Title: The Impact of Regulatory Compliance on Database Administration

Date: 05/07/2007

Time: 9:50 AM - 10:50 AM

Code: I01

Authors

This presentation was prepared by:

Craig S. Mullins

Corporate Technologist

NEON Enterprise Software, Inc.
14100 Southwest Freeway, Suite 400
Sugar Land, TX 77478
Tel: 281.491.6366
Fax: 281.207.4973
E-mail: craig.mullins@neonesoft.com

These are the bullets from the abstract:

* What are the Regulations?
* Compliance Issues
* IT Controls
* Long-term Data Retention
* Database Auditing
* Additional DBA Concerns

Regulatory Compliance

- GLB
- HIPAA
- Basel II
- Federal Information Security Management Act
- Sarbanes-Oxley

GLB: Gramm-Leach-Bliley Act

- The Gramm-Leach-Bliley Act (GLB Act), also known as the Financial Modernization Act of 1999, is a federal law enacted in the United States to control the ways that financial institutions deal with the private information of individuals.
- The Act consists of three sections:
  - The Financial Privacy Rule, which regulates the collection and disclosure of private financial information;
  - The Safeguards Rule, which stipulates that financial institutions must implement security programs to protect such information;
  - and the Pretexting provisions, which prohibit the practice of pretexting (accessing private information using false pretenses).
- The Act also requires financial institutions to give customers written privacy notices that explain their information-sharing practices.

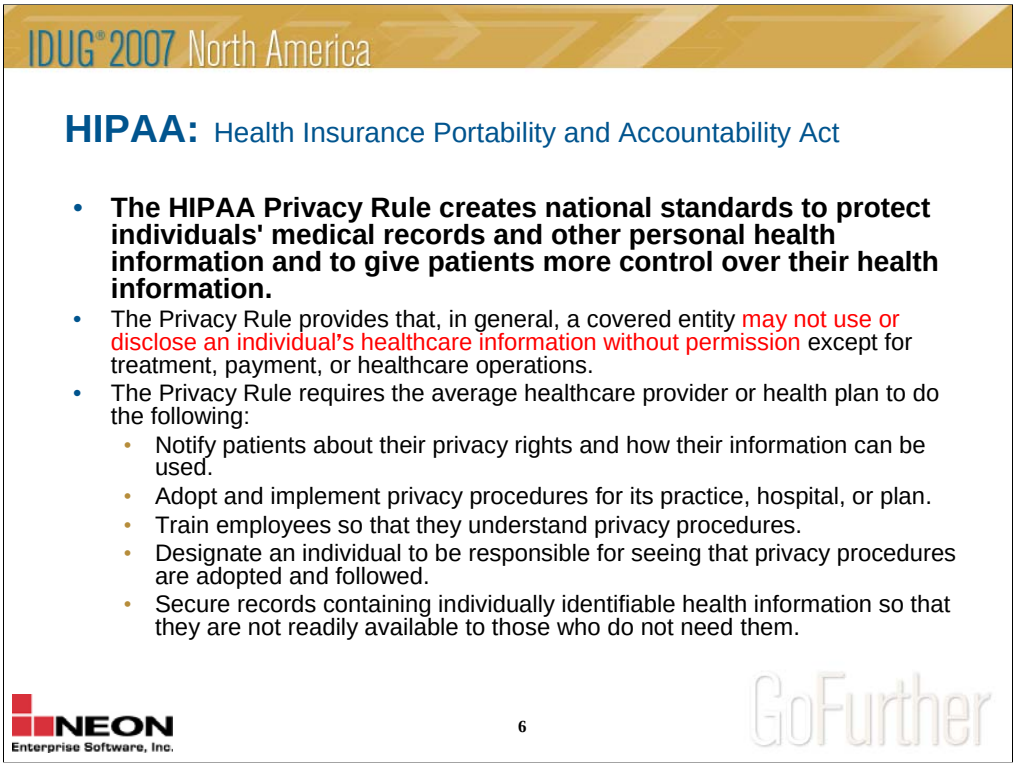

Violating the HIPAA privacy rules, 45 C.F.R. § 164.512(i) & § 164.514(c), -Research records to be shared where they are "de-identified" up to the standards of Section 164. -Where the sharing of records involves only a "limited data set." To qualify as a limited data set, "direct identifiers" such as name and SSNs must be eliminated.

An earlier accord, Basel I, adopted in 1988, is now widely viewed as outmoded as it is risk insensitive and can easily be circumvented by regulatory arbitrage.

The Basel II deliberations began in January 2001, driven largely by concern about the arbitrage issues that develop when regulatory capital requirements diverge from accurate economic capital calculations.

FISMA: Federal Information Security Mgmt Act

- The E-Government Act was passed in 2002 as a response to terrorist threats
  - Title III of the act is named the Federal Information Security Management Act (FISMA).
  - FISMA basically states that federal agencies, contractors, and any entity that supports them, must maintain security commensurate with potential risk.
  - Officials are graded on the potential effect a security breach would have on their operations.

SOX: Sarbanes-Oxley Act

- **The U.S. Public Accounting Reform and Investor Protection Act of 2002**
  - "…to use the full authority of the government to expose corruption, punish wrongdoers, and defend the rights & interests of American workers & investors."
  - Impacts auditors, corporations, and IT
    - Public companies with a market capitalization of at least $75M
    - Companies listed on a United States exchange even if they are incorporated outside of the United States
  - **The primary objectives of SOX:**
    - To **strengthen and restore public confidence** in corporate accountability and the accounting profession;
    - To **strengthen enforcement** of the federal securities laws;
    - To **improve** executive **responsibility**;
    - To **improve disclosure** and financial reporting; and
    - To **improve** the **performance of "gatekeepers."**

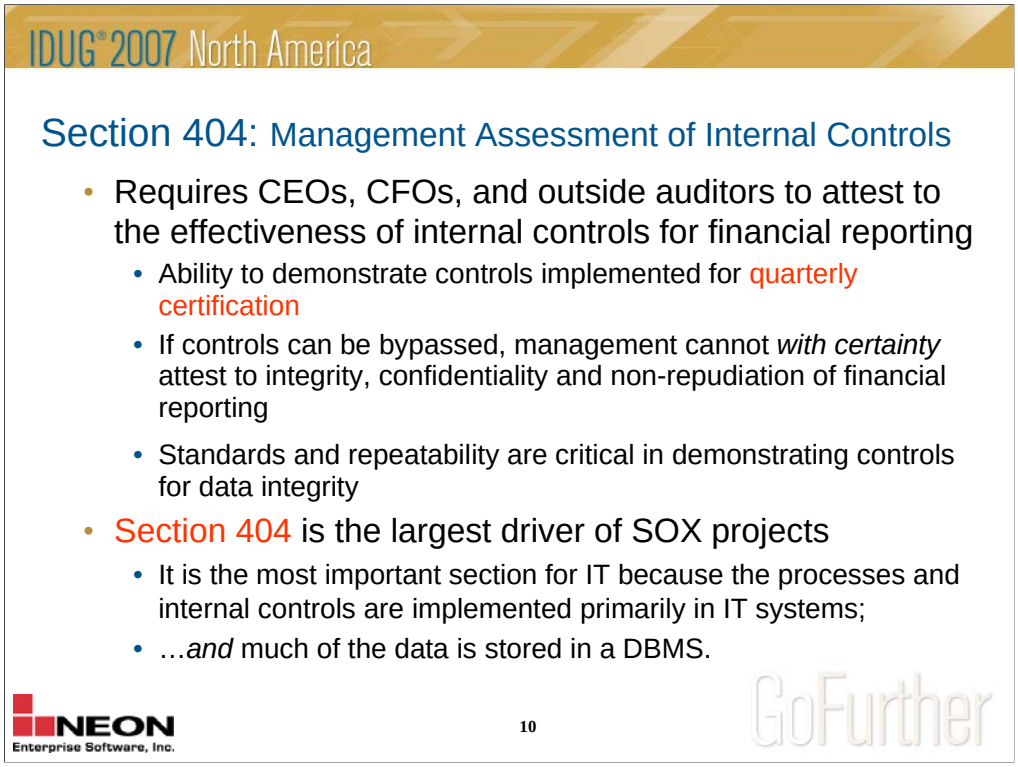

Definition of control/control activity

Safeguards or processes that mitigate risk

Processes effected by people designed to accomplish specified objectives (COSO)

Infrastructure, and other components maintain confidentiality, integrity, availability

Company Management

Assures evaluation of controls effectiveness, provides written assessment

Accepts responsibility for it; supports audit evaluation with evidence

Compliance With Sarbanes-Oxley

404 requires external auditor's opinion on effectiveness of internal controls

Ability to demonstrate controls implemented for quarterly certification

If controls can be bypassed, management cannot with certainty attest to integrity, confidentiality and non-repudiation of financial reporting

Standards and repeatability critical in demonstrating controls for data integrity

Other Regulations & Issues?

- And, there are more regulations to consider, for example:
  - the USA Patriot Act
  - Can SPAM Act of 2003
  - Telecommunications Act of 1996
  - The Data Quality Act
- And, there are more regulations to contend with; based upon your industry, location, etc.
- As well as, new regulations that will continue to be written by government and imposed over time…
- Regulations have brought to light some of the personal financial information that has been compromised (stolen)
  - More on this later…

Controlling the Assault of Non-Solicited Pornography and Marketing (Can SPAM). The law took effect on January 1, 2004. The Can Spam Act allows courts to set damages of up to $2 million when spammers break the law. Federal district courts are allowed to send spammers to jail and/or triple the damages if the violation is found to be willful. The Can Spam Act requires that businesses:

•Clearly label commercial e-mail as advertising

•Use a truthful and relevant subject line

•Use a legitimate return e-mail address

•Provide a valid physical address

•Provide a working opt-out option

•Process opt-out requests within ten business days

The Telecommunications Act of 1996, enacted by the U.S. Congress on February 1, 1996, and signed into law by President Bill Clinton on February 8, 1996, provided major changes in laws affecting cable TV, telecommunications, and the Internet. The law's main purpose was to stimulate competition in telecommunication services. The law specifies:

•How local telephone carriers can compete

•How and under what circumstances local exchange carriers (LEC) can provide long-distance services

•The deregulation of cable TV rates

Regulatory Compliance and…

- **Impact**: upper-level management is keenly aware of the need to comply, if not all of the details that involves.
- **Prosecution**: being successfully prosecuted *(see next slide)* can result in huge fines and even imprisonment.
- **Cost**: the cost of complete compliance can be significant.
- **Durability**: although there have been discussions about scaling back some laws (e.g. SOX), increasing regulations and therefore increasing time, effort, and capital will be spent on compliance.
  - *That is, the issue will not just disappear if you ignore it long enough!*
- **But**, at the end of the day, ensuring exact compliance is a gray area.

