WHITE PAPER

INTELLIGENT
BUSINESS
STRATEGIES

Information Protection on IBM System z

By Mike Ferguson
Intelligent Business Strategies
September 2010

Prepared for:

IBM

TABLE OF CONTENTS

AN INTRODUCTION TO GOVERNANCE, RISK AND COMPLIANCE

Corporate failures in that last five to ten years have resulted in a greater focus on GRC to raise the bar in striving to achiev higher quality business practices

In the last decade we have seen a significant number of major corporate failures in various countries around the world. Some of these failures have been have been due to poor management practices (also known as *governance*). This has resulted in new *compliance* legislation and regulations being introduced in an attempt to force companies to raise the bar in terms of higher quality processes and business practices. In the last few years it has been the collapse of the financial services industry that has caught the attention of many. In this case, many of the banks that have collapsed have done so because of poor credit *risk management* practices. Lending money to higher-risk customers in the sub-prime mortgage market ultimately resulted in many of these customers not being able to pay their mortgages. As a result many banks collapsed or were taken over by governments or other banks.

The focus from all of this fallout has come down to three main areas that are closely related. These are:

- Corporate governance
- Risk management
- Enterprise compliance

Together these are known as governance, risk and compliance (GRC). Wikipedia provides a diagram (see figure 1) of these three areas and how they are integrated. It also shows that a combination of strategy, people, processes and technology are needed to get the GRC problem under control.

Companies need a GRC strategy to help manage their business

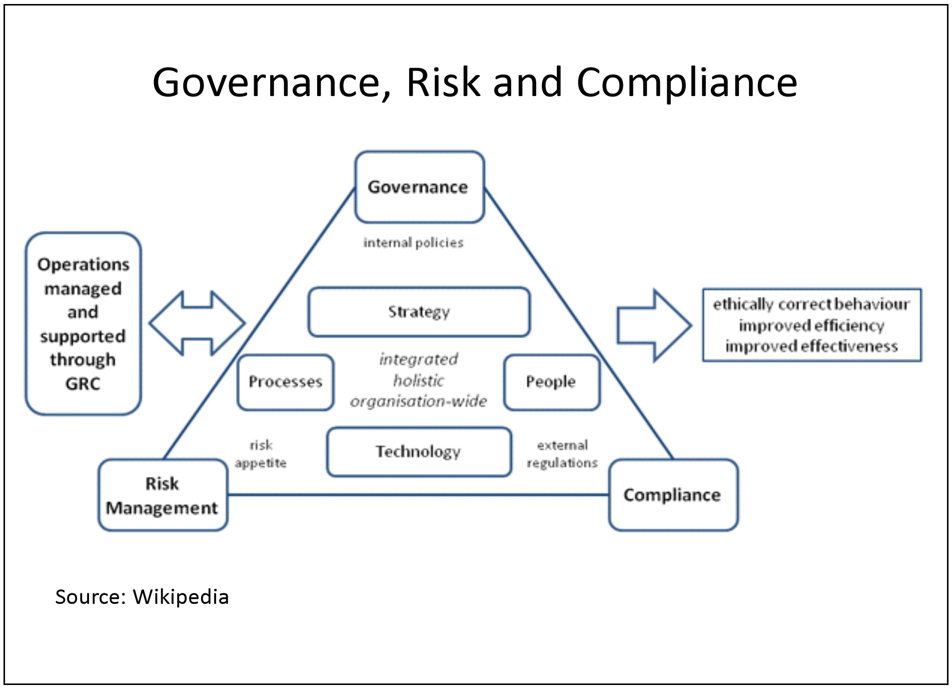

Source: Wikipedia

Figure 1

WHAT IS INFORMATION GOVERNANCE?

Governance is about accountability and control

Underpinning any GRC strategy is the dependency on information. Companies rely on structured, semi-structured (e.g. XML) and unstructured information to govern and manage a business, to manage risk, and to help the business remain compliant. Therefore, information itself must be governed. 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"

The Information Governance Council[1] incorporates three "entry points" into its definition of information governance. These are information quality, information lifecycle and information protection.

This is about introducing accountability and approval controls into information management disciplines to manage information. These disciplines include:

- Data naming and data definitions
- Enterprise metadata
- Data modelling
- Enterprise data quality
- Enterprise data integration
- Master data
- Data privacy and access security
- Enterprise content

It is also about standardisation and integration of information.

Risk management is about improving the ability to mitigate risk

Part of an information governance program is to ensure that information is *protected* to help mitigate risk. Information protection 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 is also about being able to quickly detect information risks that occur and respond in a timely manner to minimise the impact of these *events* on the business as a whole. Controlling access to information and masking sensitive data to uphold privacy is at the heart of information protection. This is central to GRC. Without information protection, there is no solid foundation on which to implement a GRC strategy.

This paper, therefore takes a closer look at information protection to understand what is included, the tasks involved enforcing that protection and the requirements associated with protecting information as it flows through an

[1] See www.infogovcommunity.com.

information supply chain. It then looks at how one platform together with its infrastructure software, supports these information protection requirements to see how it stacks up as a foundation for GRC. That platform is the IBM System z.

WHAT IS INFORMATION PROTECTION?

