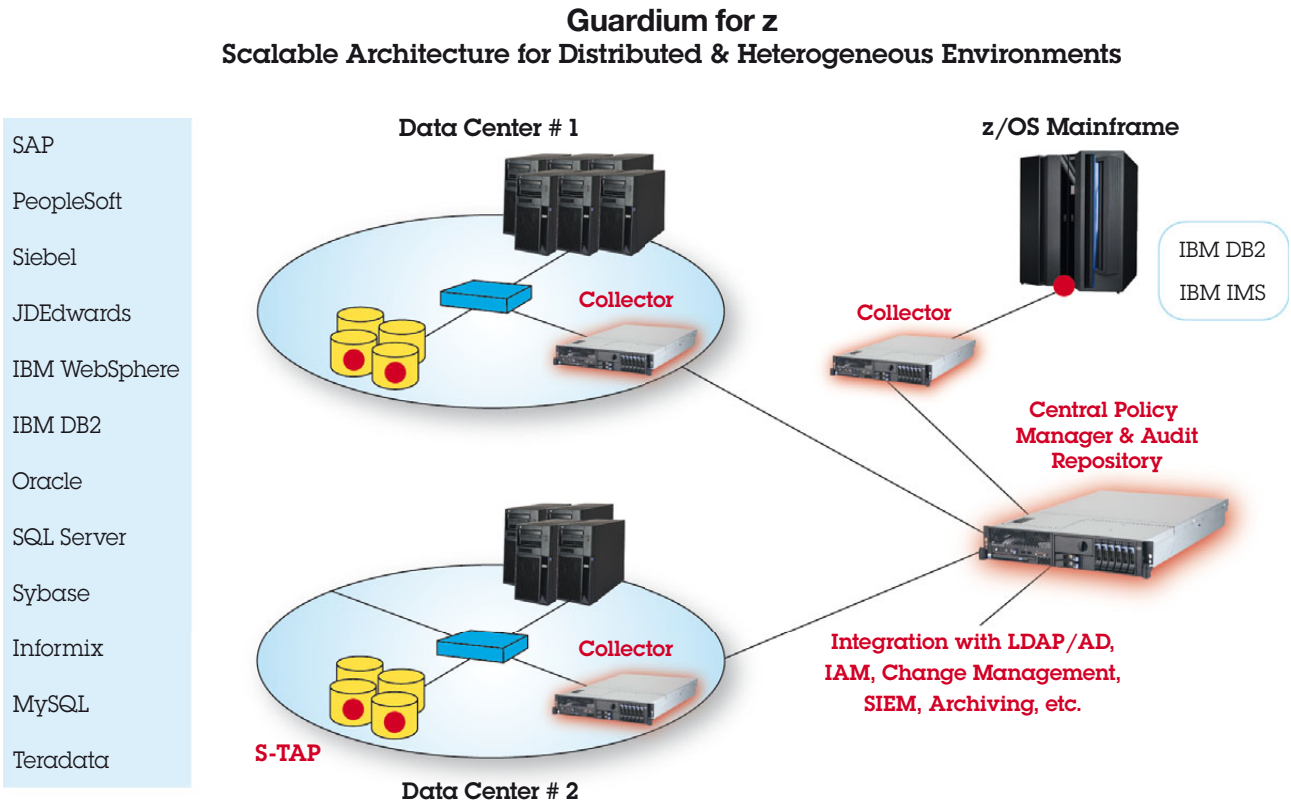

Built on a single unified console and back-end data store, InfoSphere Guardium offers a family of integrated modules for managing the entire database security and compliance life cycle.

InfoSphere Guardium for System z provides the auditing capabilities IT organizations need in order to help minimize the liability associated with growing compliance demands. It includes the ability to:

- Provide details on who made changes to data as well as where and when the changes were made.
- Offer flexible processing modes to collect data and to integrate with IBM DB2 Query Monitor.
- Free up valuable IT staff resources by allowing auditors to pursue data auditing activities with less involvement by the database administrator.
- Allow auditors to generate their own reports automatically and export the data into other applications such as spreadsheets.
- Selectively audit inserts, updates, deletes and reads in DB2 databases.
- Provide real-time alerting of potential breaches of policies.