The Reality of Prosecution

- **Enron** – In May 25, 2006 Ken Lay (former CEO) was found guilty of 10 counts against him. Because each count carried a maximum 5- to 10-year sentence, Lay could have faced 20 to 30 years in prison. However, he died while vacationing in about three and a half months before his scheduled sentencing. Jeff Skilling (another former CEO) was found guilty of 19 out of 28 counts against him, including one count of conspiracy, one count of insider trading. Skilling was sentenced to 24 years and 4 months in federal prison.
- **WorldCom** – In June 2005, 800,000 investors were awarded $6 billion in settlements; the payouts will be funded by the defendants in the case, including investment banks, audit firms, and the former directors of WolrdCom. The judge in the case noted that the settlements were "of historic proportions." Additionally, Bernard Ebbers, former CEO of WorldCom, was convicted of fraud and sentenced to 25 years in prison.
- **Tyco** – In September 2005, Dennis Kozlowski (former Tyco CEO) and Mark Swartz (former Tyco CFO) were sentenced to 8 to 25 years in prison for stealing hundreds of millions of dollars from Tyco. Additionally, Kozlowski had to pay $70 million in restitution; Swartz $35 million.
- **Adelphia** – In June 2005, John Rigas (founder and former CEO) was sentenced to 15 years in prison; his son, Tony, was also convicted of bank fraud, securities fraud, & conspiracy - he received a 20 year sentence.
- **HealthSouth** – In March 2005, Richard Scrushy, founder and former CEO, was acquitted of charges relating to 1 $2.7 billion earnings over-statement. Scrushy blamed his subordinates for the fraud…

So What Do We Need to Do?

- Implement standard controls and methods for improving:
  - Data Quality
  - Long-term Data Retention
  - Database Security
  - Database and Data Access Auditing
  - DBA Procedures
- …and this will require proper metadata management!

Compliance Is About Methods and Controls

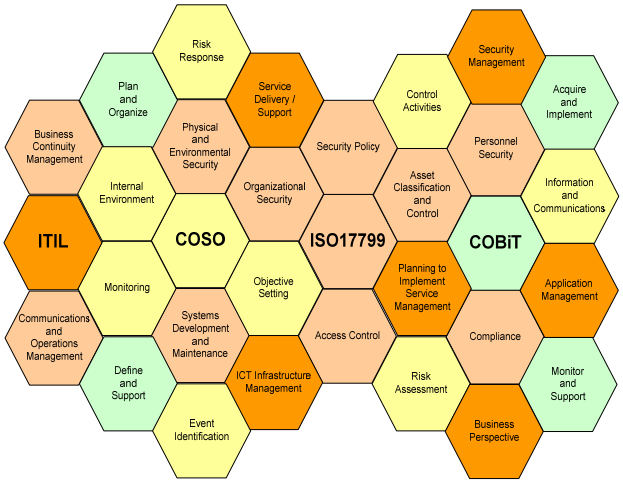

Compliance with most regulations, including Sarbanes-Oxley, is not a software issue alone but also a process issue...

A *purely* software equivalent of a silver bullet for regulatory compliance does not exist.

NEON Enterprise Software, Inc.

15

GoFurther

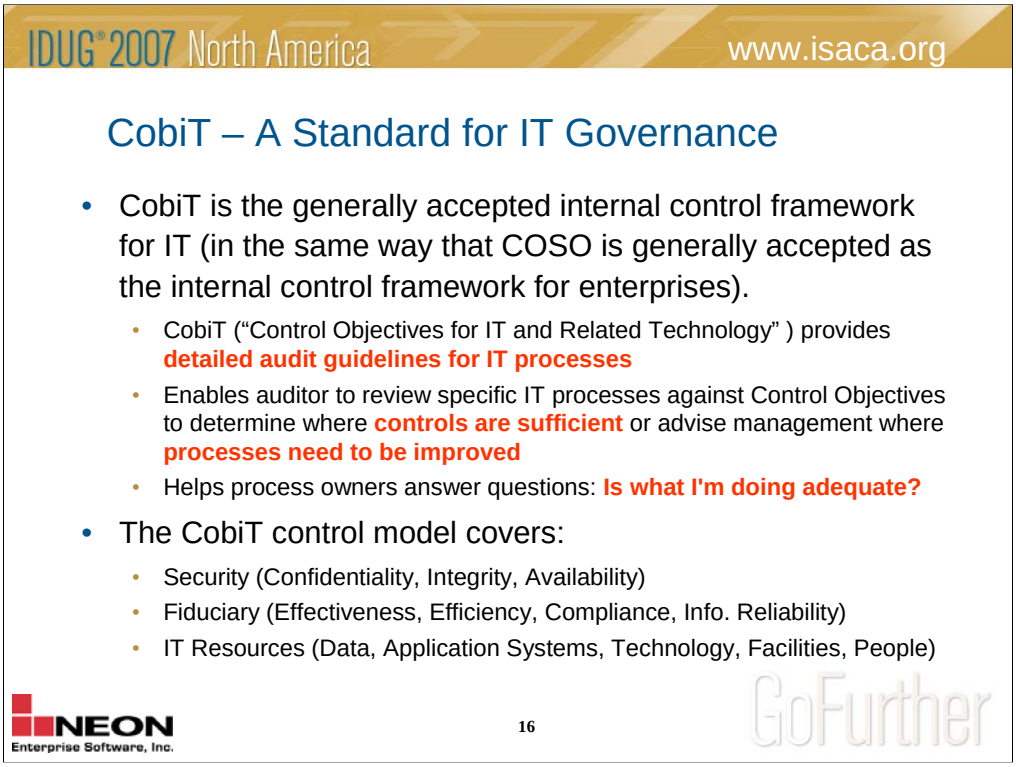

Auditors typically use standards such as CobiT for IT controls and objectives

•CobiT approaches IT controls looking at information and data that supports business requirements and associated IT resources and processes

•Don't create a non-standard approach, when you can leverage something known by an auditor

Issue #1: Data Quality

- A recent SAS/Risk Waters Group survey indicated that 93% of respondents had experienced losses of $10 million in one year…
  - And 21% of respondents said that at some point, their company suffered a loss between $10,000 and $1,000,000 in a single day.
- The prime reasons given for such losses were *incomplete, inaccurate or obsolete data*, and inadequate processes.

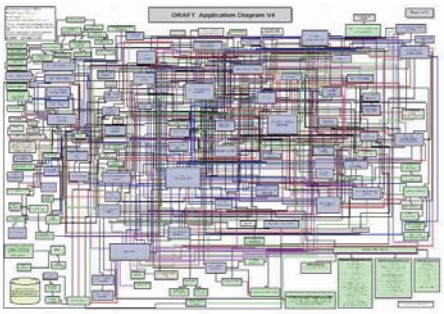

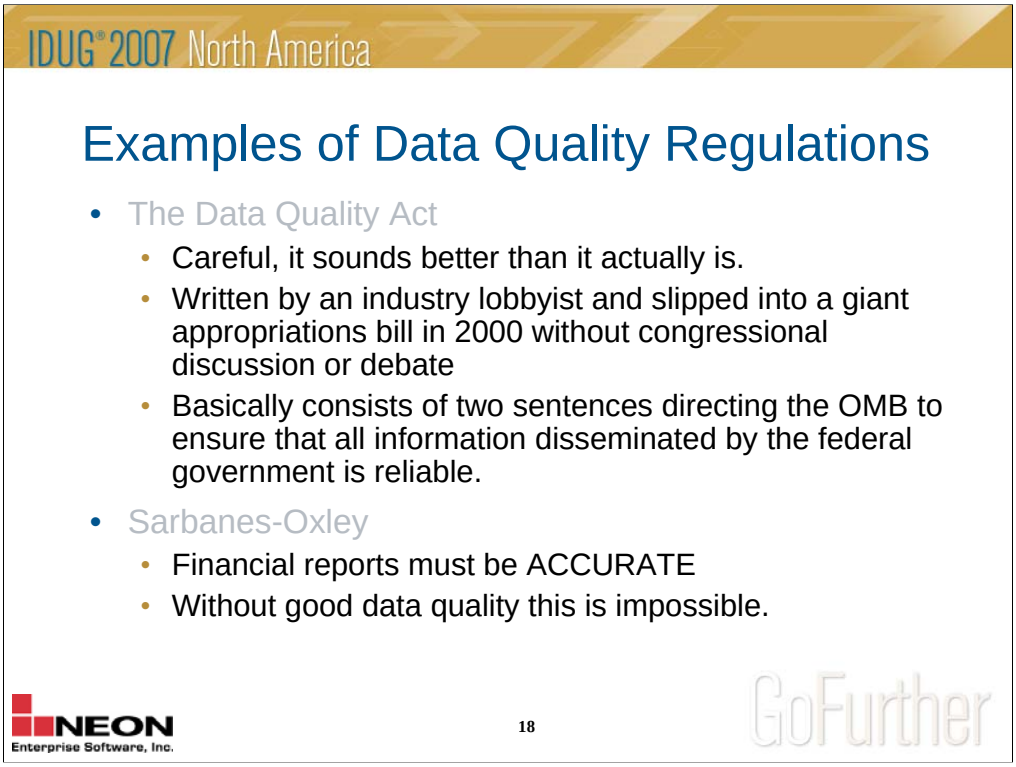

Since its inception, the Data Quality Act has been under attack as a weapon of big business, a stealthy way to keep federal agencies tied in knots over what constitutes sound science

But the Bush administration's interpretation of those two sentences could tip the balance in regulatory disputes that weigh the interests of consumers and businesses.