The purpose of an information protection program is to lower business risk

Information protection is a central part of any information governance strategy and the purpose of an information protection program is to lower business risk. This can be done by reducing the threat of data breaches, monitoring and reporting potential problems and acting before they become major issues. Information security can therefore be defined as:

"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
- Information availability

INFORMATION CONFIDENTIALITY

Information confidentiality is about implementing processes, policies and technology to prevent the disclosure of information to unauthorized individuals or systems either within or outside the enterprise.

Sensitive and confidential data needs to be protected from data breaches

The challenge in information confidentiality is to share data while protecting sensitive information. When sensitive data gets 'into the wild' it is known as a *data breach* and is of particular concern when customer data is exposed and gets into the wrong hands. Some employee data is also sensitive. Companies need to plan for data breaches and do everything possible to prevent them from happening. Data breaches can lead to fraud, major customer dissatisfaction, loss of business, liability lawsuits and regulatory penalties. They can also significantly damage a company brand and corporate reputation.

Masking sensitive data in test and production ebvironments is important

Redundant copies of sensitive and confidential data must also be taken into account

Information protected must guard against sensitive data being exposed to unauthorised personnel whether they are outside or inside the enterprise. To counter this problem, there needs to be a way of identifying, masking and encrypting sensitive data so that it is not seen by unauthorised business users. Note that unauthorised exposure to IT developers also needs to be guarded against. This issue is often related to test data. Many companies create test data sets from full copies of production data to accurately test real world scenarios. This introduces a potential exposure risk when sensitive data is involved in the testing. If several test data sets are created from production data then there will likely be multiple instances of sensitive information stored in various test environments which increases the chances of sensitive data theft by a malicious employee.

Information confidentiality can be protected by managing:

- Information access control (authentication and authorization)
- Information privacy
- Information encryption (hardware and software encryption)
- Monitoring and auditing

INFORMATION INTEGRITY

Information integrity is about implementing processes, policies and technology to prevent data being modified without authorization. Figure 2 shows how this can easily become an issue using manufacturing as an example. Each time the manufacturing order-to-cash process executes, subsets and aggregations of master and transaction data are stored in many different systems.

Sunsets and aggregations of data are typically stored in many different systems

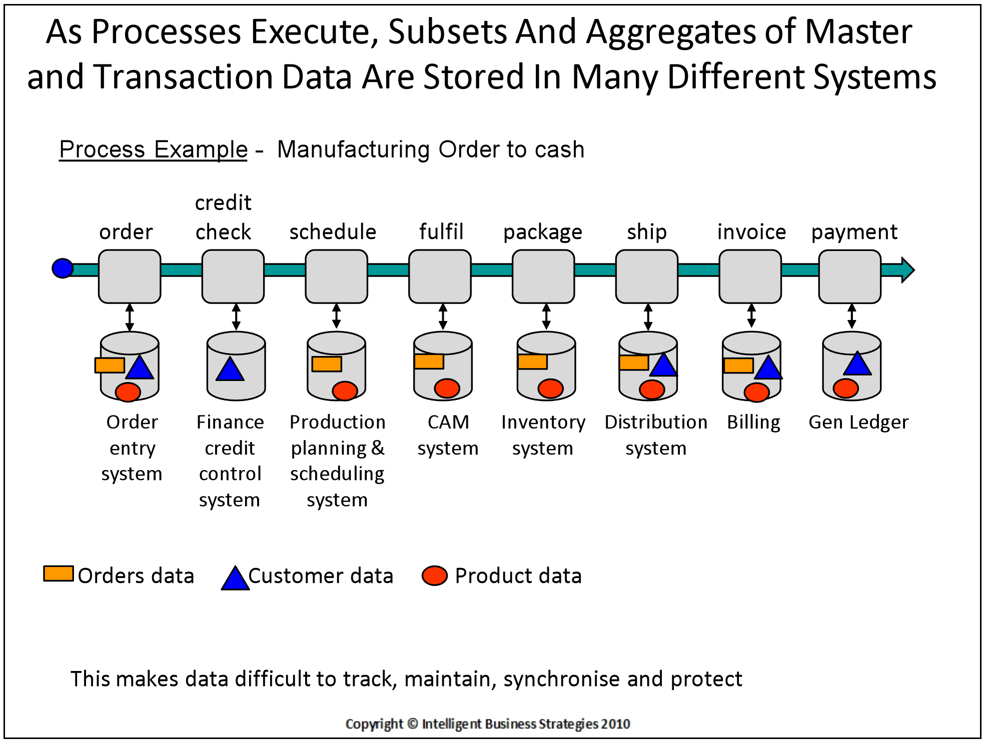

As Processes Execute, Subsets And Aggregates of Master and Transaction Data Are Stored In Many Different Systems

This makes data difficult to track, maintain, synchronise and protect

Copyright © Intelligent Business Strategies 2010

Figure 2

Data access security must also be managed to prevent unathorised maintenance of sensitive data by employees

A common authentication and authorisation mechanism simplifies management of data access security

Because subset copies of data are deposited across multiple operational systems and multiple departments, it opens up the possibility of unauthorised maintenance of data. Solving this integrity problem requires security privilege co-ordination across multiple portals, applications and database technologies. It is further complicated by that fact that many applications, even today, have their own 'baked in' proprietary authorisation models. In addition, the data that you may be trying to protect may be known by different data names in the different underlying systems that underpin a company's core business processes. Although many companies are gradually moving towards common authentication and authorisation mechanisms (e.g. corporate LDAP directory), they are still in a 'hybrid' state with a complex authentication and authorisation landscape. It is this kind of complexity that leaves to door open to data access security breaches. All kinds of risks can stem from these kinds of breaches. They include fraud, process delays due to data defects, unplanned increases in operational costs and customer dissatisfaction to name a few. Information integrity can be protected by managing

- Information access control (authentication and authorization)
- Information quality
- Monitoring and auditing

Information Availability

Information availability is about implementing processes, policies and technology so that information is available where and when it is needed. This means that the applications and databases used to process and store the information, the security controls used to protect it, and the connectivity needed to access it must be functioning correctly.