Why IBM Systems z for enterprise data needs

It is estimated that 95 percent of Fortune 1000 companies store business data on IBM System z.[6] Considering that the System z platform runs the world's banks and ATMs, and is used to reserve airline seats, buy and sell shares on the stock market, track the world's parcels and more, most people interact frequently with processes involving data hosted on a System z platform without even realizing it.

While some solutions offer weeks or months of uptime, the IBM System z platform offers uptime in terms of years. It provides disaster recovery configurations and is designed to deliver 99.999 percent application availability where Parallel Sysplex is implemented, minimizing planned downtime as well as downtime associated with equipment failure or the complete loss of a data center.

It also employs some of the most advanced security technologies in the industry—hardware, software and operating system encryption solutions, access control management, and extensive auditing features—helping organizations meet rigid regulatory requirements.

For many organizations, Systems z is the ideal hub for information governance solutions. It is the optimal platform for workload consolidation, both traditional and new. Deploying an organization's data, transactions, processes, business applications and analytic software on the resilient System z platform can help ensure high levels of business integrity.

Virtualization with System z: The power to simplify

Virtualization on System z allows organizations to do far more work and to achieve greater cost savings on one server footprint. Rather than every application having its own servers for production and development, virtualized environments share resources and simplify operations using as few servers as possible.

Studies have shown that the IBM zEnterprise™ System helps achieve greater savings, operational simplification and reliability by optimizing and consolidating resources. For example, organizations can use this system to:

- Consolidate workloads up to 60 percent faster and at a 33 percent lower price using zEnterprise Integrated Facility for Linux (IFL) processors.
- Increase energy savings by up to 75 percent as the system scales.[7]
- Reduce acquisition costs by up to 70 percent[8] and boost staff productivity by up to 70 percent[9] compared to virtualized x86 alternatives.

As a result, organizations can reduce data center sprawl, electricity costs and floor space requirements. Virtualization also helps organizations save money and grow business by:

- Aligning IT resources with business results.
- Enabling assets to be purchased and provisioned as they are needed, rather than as a hedge against workload spikes.
- Allocating additional capacity as soon as the demand presents itself, rather than losing business while waiting for new systems to come online.

With virtualization, organizations have the choice: continue to grow distributed servers in racks and footprints taking up an increasing amount of floor space, or grow in a concentrated way on System z and save both space and expense.

For more information

To learn more about how IBM information governance solutions can help lower costs and identify revenue opportunities visit: **ibm.com**/software/data/db2imstools/solutions/data-governance.html

[1] Davidow, William H., Marketing High Technology: The Insider's View, Free Press, 1986.

[2] Davis, Judith R., "Information Governance as a Holistic Approach to Managing and Leveraging Information," BeyeNetwork, 2010. Prepared for IBM. ftp://public.dhe.ibm.com/software/os/systemz/IBM_Information_Governance_Survey_Report.pdf

[3] "The Total Economic Impact™ of IBM Optim Integrated Data Management Solutions: Multicompany Study," Forrester Consulting, October 2009. Prepared for IBM. public.dhe.ibm.com/software/data/sw-library/data-management/optim/papers/forrester_optim_total_economic_impact.pdf

[4] "Fourth Annual US Cost of Data Breach Study: Benchmark Study of Companies," Ponemon Institute, January 2009. www.ponemon.org/local/upload/fckjail/generalcontent/18/file/2008-2009%20US%20Cost%20of%20Data%20Breach%20Report%20Final.pdf

[5] "2009 Annual Study: Cost of a Data Breach: Understanding Financial Impact, Customer Turnover and Preventive Solutions," Ponemon Institute, January 2010. www.ponemon.org/local/upload/fckjail/generalcontent/18/file/US_Ponemon_CODB_09_012209_sec.pdf

[6] "Business resilience: Ensuring continuity in a volatile environment," Economist Intelligence Unit, 2007. Sponsored by ACE, IBM and KPMG. **ibm.com**/services/us/bcrs/pdf/wp_business-resilience.pdf

[7] Based on zEnterprise comparison to virtualized x86 alternatives

[8] Based on three-year acquisition costs for large-scale, enterprise-class workloads

[9] Based on life cycle management testing of large-scale virtual server environment conducted by IBM

IBM Systems and Technology Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
April 2011

IMW14561-USEN-00