Data Quality: Is it Really a Major Concern?

- How good is your data quality?
  - The cost of poor quality is usually hidden and not obvious to those not looking for it.

    Source: Jack Olsen, *Data Quality: The Accuracy Dimension*, (Morgan Kaufmann).

- Estimates show that, on average, data quality is suspect:
  - Payroll record changes have a 1% error rate;
  - Billing records have a 2-7% error rate, and;
  - The error rate for credit records: as high as 30%.

    Source: Thomas C. Redman, *Data Quality: Management and Technology*, (Bantam Books).

- Similar studies in Computer World and the Wall Street Journal back up the notion of overall poor data quality.
  - W.M. Bulkeley, "Databases Are Plagued by Reign of Error," The Wall Street Journal, 26 May 1992, B2.
  - B. Knight, "The Data Pollution Problem," ComputerWorld, 28 September 1992, 81-84.

The High Cost of Poor Quality Data

- "Poor data quality costs the typical company at least ten percent (10%) of revenue; twenty percent (20%) is probably a better estimate."

Source: Thomas C. Redman, "Data: An Unfolding Quality Disaster", *DMReview Magazine*, August 2004

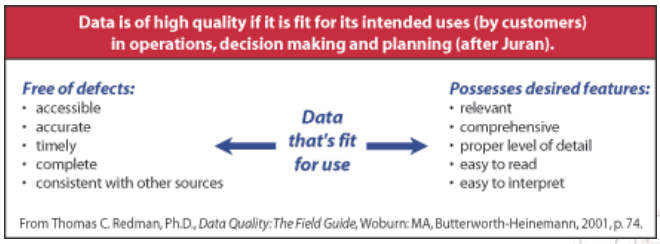

NEON Enterprise Software, Inc.

20

GOFurther

Juran, Joseph M. and A. Blanton Godfrey, Juran's Quality Handbook, Fifth Edition, p. 2.2, McGraw-Hill, 1999.

Database Constraints

- By building constraints into the database, overall data quality may be improved:
  - Referential Integrity
    - Primary Key
    - Foreign Key(s)
  - UNIQUE constraints
  - CHECK constraints
  - Triggers
  - *Domains*

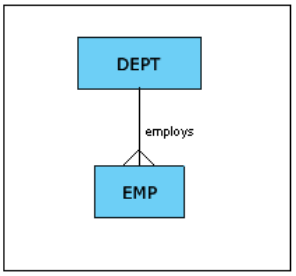

Data Profiling and Mapping

- With data profiling, you can:
  - Discover the quality, characteristics and potential problems of information before beginning data-driven projects
  - Reduce the time and resources required to find problematic data
  - Allow business analysts and data stewards to have more control on the maintenance and management of enterprise data
  - Catalog and analyze metadata and discover metadata relationships

Issue #2: Long-Term Data Retention

- The average installed storage capacity at Fortune 1000 corporations has grown from 198TB to 680TB in less than two years.
- This is a growth rate of more than 340% as capacity continues to double every 10 months.

Source: TIP (TheInfoPro), 2006 www.theinfopro.net

NEON Enterprise Software, Inc.

23

GoFurther

Demand for data archiving will increase. Exploding data volumes and long-term data retention policies driven by regulatory requirements are creating the need for an expanded data archiving solution. Third-party vendors will continue to offer enhanced features to support more packaged applications, in addition to improved integration and automation of archiving solutions.

Source: Trends 2006: Database Management Systems, Forrester Research

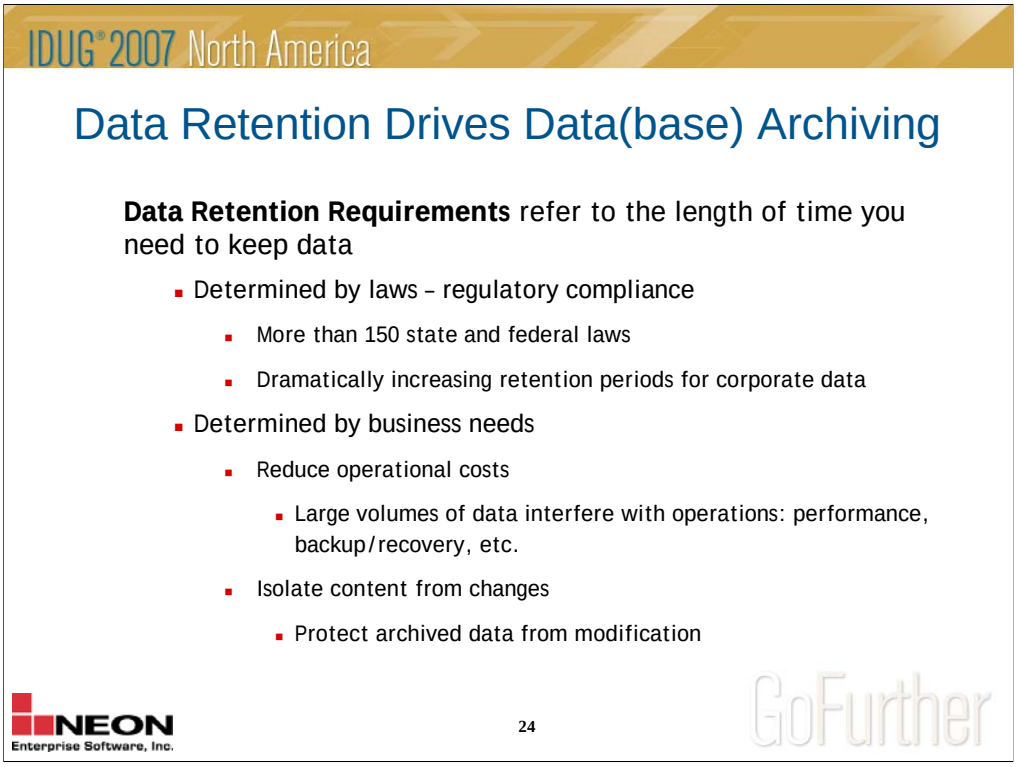

OK, now let's take a closer look at the requirements for data retention. First thing to keep in mind is that laws NEVER say you have to archive data, just that you retain data. So if you can keep it in the operational database, OK. But if not, you have to archive it.

<u>Data Archiving</u> is a process used to move data from the operational database to another data store to be kept for the duration of the retention period when it is unacceptable to keep the data in the operational database for that long. So, why might it be unacceptable?

- large volumes of data interfering with operations (the more data in the operational database, the slower all processes may run)

- need for better protection from modification

- need for isolation of content from changes

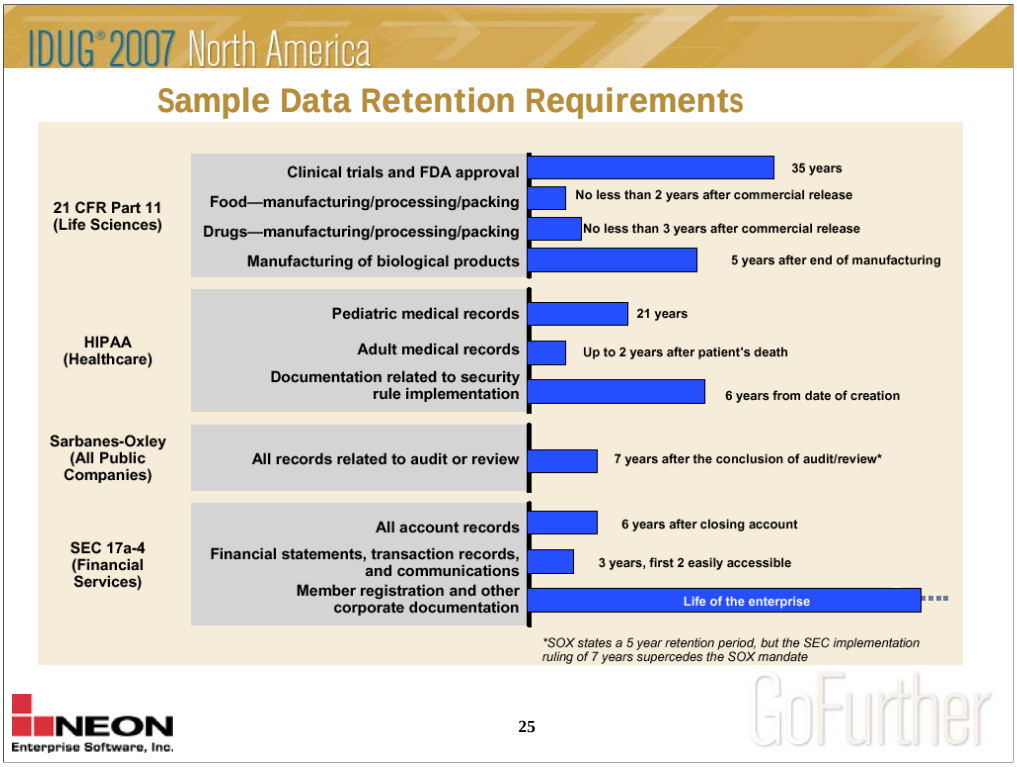

Here we have some examples of regulations that impact data retention.

This is just a sampling of the more than 150 different regulations (at the local, state, national, and international levels) that impact data retention.

Regulations Drive Retention and Discovery

Source: www.domains.cio.com/symantec/wp/ebs_**ediscovery**_ev_wp.pdf

E-Discovery

- Electronic evidence is the predominant form of discovery today. (Gartner,Research Note G00136366)
- Electronic evidence could encompass anything that is stored anywhere. (Gartner,Research Note G00133224)
- When data is being collected (for e-discovery) it is imperative that it is not changed in any way. Metadata must be preserved… (Gartner,Research Note G00133224)
- Gartner Strategic Planning Assumption: Through 2007, more than half of IT organizations and in-house legal departments will lack the people and the appropriate skills to handle electronic discovery requirements (0.8 probability). (Gartner,Research Note G00131014)

According to Gartner: As the custodians of corporate information systems, IT departments must understand the legal ramifications surrounding the information stored in those systems. Understanding the legal process of discovery as it pertains to electronically stored information is essential.