Being able to backup and recover master data and transaction data entities irrespective of location of the data acrosss the enterprise would significantly reduce risk

The issue with availability is to do with data backup, data retention and data recovery. If data is not recoverable it can have a massive impact on the company's ability to operate. Processes can break and come to a halt which can severely impact costs, revenue and customer confidence. While this may seem a straightforward issue to solve, it is often complicated by the fact that subsets of data may have proliferated across departments and application systems as business processes execute as shown in Figure 2. Therefore guaranteeing availability of certain types of data may be far more complex than originally anticipated. Maintaining data relationships and data integrity across all systems in the enterprise is critical to continuous operation and availability. In addition, data retention is important for operational and compliance reasons. It follows that so called 'enterprise backup' and 'enterprise recovery' may be needed in the most severe cases. Being able to backup, retain and recover data irrespective of the location of that data is therefore desirable. While this is not typically happening today, there are most certainly aspirations in many companies to be able to do this as it would significantly simplify current practices while also reducing risk. The challenge is to be able to identify where all the data (and corresponding data relationships) are located in the enterprise.

Information availability can be protected by managing

- Data retention
- Data backup, archive and recovery
- High availability of systems to prevent service disruptions due to power outages, hardware failures, and system upgrades.

Information Protection Strategy

Key metrics are used to measure the success of an information protection program

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. These are associated with how the company identifies and controls 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. These changes also need to be monitored and be auditable.

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

Companies also need processes and 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 be able to respond in a robust way if they occur.

INFORMATION PROTECTION SCOPE

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

In any business, 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. In addition, it may flow into a data warehouse 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, product, order 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

Information Protection Implementation Tasks

For any information protection initiative to be successful it should be part of a wider information governance program that formally defines the data to be governed and then sets about defining policies to control information quality, privacy, access security, maintenance, retention, backup and recovery. Once this is done we can then start to make use of data management and security technologies to help implement the policies to govern and protect the information.

Defining Information To Be Governed

Companies need to identify 'at risk' data

The first task therefore is to define the information to be governed. This can be done incrementally by focussing on specific master data entities (e.g. customer, product, asset, suppliers…) transaction types (e.g. orders) used in the enterprise and then building out to gradually govern all master data and transaction data. Data is defined using a common set of data names, definitions and integrity constraints sometimes referred to as a shared business vocabulary (SBV).

Having a shared business vocabulary makes it easier to identify 'at risk' data

The SBV is held in a business glossary which is accessible by both business users and IT professionals. The business glossary provides business users with a place to go where they can see what data exists in the enterprise, what it means, who owns it and what policies have been applied to it.
In terms of definitions, each data item should have:

Every data item needs to be defined

- A data name (term) and data definition
- A description
- Integrity constraints (policies)
- Approved synonyms
- Related terms
- Languages it must be rendered in
- A sensitivity classification e.g. low, medium, high
- An assigned data steward
- A set of information governance policies that include
  - Data quality policies
  - Data privacy policies
  - Data security policies
  - Data retention policies
  - Data lifecycle policies around who is authorized to create, read, update and delete
- Where the data item is used e.g. tables, files, reports

Policies need to be defined so that integrity, data quality, security and privacy can be managed

Also, a history of who changed names, definitions, integrity constraints and policies should be kept together with an audit of when that happened and what the before and after versions were

Highlighting Information To Be Protected

Once master and transaction data items have been defined within the SBV, we then need to identify what subset of data items within the SBV need to be protected. The best way to do this is to highlight data items that are potentially 'at risk' from a particular threat. This would include

- Sensitive or confidential data items that could be subject to breaches in data privacy
- Restricted data items that could be accessed by unauthorised users
- Data items used in multiple applications that need to be kept in sync otherwise the lack of cross system integrity would disrupt business operations
- Data items considered critical to business operation and decision making which may cause significant business problems if they were not available to business users at the time they need it
- Data items that need to be retained for compliance reasons

'At risk' data can be associated with threats to indicate what could go wrong if it is not fully protected

For example, customer data may be considered sensitive and therefore potentially 'at risk' in terms of a data privacy breach or unauthorised access. Similarly, employee data items like Social Security Number and Salary may be considered confidential.

Understanding the business impact of data risks helps prioritise information protection efforts

It is also important to understand the business impact of data risks. For example is the impact financial? Is it operational and if so what resources are needed to fix the problem? Does it result in regulatory compliance violations or does it impact the company brand and reputation?

