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

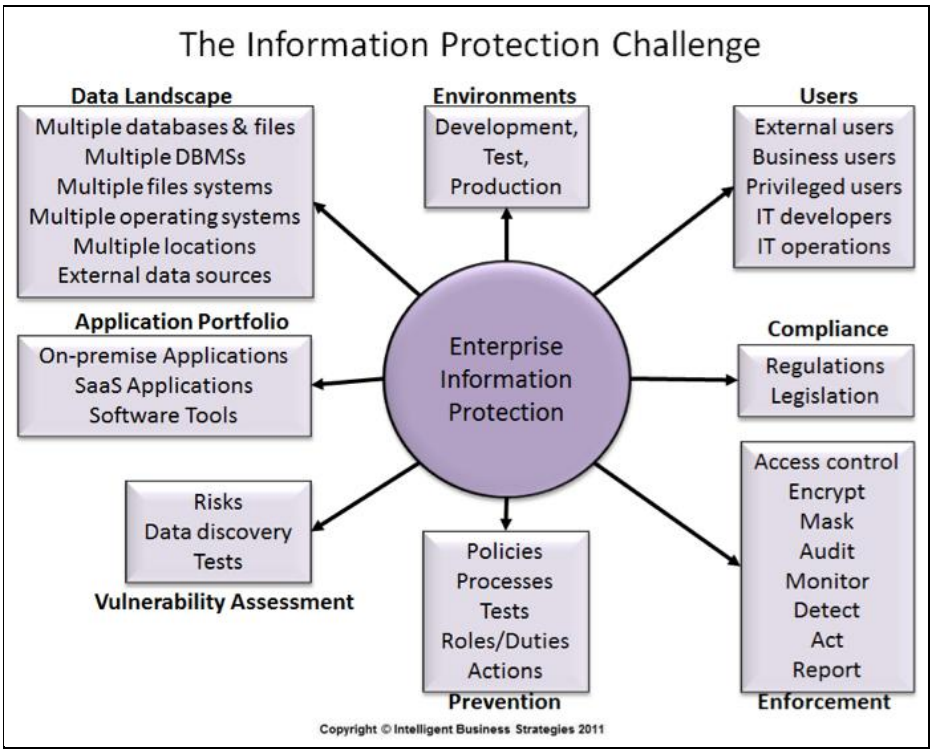

Figure 1

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

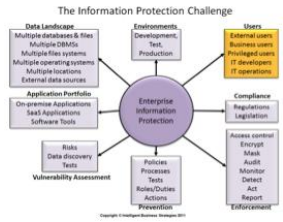

- It should be possible for enterprise information audit and protection software to be able to distinguish between:
  - External users
  - Internal business end users
  - Privileged users e.g. database administrators
  - IT developers
  - IT operations personnel

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

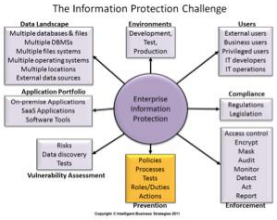

These requirements define the capabilities enterprise information audit and protection software need to provide to minimise exposure to data risks across DBMSs and file systems in the enterprise. The requirements are as follows:

- 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:

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

All types of VSAM dataset can be protected

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