E-Discovery and the FRCP

- Federal Rules of Civil Procedure, Rule 34b
  - *Took effect December 2006*
  - A party who produces documents for inspection shall produce them . . . as they are kept in the usual course of business..."
  - The amended rules state that requested information must be turned over within 120 days after a complaint has been served.
- So data stored in database systems must be able to be produced in electronic form.

NEON Enterprise Software, Inc.

27

GoFurther

The regulations dictate the need to retain the data, legal discovery rules dictate how that data must be produced in a court of law. The two work hand-in-hand.

Today, most discovery requests are for electronic data, not paper documents.

The FRCP changes which took effect in December impact how retained data must be produced for legal purposes.

- Most organizations are not prepared to handle an e-discovery request. Requires marriage of IT/business/legal.

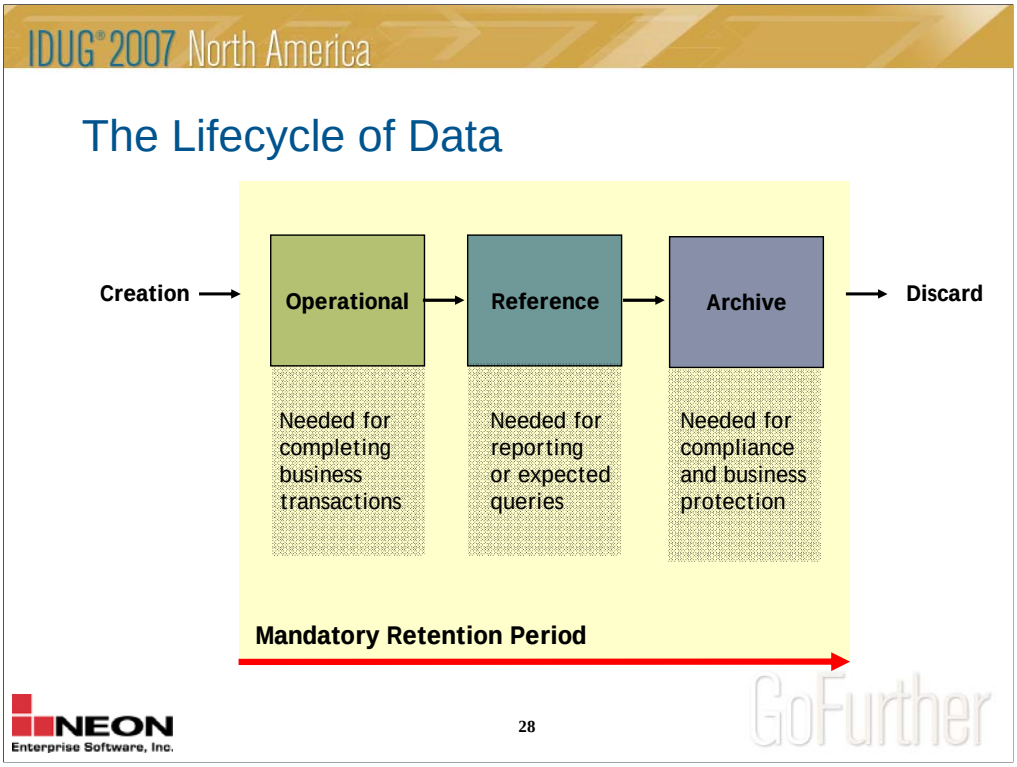

Here is what we like to call the information life cycle. You may recognize these steps as something else in your organization, but the basic steps are common.

This slide delineates the various states of data over its useful life. This could be three separate databases or a single database or any combination thereof.

Don't think about data warehousing in this context – here we are talking about the single, official store of data.

We start out, of course, with creation of data.

After creating the data it moves into its operational state. This is where it serves it primary business purpose. Transactions are enacted upon data in this state.

Then we move to the reference state. Data in this form is still needed for reporting and querying purposes.

After this state we move into an area where the data is no longer needed for completing business transactions and the chance of it being needed for querying and reporting is small to none. However, the data still needs to be saved for regulatory compliance and other legal purposes, particularly if it pertains to a financial transaction.

Finally, after a period of time, the data is no longer needed at all and it can be discarded. This actually should have a much stronger emphasis: the data must be discarded. In many cases the only reason old data is being kept is to enable lawsuits. When there is no legal requirement to maintain such data, companies demand that such data be destroyed – why enable anyone to sue you if it is not a legal requirement?

What is the "Solution?"

- **Keep Data in Operational Database?**
- **Store Data in UNLOAD files** (or backups)**?**
- **Move Data to a Parallel Reference Database?**
- **Move Data to a Database Archive!**

OK, so if we have to archive data, how can we go about doing it?

Well, there is always bullet number 1. If you can store it in the operational database then you save a lot of work. But it does not work well for large amounts of data or very long retention periods. And the more data in the database the less efficient transaction processing and reporting will be.

So we could take a simple approach of using UNLOAD files. But this is not really useful in today's environment. It assumes the same data schema will be used if you want to ever RELOAD the data (and that is not likely over long periods of time). When they run into audit situations, where they're asked to produce transaction records spanning multiple years, they have to slog through the work of restoring the data. This is a very labor-intensive, manual process. And costly… some sites have spent hundreds of thousands of dollars to respond to audit requests when they have to resort to such manual data restoration. It might be feasible for 2 to 3-year timespans, but not 20 or more years.

In situation number 3, we consider cloning and moving data. But this has a lot of the same problems as UNLOADing the data. Think about what happens as the database schema changes – and how you'd go about accessing the data?

Finally, we come to the proper approach for retaining large amounts of data for long periods of time – database archiving.

Components of a Database Archiving Solution

Database Archiving: The process of removing selected data records from operational databases that are not expected to be referenced again and storing them in an archive data store where they can be retrieved if needed.

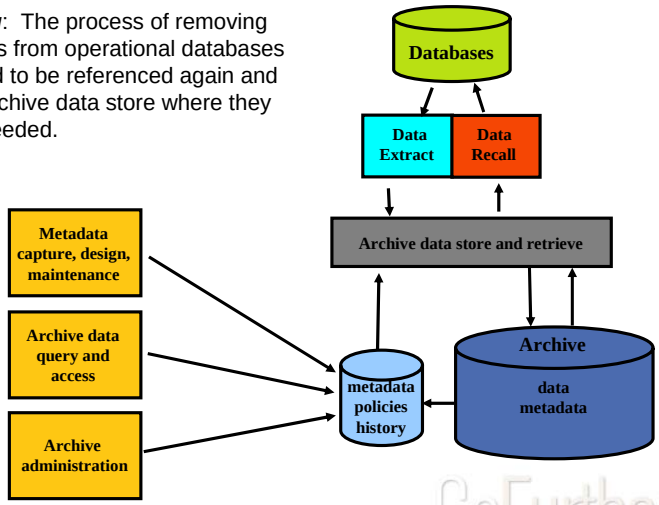

NEON Enterprise Software, Inc.

30

GoFurther

<u>Database Archiving</u> is part of a larger topic, namely Data Archiving. There are many types of data that needs to be archived to fulfill regulatory, legal, and business requirements. Perhaps the biggest driver for archiving has been e-mail. At any rate, each type of data requires different archival processing requirements due to its form and nature. What works to archive e-mail is not sufficient for archiving database data, and so on.

This diagram depicts the necessary components of a database archiving solution. Starting with the databases down the left side is the Extract portion, and up the right side is the data recall portion. Several sites have been very vocal about needing this capability, but I think it is a dubious requirement. If you can query from the archive recalling data into the operational database is not really a necessity, it it? And what if the operational database changes (e.g. Oracle to SAP)?

The whole process requires metadata to operate. You must capture, validate and enhance the metadata to drive the archive process. You need to know the structure of the operational database and the structure of the archive.

There is also the case of the metadata about data retention for the archive. This policy-based metadata must be maintained and monitored against the archive, for example, to determine if data needs to be discarded.

And, as you can see off to the bottom left we also need to be able to query the archived data. This will not necessarily be the most efficient access because of differences in the metadata over time, but for queries against archived data performance is not a paramount concern.

Finally, we will also have on-going maintenance of the archive: security, access audit, administration of the structures (REORG), backup/recovery, etc.

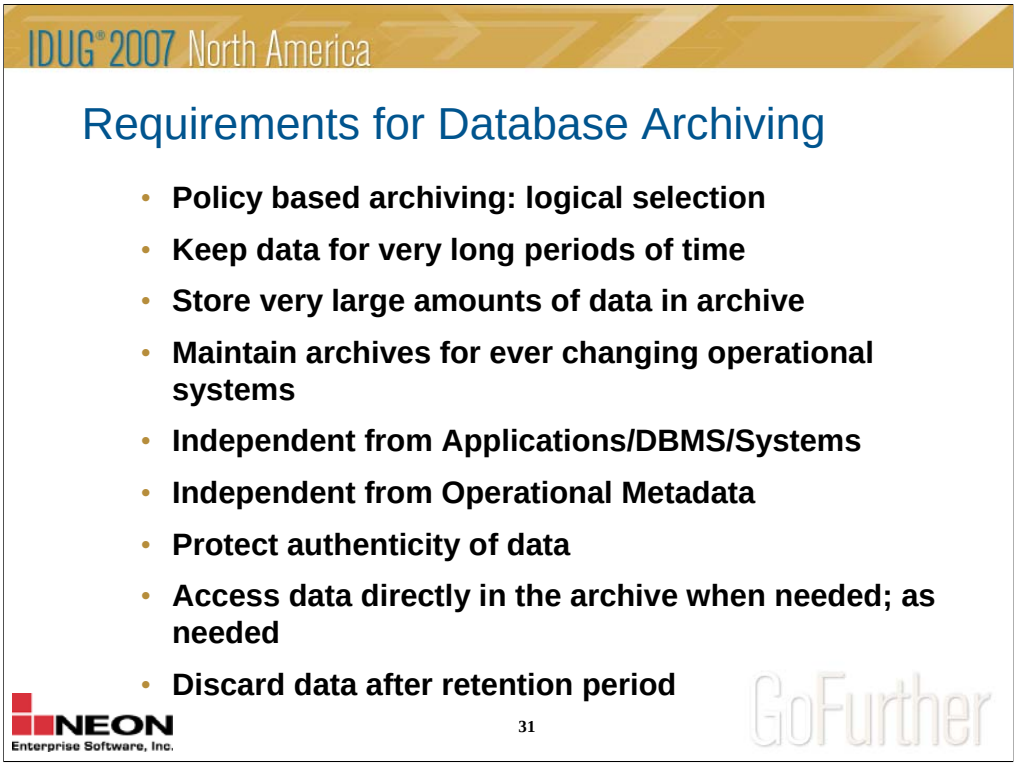

OK, so let's take a look at the actual functionality needed to support database archiving. This slide outlines the various features and functions needed:

Additional summary points:

Keeping data in operational systems is a bad idea, as is putting data in UNLOAD or backup files, putting data in a parallel references database, and/or using a DBMS to store the archive does not work

Database archiving requires a great deal of data design

Establishing and maintaining metadata

Designing how data looks in the archive

Achieving application independence

Database archives must be continuously managed

Copying data for storage problems (e.g. media rot)

Copying data for system changes

Copying data for data encoding standard changes

Logging, auditing, and monitoring

Archive events

Partition management

Accesses

New requirement: Database Archivist staff

Issue #3: Data and Database Security

- **75% of enterprises do not have a DBMS security plan.**

  Source: Forrester Research, "Your DBMS Strategy 2005"

- **A single intrusion that compromises private data can cause immense damage to the reputation of an enterprise — and in some cases financial damage as well.**

  Source: Forrester Research, "Comprehensive Database Security…"

- **Database configurations are not secure by default, and they often require specialized knowledge and extra effort to ensure that configuration options are set in the most secure manner possible.**

  Source: Forrester Research, "Comprehensive Database Security…"

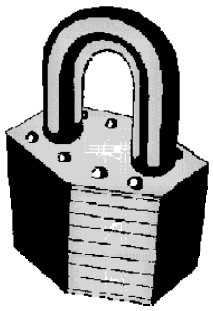

NEON Enterprise Software, Inc.

32

GoFurther

The Forrester Research report from which these clips are taken is titled:

"Comprehensive Database Security Requires Native DBMS Features And Third-Party Tools"

Published: March 29, 2005

Obstacles to Database Security

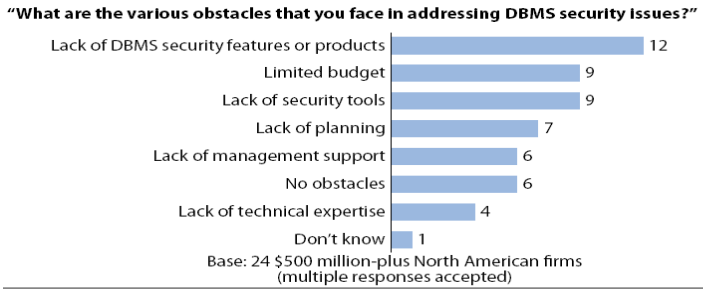

Source: Forrester Research,
"Comprehensive Database Security…"

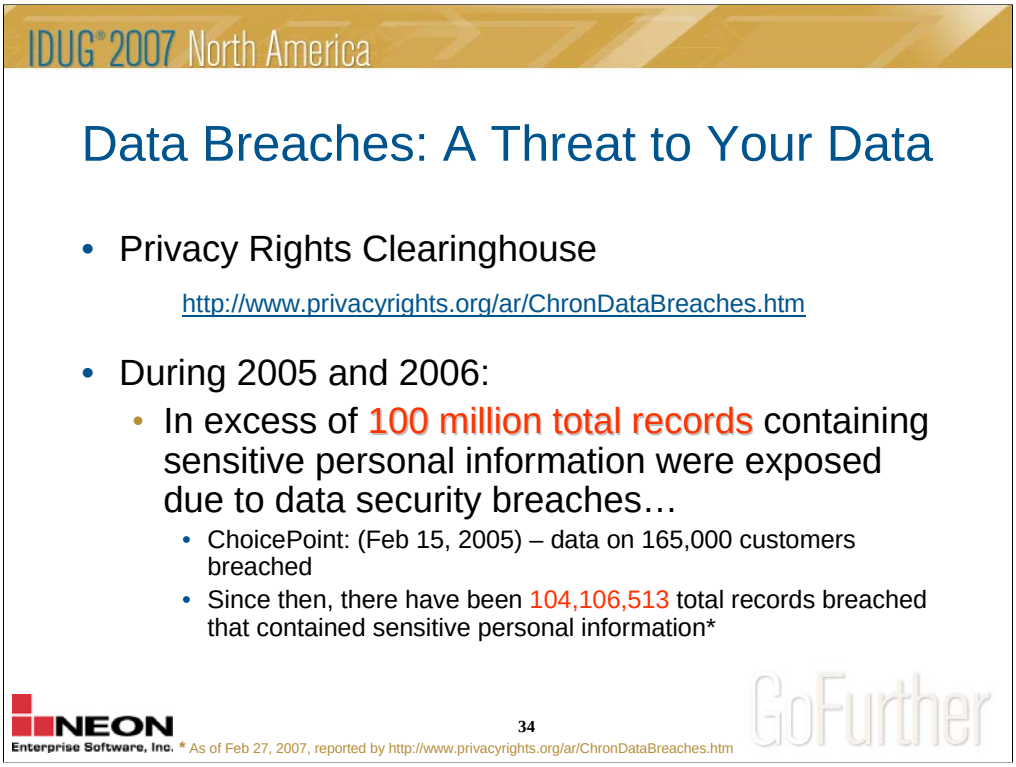

The Privacy Rights Clearinghouse began keeping records on February 15, 2005 – the date of the ChoicePoint breach.

104,106,513 total

Database Security Issues in a Nutshell

- Authentication
  - Who is it?
- Authorization
  - Who can do it?
- Encryption
  - Who can see it?
- Audit
  - Who did it?

Database Security Issues

- Who has access to high-level “roles”:
  - Install SYSADM, SYSADM, SYSCTRL, DBADM, etc.
- Auditors will have issues with this, but it is difficult, if not impossible, to remove.
- Do any applications require DBA-level authority? Why?
- Security monitoring
  - Audit from the log; from an appliance.
  - Additional tools

SYSADM and Install SYSADM

- Install SYSADM bypasses the DB2 Catalog when checking for privileges.
  - So, there are no limits on what a user with Install SYSADM authority can do.
  - And it can only be removed by changing DSNZPARMs
  - Basically, needed for catalog and system "stuff"
- SYSADM is *almost* as powerful, but:
  - It can be revoked
  - Biggest problem for auditors: SYSADM has access to all data in all tables.

One or two IDs are assigned installation SYSADM authority when DB2 is installed. They have all the privileges of SYSADM, plus:

•DB2 does not record the authority in the catalog. Therefore, the catalog does not need to be available to check installation SYSADM authority. (The authority outside of the catalog is crucial. For example, if the directory table space DBD01 or the catalog table space SYSDBAUT is stopped, DB2 might not be able to check the authority to start it again. Only an installation SYSADM can start it.)

•No ID can revoke installation SYSADM authority. You can remove the authority only by changing the module that contains the subsystem initialization parameters (typically DSNZPARM).

IDs with installation SYSADM authority can also perform the following actions:

•Run the CATMAINT utility

•Access DB2 when the subsystem is started with ACCESS(MAINT)

•Start databases DSNDB01 and DSNDB06 when they are stopped or in restricted status

•Run the DIAGNOSE utility with the WAIT statement

•Start and stop the database that contains the ART and the ORT

Suggestions

- Limit the number of SYSADMs
  - And audit everything those users do.
- Consider associating Install SYSADM with a RACF group
  - Authids that absolutely need Install SYSADM can be connected to the RACF group using secondary authids.
  - *similar issues with SYSOPR and Install SYSOPR*

Data Encryption Considerations

- California's SB 1386 protects personally identifiable information; obviously it doesn't matter if encrypted data becomes public since it's near impossible to decrypt.

Source: "Encryption May Help Regulatory Compliance"
Edmund X. DeJesus, SearchSecurity.com

- Types of encryption
  - At Rest
  - In Transit
- Issues
  - Performance
    - Encrypting and decrypting data consumes CPU
  - Applications *may* need to be changed
    - See next slide for DB2 V8 encryption functions

URL for the sourced quote above:
http://searchsecurity.techtarget.com/originalContent/0,289142,sid14_gci991508,00.html

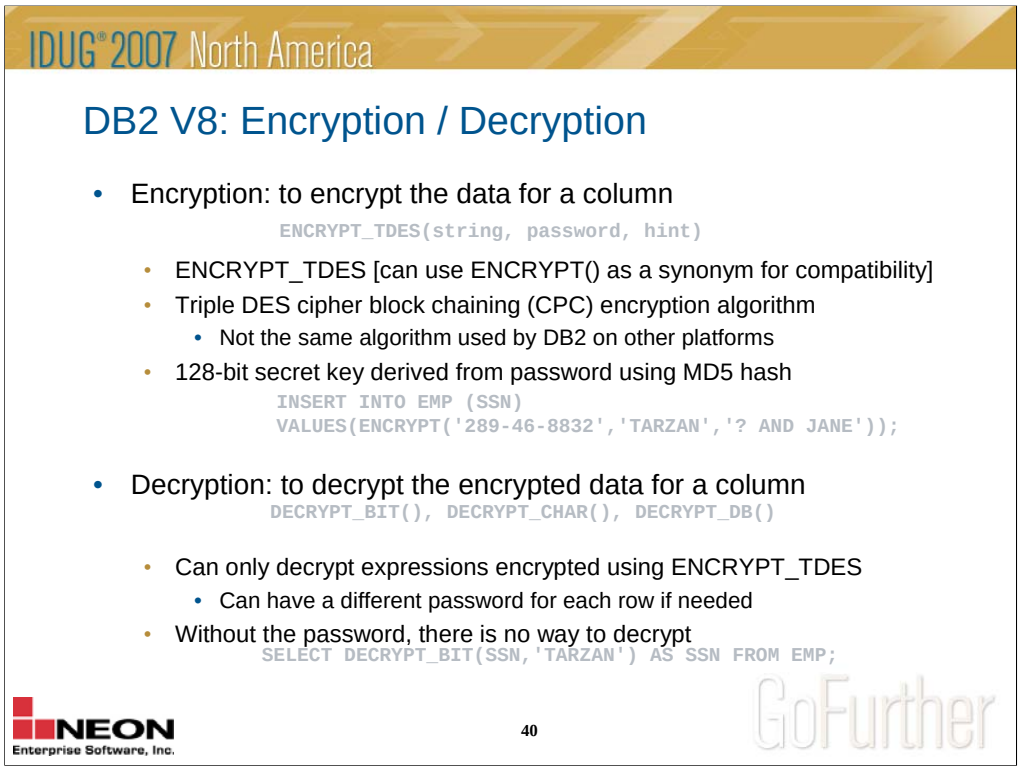

DB2 V8 offers new functions that allow you to encrypt and decrypt data at the column level. Because you can specify a different password for every row that you insert, you can really encrypt data at the "cell" level in your tables. If you use these functions to encrypt your data, be sure to put some mechanism in place to manage the passwords that are used to encrypt the data. Without the password, there is absolutely no way to decrypt the data.