DEFINING POLICIES TO PROTECT HIGHLIGHTED INFORMATION

Having identified data in need of protection and the types of threats that this data may be subject to, the next step is to define the information protection policies needed to reduce the possibility of protection violations occurring. Several types of information protection policies may need to be defined for each 'at risk' data item. These include:

Policies need to be defined for each 'at risk' data item so that information can be protected

- Data privacy policies
- Data lifecycle policies around who is authorized to
    - Create
    - Read
    - Update
    - Delete
- Data synchronisation policies
- Data retention and archive policies
- Data backup and recovery policies

These policies may also be restricted by scoping limitations to limit their information protection capability e.g. to allow certain information protection policies to be set for data in a specific application used in a specific business process by people in a specific organisational area. However if data spans organisational units and systems then enterprise wide common information protection policies should be enforced.

Determining Where Information Is Located In Order To Protect It

Data 'at risk' needs to be located if information protection is to be successfully managed

Having identified the data items in the SBV that need to be protected, and defined policies to be applied to those items to protect them, the next challenge is to identify where that information actually resides in the enterprise. This includes identifying all redundant copies of it whether they are in production systems or test systems. This is necessary because if we don't know where data is located, we cannot fully protect it. We need to locate this data and then map it back to the SBV in order to determine:

- Whether or not data privacy is enforced on confidential and sensitive data items
- If authorised access to restricted data items is enforced across the enterprise
- If data synchronisation is working across all systems where that data needs to be available for business use and if the synchronization is complete
- If critical data is backed up and recoverable across all systems that use it

Data discovery technology helps to locate 'at risk' data quickly

The key point here is that a data discovery and mapping exercise is needed to find data across the enterprise that needs to be protected. This task may be done manually or automatically. If done manually the discovery task will be expensive and very time consuming. Therefore vendors that provide an automated data discovery tool to locate data and discover data relationships across multiple systems in the enterprise offer a distinct advantage. By automatically mapping located data to the SBV in the business glossary, it becomes possible for these tools to locate data items flagged in the SBV as sensitive, confidential, restricted and business critical also to determine if information protection policies are enforced.

Implementing Information Protection Policies

Having identified where the data that needs to be protected actually is located across the enterprise, the next step is to make use information governance and information management processes and technologies to protect it so that policy violations do not occur.

This can be done by

1. Applying information protection policies to discovered at risk data items to reduce protection violations
2. Identifying the source or cause of each information protection policy violation and making changes to reduce the chances of further occurrences
3. Monitor 'protected data' to see if policies and procedures are working

Applying Information Protection Policies

Tools in a data management technology platform can be used to apply and enforce information protection policies

Looking at these in more detail, the first is the application of information protection policies to discovered data that has been highlighted as being at risk. The way this is done is to share metadata between tools in a data management platform. Figure 3 shows the tools needed to enforce information protection. They include:

- A business glossary
- A data and data relationship discovery tool
- A data privacy masking tool
- Data encryption software
- Data retention and archive tools
- Data backup and recovery tools
- User access authentication and authorisation tool to manage single sign-on and authorizations across multiple systems.

A suite of integrated tools is needed to protect information

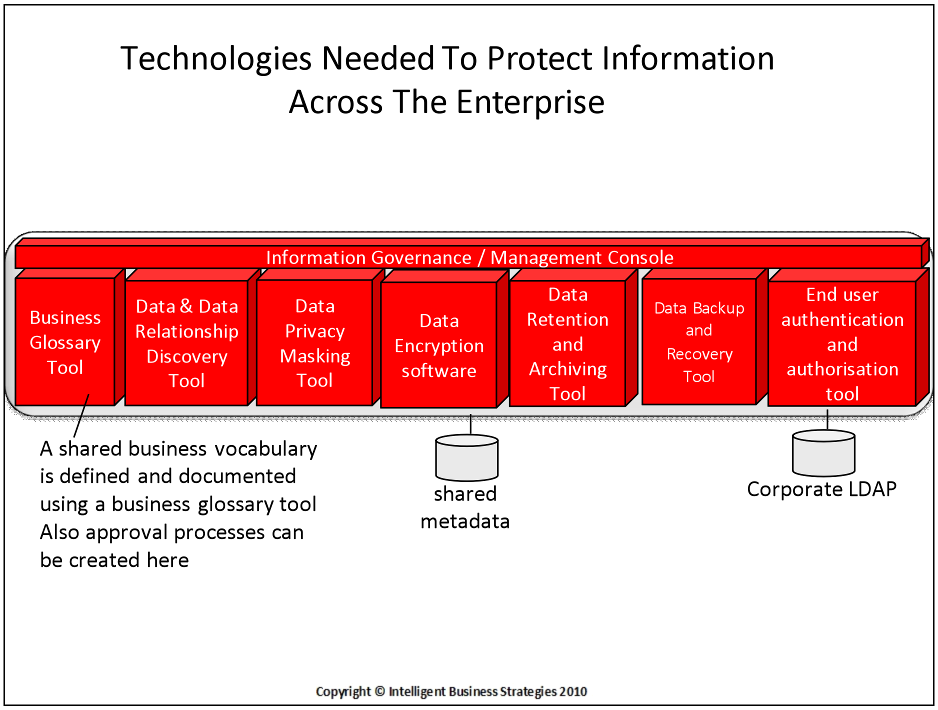
Technologies Needed To Protect Information Across The Enterprise
Information Governance / Management Console
Business Glossary Tool
Data & Data Relationship Discovery Tool
Data Privacy Masking Tool
Data Encryption software
Data Retention and Archiving Tool
Data Backup and Recovery Tool
End user authentication and authorisation tool
A shared business vocabulary is defined and documented using a business glossary tool Also approval processes can be created here
shared metadata
Corporate LDAP
Copyright © Intelligent Business Strategies 2010

Figure 3

The key to this is the metadata sharing between the data discovery tool and the other information protection tools.