To assist you in remembering the password, you have an option to specify a hint (for the password) at the time you encrypt the data. The SQL example in the middle of the screen shows an INSERT that encrypts the SSN (social security number) using a password and a hint. Now, in order to retrieve the row you must use the DECRYPT function supplying the correct password. This is shown in the SELECT statement at the very bottom of the slide. If we fail to supply a password, or the wrong password, the data is returned in an encrypted format that is unreadable.

The result of encrypting data using the ENCRYPT function is VARCHAR -- FOR BIT DATA.

(CAN SKIP THIS → The encoding scheme of the result is the same as the encoding scheme of the expression being encrypted.)

When encrypting data keep the following in mind. The encryption algorithm is an internal algorithm. For those who care to know, it uses Triple DES cipher block chaining (CBC) with padding and the 128-bit secret key is derived from the password using an MD5 hash.

When defining columns to contain encrypted data the DBA must be involved because the data storage required is significantly different. The length of the column has to include the length of the non-encrypted data + 24 bytes + number of bytes to the next 8 byte boundary + 32 bytes for the hint.

The decryption functions (DECRYPT_BIT, DECRYPT_CHAR, and ECRYPT_DB) return the decrypted value of the data. You must supply the encrypted column and the password required for decryption. The decryption functions can only decrypt values that are encrypted using the DB2 ENCRYPT function.

Some additional details on encryption are provided on slide 79 but let's go on to slide 80 now.

You may have noticed that there is one encryption function but three decryption functions. The result of the function is determined by the function that is specified and the data type of the first argument. The next slide provides details that you can review later. Let's go on to slide 80.

Built-in functions for encryption and decryption require cryptographic hardware in a cryptographic coprocessor, cryptographic accelerator, or cryptographic instructions. You must also have the z/OS Cryptographic Services Integrated Cryptographic Service Facility (ICSF) software installed.

[[Encrypted data can be decrypted only on servers that support the decryption functions that correspond to the ENCRYPT_TDES function. Therefore, replication of columns with encrypted data should only be to servers that support the decryption functions.

Issue #4: Auditing

- **In a world replete with regulations and threats, organizations have to go well beyond securing their data. Essentially, they have to perpetually monitor their data in order to know who or what did exactly what, when and how – to all their data.**

  Source: Audit the Data – or Else:
  Un-audited Data Access Puts Business at High Risk.
  Baroudi-Bloor, 2004.

- **HIPAA, for example, requires patients to be informed any time someone has even *looked* at their data.**

NEON Enterprise Software, Inc.

41

GoFurther

Auditing helps to answer questions like “Who changed data?” and “When was the data changed?”

and “What was the old content was prior to the change?” Your ability to answer such questions is

very important for regulatory compliance. Sometimes it may be necessary to review certain audit

data in greater detail to determine how, when, and who changed the data.

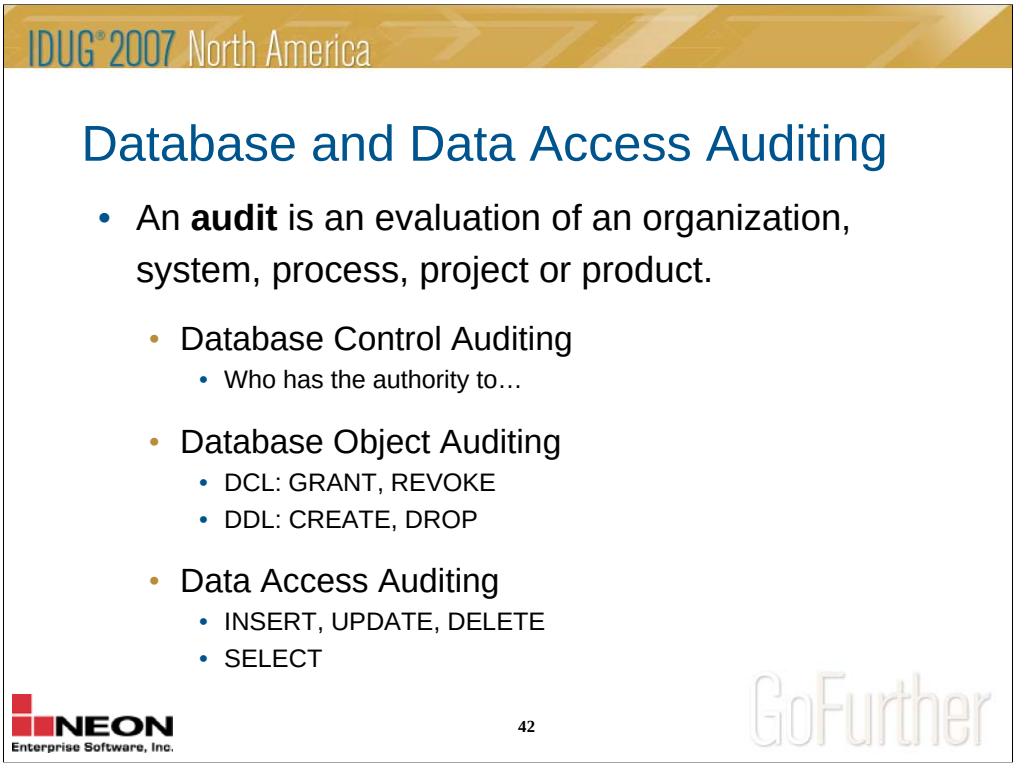

Auditing tools should not only capture and present to you the audit data, but they should also alert you on activities and quickly identify records that are needed for an audit.

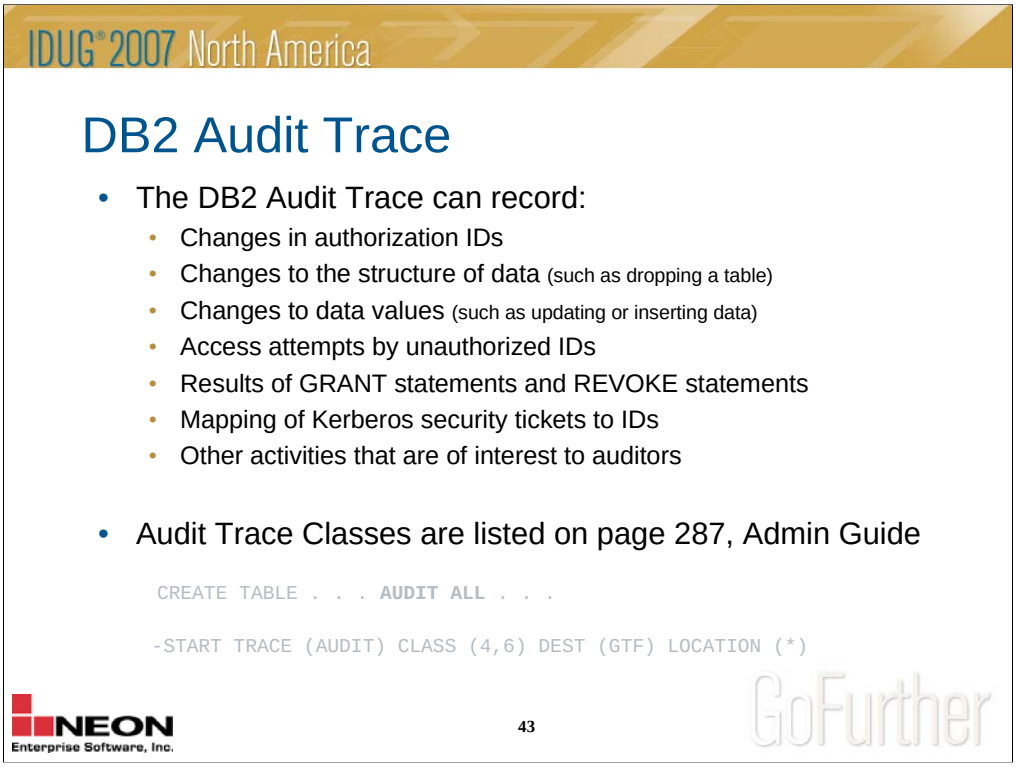

Limitations of the audit trace The audit trace does not record everything, as the following list of limitations indicates:

•Auditing takes place only when the audit trace is on.

•The trace does not record old data after it is changed (the log records old data).

•If an agent or transaction accesses a table more than once in a single unit of recovery, the audit trace records only the first access.

•The audit trace does not record accesses if you do not start the audit trace for the appropriate class of events.

•The audit trace does not audit some utilities. The trace audits the first access of a table with the LOAD utility, but it does not audit access by the COPY, RECOVER, and REPAIR utilities. The audit trace does not audit access by stand-alone utilities, such as DSN1CHKR and DSN1PRNT.

•The trace audits only the tables that you specifically choose to audit.

•You cannot audit access to auxiliary tables.

•You cannot audit the catalog tables because you cannot create or alter catalog tables.

Using the Log?

- How to review the changes made to large financial databases… or to any databases under the purview of regulations?
  - The transaction log(s) capture ALL changes made to ALL data.