Automated data discovery seeks to find complete data entities across heterogeneous systems e.g. customer

Automated data discovery seeks to find complete business data objects across heterogeneous systems. An example of such an object is customer. This means identifying all customer data and data relationships across the enterprise irrespective of location. This is shown in Figure 4.

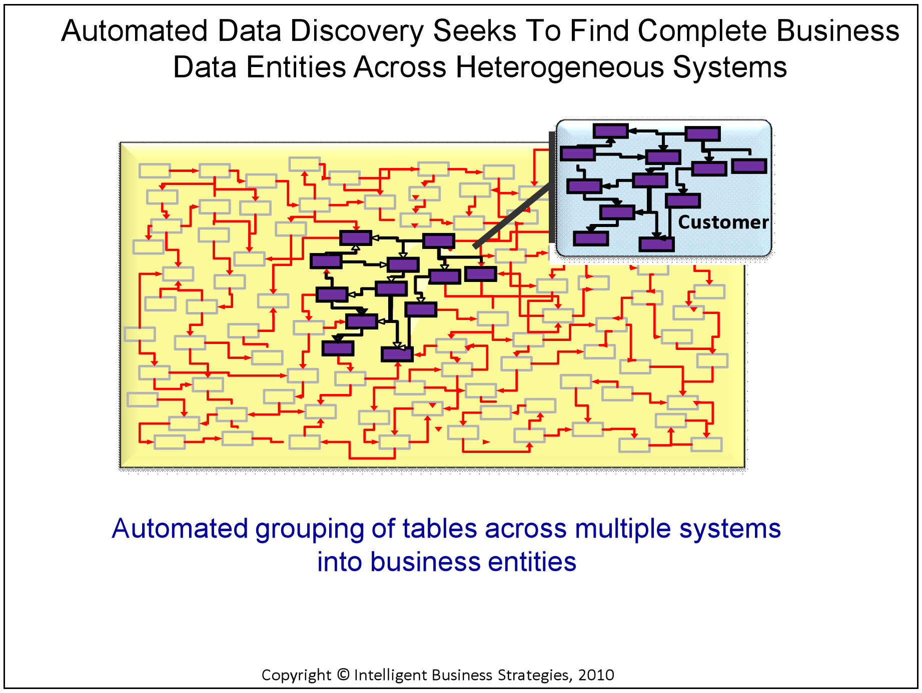

Automated Data Discovery Seeks To Find Complete Business Data Entities Across Heterogeneous Systems

Automated grouping of tables across multiple systems into business entities

Figure 4

This means that it becomes possible to know where all 'at risk' data items associated with a discovered data entity are located

By finding complete business data entities (e.g. customer, product, order) it becomes possible to protect information at the entity level irrespective of where the data is located. For example, customer data privacy, customer data retention, customer data access security, and customer data recovery. This is because data discovery has already identified where the data is and by mapping it to an SBV it becomes possible to identify which information protection policies (privacy, retention, authorised access etc.) that are 'in play' for a particular business data entity.

By sharing metadata generated during data discovery with other tools it becomes possible to apply already defined information protection policies to manage risks. This is shown in Figure 5.

To protect confidentiality and prevent beaches in data privacy, metadata about discovered sensitive and confidential data items can be shared with data privacy tools to allow them to mask this data in the systems where they are located and also to manage the generation of test data sets.

To protect information integrity metadata can be shared with multiple technologies to co-ordinate consistently application of end user access privileges for complete business data entities level irrespective of the location of the data. This together with single sign-on and the use of a corporate LDAP directory can be used for enterprise management of authentication and authorisation.

To protect information availability, it becomes possible to archive complete data entities in accordance with data retention policies because data discovery has identified all data items and data relationships across heterogeneous systems in the enterprise irrespective of location that need to be retained. The same is true for data backup and recovery, it becomes possible to back up and recover complete data entities irrespective of location.

By sharing the location of all 'at risk' data with other information management tools, it becomes possible to apply information protection policies to data across the enterprise