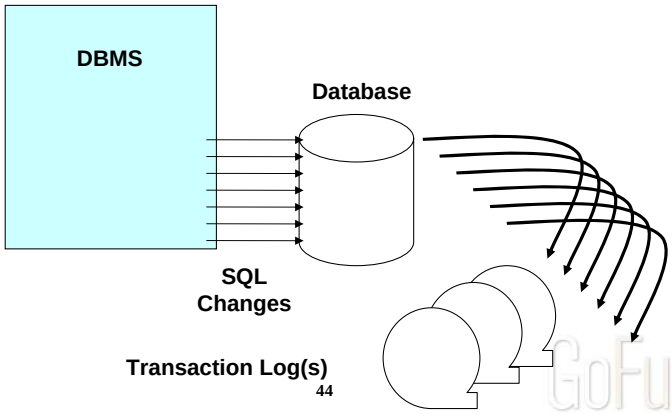

NEON Enterprise Software, Inc.

44

GoFurther

The log is the database – the database tables are just optimized access paths to the current state of the log.

And is all your audit trail being written to write-once media, so that it can't be manipulated to hide the culprit's tracks?

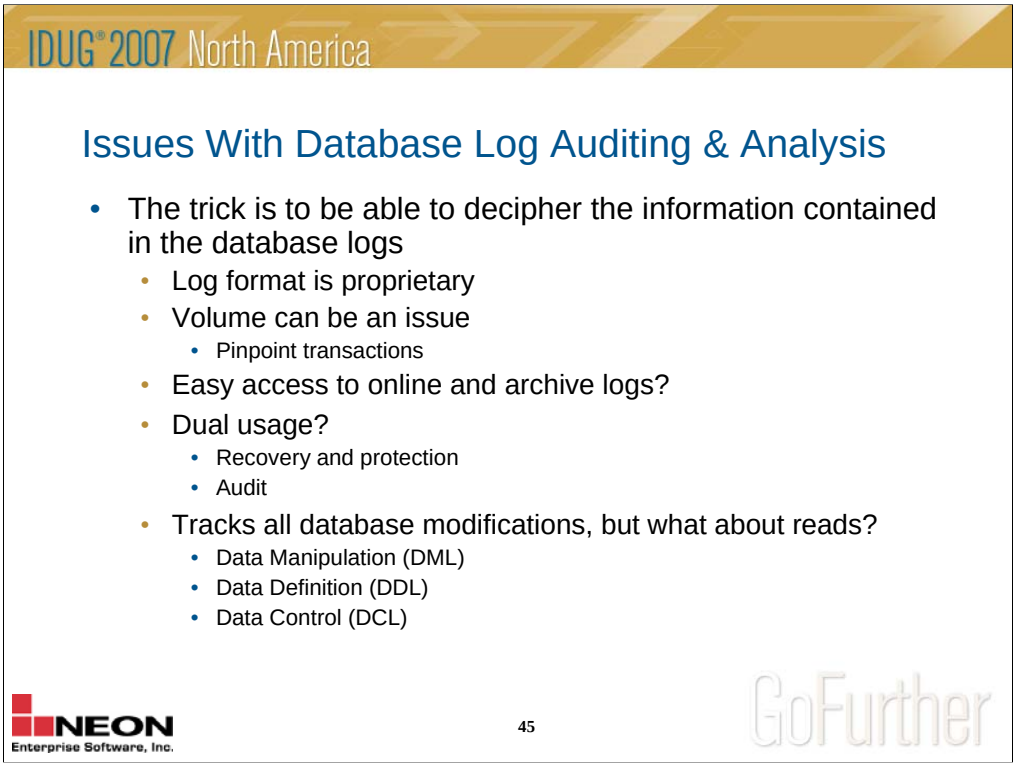

Using the log can be useful to confirm approved modifications and identify unauthorized changes. The log can show object changes that might affect financial reporting, unauthorized transactions (Who? What? Detect, report, and correct . . .) And it has the additional benefits of not requiring additional hardware and potentially being able to ensure SQL-based data recoverability.

But, some companies may discover that scanning the entire log with any of the currently available anaylzers could take more than 24 hours (think about large data sharing groups). I think we should work to charge the scanning costs to the accounting department, including the hardware upgrades that will undoubtedly be necessary to support such an enormous process. It might bring some sanity to the discussion.

Bottom Line

- Database auditing software can produce multiple userful reports on database *modification* activities
  - Reads are problematic and are not stored on the database transaction log
  - Auditing software can resolve many of the issues with transaction log based auditing
  - Many of the tools today that purport to work with mainframe data only “audit” network access
    - That is, they are network sniffers and will not show, for example, CICS transactions that access DB2 for z/OS

Issue #5: DBA Procedures

- Are changes made to your database environment using standard, repeatable methods and in a controlled manner?
- Are your databases properly backed up and recoverable within the timeframe mandated by your business (and any regulations)?

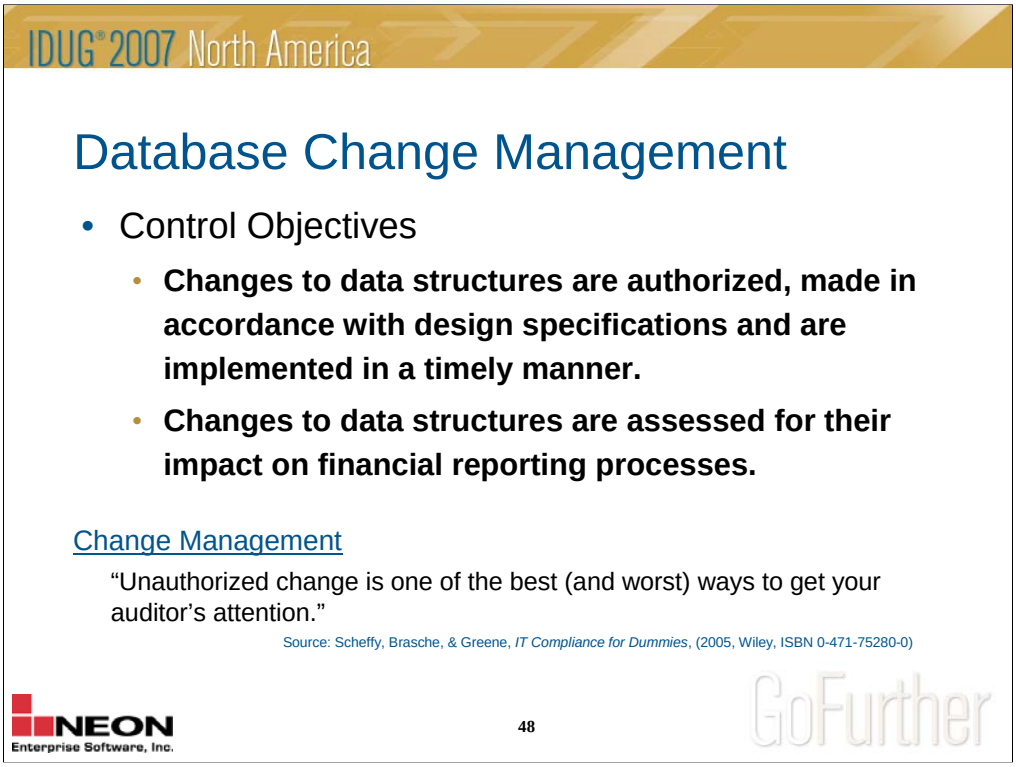

Objective Summary:

Changes made to the data structures must be done using an authorized process. All impacts and deltas are tracked and reported to ensure that there are no unauthorized changes that could invalidate the financial reporting systems

Solution Summary:

Workflow and task approval process

Track all changes to data structures and catch those that are made outside of the authorized change approval process

Produce delta of changes between releases to provide complete summary of data structure changes for compliance reporting

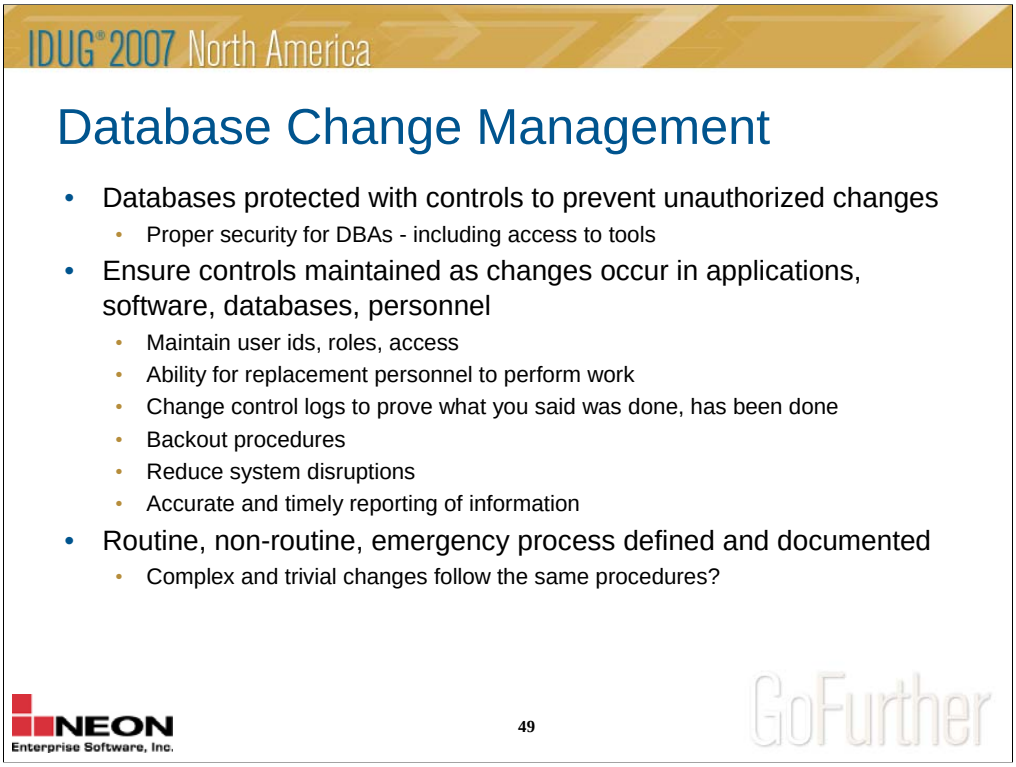

It is possible to develop a process using change management tools to automatically track all structure changes to the database using the change auditing and comparison features of such tools. By capturing all deltas between successive upgrades/changes, a record of what changed and when is available to the auditors.