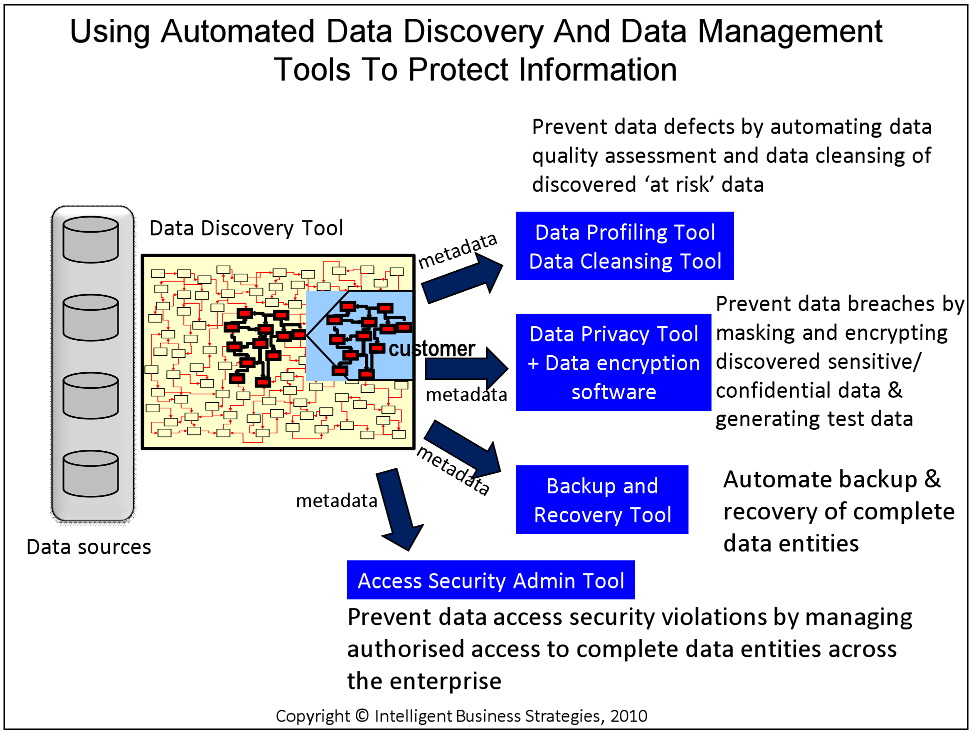

Using Automated Data Discovery And Data Management Tools To Protect Information

Prevent data defects by automating data quality assessment and data cleansing of discovered 'at risk' data

Prevent data breaches by masking and encrypting discovered sensitive/ confidential data & generating test data

Automate backup & recovery of complete data entities

Prevent data access security violations by managing authorised access to complete data entities across the enterprise

Copyright © Intelligent Business Strategies, 2010

Figure 5

Sourcing Data Risks

In addition to being able to apply policies and controls to protect information, it is also very important to be able to determine the source of information protection policy violations to prevent future occurrences of these. The source (cause) of these events can be an employee, an external user (e.g. customer, partner, supplier), an application or script running inside the enterprise or an application or script running outside the enterprise e.g. in the cloud, or on a customer or partner system.

Data risks can be sourced and changes made to reduce risks even further

Once the source has been identified, changes to processes, training and testing procedures may be needed to help reduce policy violations. It may also be the case that an information security council is introduced with a formal approval process to govern information protection to tie this to internal controls.

Monitoring and Auditing To Protect Information

Monitoring and auditing information protection policy violations helps companies stay in control

The final step in information protection is being able to monitor and audit activity to understand the frequency and severity of any information protection policy violations. This means that business data owners need to have access to reports and dashboards that highlight problems. Auditors also need to be able to carry out their roles and responsibilities to check if legislation and regulatory compliance requirements are being enforced and to see if any actions need to be taken with respect to follow up on policy violations.

INFORMATION PROTECTION ON IBM SYSTEM Z

Having understood what information protection is and the tasks involved in implementing it, this section of the paper looks at how one hardware platform stacks up in terms of helping an organisation enforce information protection. That platform is the IBM System z platform.

THE IBM SYSTEM Z SERVER

The IBM System z platform has a track record in running mission critical systems

IBM System z is a modern, robust platform for mainstream operational and analytical computing. This modern mainframe server has a strong history of running some of the most high volume transaction processing systems in the world including some very large e-commerce systems. Therefore there are a large number of core operational systems running on z/OS. One of the reasons for this is its track record on high availability. IBM System z also includes a powerful workload management capability to managing complex mixed workloads.

The zEnterprise platform processors can support more workload sets andmore in-memory data

The latest generation of IBM System z is the zEnterprise platform. The zEnterprise processor technology is built on IBM's CMOS mainframe heritage with up to 24 5.2GHz quad core processors arranged in a shared cache symmetric multi-processor (SMP) architecture with and up to 3 terabytes of memory capacity[2]. Large on-chip cache memory, redundant array of independent (RAIM) memory technology (for high availability) and a 64-bit z/OS operating system making it perfectly capable of supporting more workload sets and in-memory data

zEnterprise runs z/OS and z/Linux side-by-side

zEnterprise runs z/OS and zLinux operating systems side-by-side on the same server and can run multiple virtual Linux servers. This allows it to run software natively on z/OS and on zLinux with interoperability between the two environments. zEnterprise is also capable of running multiple z/OS instances in a parallel sysplex configuration to offer scalability and availability.

WHAT DOES SYSTEM Z PROVIDE TO HELP PROTECT INFORMATION?

The IBM System zEnterprise platform offers both hardware and software to help support information protection.

Hardware Capabilities Supporting Information Protection

The zEnterprise platform offers both hardware and software encryption

zEnterprise supports both standard cryptographic hardware and optional cryptographic features. The cryptographic hardware features available include:

- Central processor assist for cryptographic functions (CPACF)
- Configurable Crypto Express3 (CEX3)
- Trusted Key Entry (TKE) workstation and smart card reader

The zEnterprise hardware encryption expedites the encryption of data in motion

[2] Up to 1TB per LPAR

CPACF is an encryption accelerator facility on zEnterprise quad-core processors, which is designed to provide high-speed cryptography. Its primary function is to accelerate the encrypting and decrypting of SSL transactions and VPN-encrypted data transfers. The assist function uses a special instruction set for symmetrical clear key cryptographic encryption and encryption operations. CPACF offers more protection and security options with Advanced Encryption Standard (AES) 192 and 256 and stronger hash algorithm with Secure Hash Algorithms SHA-512 and SHA-384.

Hardware encryption also offers high speed digital signing

Crypto Express3 is an optional feature on the zEnterprise that complements the CPACF. It is suited to applications requiring high-speed, security-sensitive, RSA acceleration, cryptographic operations for data encryption and digital signing, secure management, and use of cryptographic keys, or custom cryptographic applications. These can include financial applications such as PIN generation and verification in automated teller and point-of-sale (POS) transaction servers, remote key loading of ATMs and POS terminals, Web-serving applications, Public Key Infrastructure applications, smart card applications, and custom proprietary solutions.

The optional Trusted Key Entry (TKE) workstation provides security-rich local and remote key management, providing authorized personnel a method of operational and master key entry, identification, exchange, separation, and update. Recent enhancements include support for the AES encryption algorithm, audit logging, and an infrastructure for payment card industry data security standard (PCIDSS), as well as:

- EEC Master Key Support
- CBC Default Settings Support
- TKE Audit Record Upload Configuration Utility Support
- USB Flash Memory Drive Support
- Stronger PIN Strength Support
- Stronger Password Requirements for TKE Passphrase user Profile Support

Support for an optional Smart Card Reader attached to the TKE workstation allows for the use of smart cards that contain an embedded microprocessor and associated memory for data storage. Access to and the use of confidential data on the smart cards is protected by a user-defined personal identification number (PIN).

System z Software Used in Information Protection

The following IBM information management software products that support information protection run on the IBM zEnterprise platform

IBM Information Management Products on System z Required to Implement Information Protection	Usage in Protecting Information
IBM RACF	Managing role based access to data and services
IBM DB2 10 DBMS for Z/OS	Data encryption for DB2 data

IBM System Z also includes a suite of software tools to help protect information

Data Encryption for IMS and DB2 Databases	Data encryption for DB2 and IMS data using System Z hardware
IBM InfoSphere Business Glossary	Manage common data definitions for master data and transaction data and allow business users to highlight information to be protected
IBM InfoSphere Discovery	Discovery of data and data relationships to identify the location of the data to be protected
IBM Optim for z/OS	Managing Data Privacy, information retention and archive processing
IBM Optim Data Redaction	Protects sensitive unstructured data contained in documents & forms from unintentional disclosure
IBM InfoSphere Guarduim	Real-time database activity monitoring Monitor privileged users e.g. DBAs Monitor enterprise application users for fraud Enforce database change control Prevent database leaks

How Can Information Software Be Used To Protect Information on System z?

The above table represents just some of the information management tools that run in the IBM System z stack. These are the components specific to implementing information protection and are representative of the components discussed in Figure 3. Taking into account the information protection implementation tasks already discussed, these components can be used as follows:

Integrated authentication, authorization and single sign-on means access control is managed

- IBM Resource Access Control Facility (RACF) lets you decide which resources you want to protect and which users need access to them. RACF provides the functions that let you:
  - Identify and verify system users
  - Identify, classify, and protect system resources
  - Authorize the users who need access to the resources you've protected
  - Control the means of access to these resources
  - Log and report unauthorized attempts at gaining access to the system and to the protected resources
  - Administer security to meet your installation's security goals

  This includes both authentication and authorisation. It can also be combined with IBM Tivoli Access Manager as part of an enterprise single sign-on architecture. Both application services and information resources can be protected.

- DB2 for z/OS helps to protect information in a number of ways

Database encryption is supported at the column and value level

  - It provides built-in data encryption and decryption functions that you can use to encrypt sensitive data, such as credit card numbers and medical record numbers. You can encrypt data at the column or value level. To do this you must install the Integrated Cryptographic Service Facility to use the built-in functions for data encryption. When you use data encryption, DB2 requires the correct password to retrieve the data in a decrypted format. If an incorrect password is provided, DB2 does not decrypt the data. The ENCRYPT keyword encrypts data. The DECRYPT_BIT, DECRYPT_CHAR, and DECRYPT_DB keywords decrypt data. These functions work like other built-in functions. Note that built-in encryption functions work for data that is stored within DB2 subsystem and is retrieved from within that same DB2 subsystem. The encryption functions do not work for data that is passed into and out of a DB2 subsystem. This task is handled by DRDA® data encryption, and it is separate from built-in data encryption functions

DB2 can protect sensitive data from priveleged users

  - The latest release of DB2 can protect sensitive data from privileged users. It does this by providing a new SYSADM authority without data access. There is also a separate authority to perform security related tasks
  - Also, DB2 for z/OS allows the EXPLAIN query plan facility to work without execute privilege or ability to access data

- Data Encryption for IMS and DB2 Databases runs as an exit. The exit code invokes the System z, Crypto hardware to encrypt data for storage and decrypt data for application use, thereby protecting sensitive data residing on various storage media. Data Encryption for IMS and DB2 Databases is implemented by using the IMS Segment Edit/Compression exit and the DB2 EDITPROC

- InfoSphere Business Glossary is part of the InfoSphere Foundation Tools. This product allows business users to define information to be governed. Data is defined using a common set of data names, definitions and integrity constraints sometimes referred to as a shared business vocabulary (SBV). Once data is defined here, it becomes possible to look in the business glossary to highlight what information needs to be protected

InfoSphere Discovery can be used to identify information to be protected across the enterprise

- InfoSphere Discovery is also part of the InfoSphere Foundation Tools suite of products. Using InfoSphere Discovery it becomes possible to automatically discover where in the enterprise that data is located that needs to be protected. InfoSphere Discovery uses a cross-profiler, unified schema builder and transformation analyzer to finding complete business data entities (e.g. customer, product, order) as described in Figure 4. This includes the discovery of sensitive data even if it has been transformed in some way. InfoSphere Discovery can also pass metadata to other tools running on the IBM System z stack as discussed in Figure 5. This capability makes it potentially possible to implement information protection at the data entity level irrespective of where the data is located. It is this capability that offers enormous productivity gains in managing information protection across the enterprise. For example, customer data privacy, customer data retention, customer data access security, and customer data archive are all potentially possible

- Looking at Figure 5, IBM Optim Data Privacy can receive metadata from InfoSphere Discovery and then apply data privacy policies to mask discovered sensitive data that needs to be protected across the enterprise. This includes masking test data.

IBM Optim can mask sensitive data to protect privacy

In addition IBM Optim can also manage data retention by archiving discovered data at the data entity level. For example it becomes possible to archive orders data that is over three years old. This capability is particularly useful when data has to be retained for compliance reasons. IBM Optim allows this type of requirement to be easily managed.

Protecting sensitive data in unstructured documents and forms can also be managed

- In addition to protecting structured data, IBM Optim Data Redaction is new product that offers the ability to protect sensitive data buried in documents and forms. This product finds and removes sensitive data and metadata from documents and therefore reduce the cost of compliance

- Finally there is InfoSphere Guardium. This product creates a continuous, fine-grained audit trail of all database activities, including the "who, what, when, where, and how" of each transaction. InfoSphere Guardium uses agent technology to monitor and audit database activity in real-time across a range of DBMS products that include:
  - Oracle
  - Microsoft SQL Server
  - IBM DB2 (Windows, Unix, z/Linux)
  - IBM DB2 for z/OS
  - IBM DB2 for iSeries (AS/400)
  - IBM Informix
  - MySQL
  - Sybase ASE
  - Sybase IQ
  - Teradata

InfoSphere Guardium can monitor and audit database activity across heterogeneous databases in real-time

Guardium provides the capability to proactively identify unauthorized or suspicious activities by continuously tracking all database actions. In addition malicious or unapproved activity by DBAs, developers and outsourced personnel can be detected or blocked without relying on native logs, triggers or other DBMS-resident mechanisms. Pre-configured reports and automated workflows (electronic sign-offs, escalations, etc.) are also available. These make it possible to simplify compliance processes and continuously monitor information protection.

CONCLUSIONS

Information protection is currently implemented in a fractured way in most organisations

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. Yet the increasing threat of internet fraud and the mounting pressures brought about by stricter regulatory compliance and risk management has led many companies to start looking for integrating end-to-end solutions to solve this problem.   The challenge is to protect information no matter where it resides or flows to in the enterprise.

While no vendor in the market can claim they have a complete end-to-end fully integrated information protection solution, IBM zEnterprise platform, goes a long way to offering one. It has components available to define information to be governed, to highlight what needs to be protected, to discover where that data is in the enterprise and to apply policies to protect it across both structured and unstructured information. Hardware encryption allows high performance encryption to be applied while data is in motion and database encryption allows it to be applied while data is at rest.

IBM System z removes much of the risk associated with sensative data

The integration of InfoSphere Discovery with Optim Data Privacy combined with Optim Data Redaction and InfoSphere Guardium on the IBM System Z platform offers very comprehensive information protection and removes much of the risk associated with sensitive data. IBM Optim also allows data across the enterprise to be archived while protecting integrity as part of a data retention program as data grows. Application retirement can also be managed in this way.

In addition IBM System z offers integrated authentication and authorisation in terms of managing access control as part of an enterprise single sign-on program and can monitor and audit all database activity across the enterprise in real-time to enforce policy.

Information confidentiality, information integrity and information availability are all protected

All in, the IBM System z stack offers integrated technologies that help protect information confidentiality, information integrity and information availability and is well worth considering for any company looking for solution to information risk management.

ABOUT INTELLIGENT BUSINESS STRATEGIES

Today, successful companies are those that can absorb new information technologies and use them effectively in their businesses. But faced with so many new technology developments how can IT and business users possibly keep up? Intelligent Business Strategies is an IT research and consulting company whose goal is to help companies understand and exploit new developments in business intelligence, analytical processing, data management and enterprise business integration. Together, these technologies help an organisation become an *intelligent business*.

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 29 years of IT experience, Mike has consulted for dozens of companies on business intelligence, data management and enterprise business integration. He has spoken at events all over the world and written numerous articles. Mike is a resident expert on the B-EYE-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 and European Managing Director of Database Associates. He teaches popular master classes in Operational Business Intelligence and Performance Management, Enterprise Data Governance, Master Data Management and Enterprise Business Integration

INTELLIGENT
BUSINESS
STRATEGIES

Intelligent Business Strategies
2nd Floor, Springfield House
Water Lane, Wilmslow
Cheshire SK9 5BG
England
Telephone: (+44)-1625-520700

Internet URL: www.intelligentbusiness.biz
E-mail: mferguson@intelligentbusiness.biz

Information Protection on IBM System z