Catalog/dictionary visibility tools offer the ability to keep a complete log all database change activity performed by DBAs/developers.

Additional thoughts from "SOX Requirements for DBAs in Plain English by James McQuade" http://www.dbazine.com/ofinterest/oi-articles/mcquade2:

SOX requirements for Database Change Management in plain English:

> Changes to the database are widely communicated, and their impacts are known beforehand.
>
> Installation and maintenance procedure documentation for the DBMS is current.
>
> Data structures are defined and built as the designer intended them to be.
>
> Data structure changes are thoroughly tested.
>
> Users are apprised, and trained if necessary, when database changes imply a change in application behavior.
>
> The table and column business definitions are current and widely known.
>
> The right people are involved throughout the application development and operational cycles.
>
> Any in-house tools are maintained and configured in a disciplined

And it all Requires…

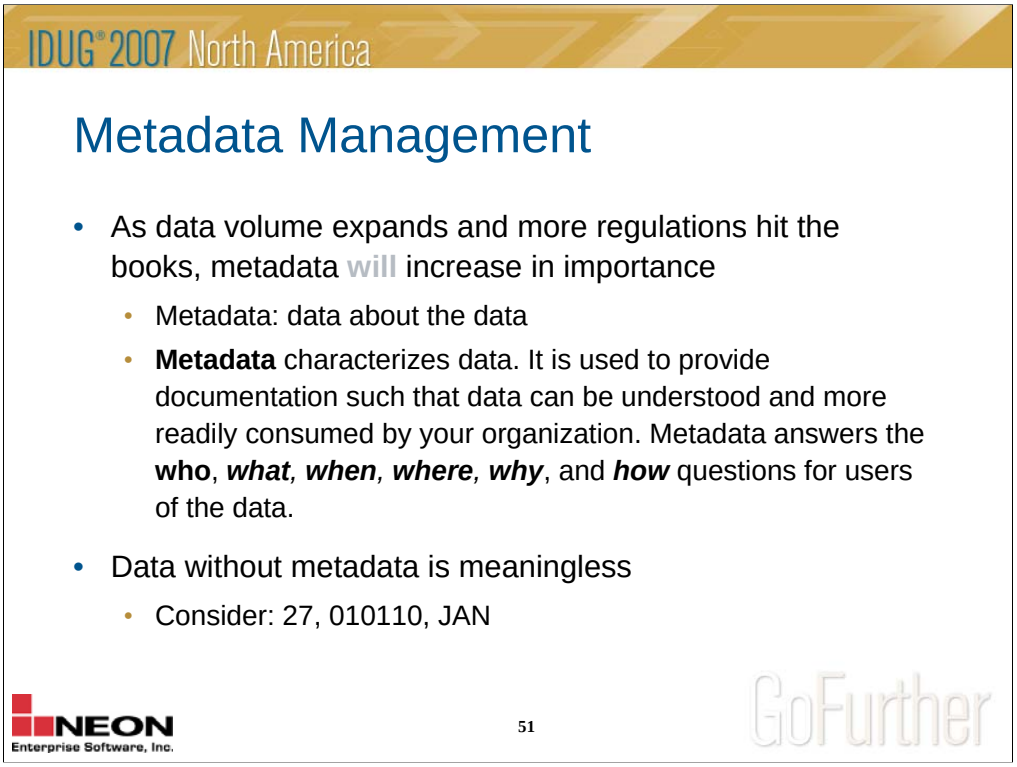

Antiques Roadshow metaphor

27 – number in base 10? Octal (which would be 23 in base 10)? An age? An IQ? Dollar amount?

010110 - binary number? a date, January 1, 1910? January 1, 2010? something else?

JAN - woman's name? man's name? first month of the year? An acronym? something else?

Data Categorization

- Data categorization is critical
  - Metadata is required to place the data into proper categories for determining which regulations apply
    - Financial data → SOX
    - Health care data → HIPAA
    - Etc.
  - Some data will apply to multiple regulations
  - Who does this now at your company? Anyone?

Data must be categorized to ensure that it is treated appropriately for regulatory compliance.

The who, what, where, when, and why for each piece of data is what determines how that data is governed:

Audit, Protection, Retention, etc.

So, metadata increases in importance!

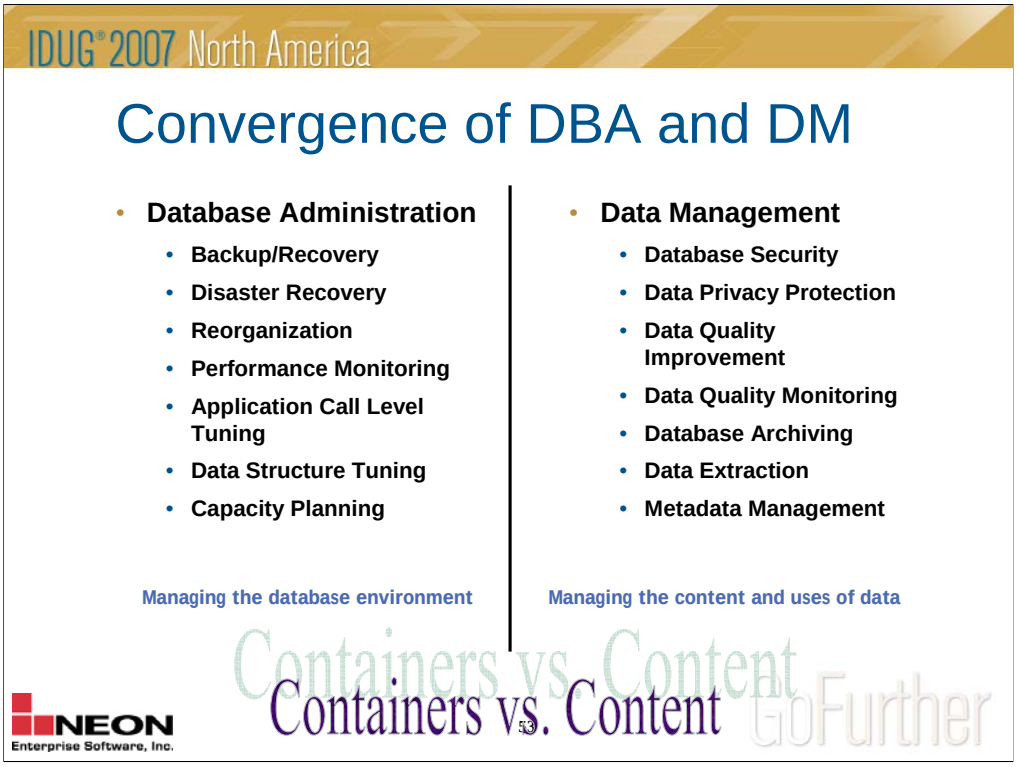

More and more, the DBA will be transitioning to take on more non-traditional tasks dealing with data management. With the impending and continuing impact of autonomics, DBAs may be freed from some of the traditional DBA tasks and therefore have time to concentrate on expanding their role into the world of "DM"…

Data Management

- Tasks definitions are emerging
- No standard Job Titles or Descriptions
- Aligned with business not just IT
- Little Vendor Support
- DBMS architectures built without consideration of DM
- IT management has not been supportive (NMP)
- Executive management has not been supportive
- Companies have accrued many penalties for not paying attention to DM requirements

DBA

- Very well defined tasks
- Very well defined Job Title and Description
- Functions fall entirely in IT
- Overwhelming vendor support
- DBMS architectures fully supportive
- Must be done well to support efficient operational environment

Containers versus Content

(we break down the content side further on the next slide)

So, What is the Impact on DBAs?

- Data Management vs. Database Administration
  - Data Governance
- Business acumen vs. bits-and-bytes
  - Can't abandon technical skills, but must add more soft skills to your bag of tricks
- Communication & Inter-organizational Cooperation
  - IT, Lines of Business, Legal

Increasing governmental regulations will cause DBAs to become more business focused as the are drafted to support compliance efforts. Regulations, IMHO, are a good thing. Sometimes they can cause people & companies to do things that they never would've done in the first place. SEAT BELTS example (they didn't even exist, then no one used them, then law passed, now people are afraid of a $50 citation so they use them – wouldn't you think fear of death would be worse than fear of losing $50? Evidently not.)

Additionally, inter-organization communication will be imperative. More interaction with business users is a must, and DBAs are accepting of this. But more interaction with the legal department will also be needed. How often do DBAs do that today? Almost never – only when they need a contract approved.

But, Don't Despair…

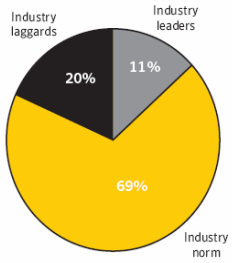

	# of overall deficiencies	# of significant and material deficiencies	% of compliance deficiencies considered significant and material
Industry leaders	25	2	8%
Industry norm	36	6	16.6%
Industry laggards	75	35	46.6%

Figure 1. Performance results for regulatory compliance
Source: SecurityCompliance.com, 2006

- …by 2008, less than 10% of organizations will succeed at their first attempts at data governance because of cultural barriers and a lack of senior-level sponsorship. Source: Gartner, Inc. as reported by the web portal SearchDataManagement

Web References

- Industry Organizations and References
  - www.isaca.org
  - www.coso.org
  - www.aicpa.org/news/2004/2004_0929.htm
  - www.auditnet.org/sox.htm
  - www.itgi.org
  - www.sox-online.com/coso_cobit.html
  - www.bis.org/publ/bcbs107.htm - (Basel II)
  - www.snia-dmf.org/100year

Recommended Reading

- *Implementing Database Security and Auditing* by Ron Ben Natan (2005, Elsevier Digital Press, ISBN: 1-55558-334-2)
- *Manager's Guide to Compliance* by Anthony Tarantino (2006, John Wiley & Sons, ISBN: 0-471-79257-8)
- *IT Compliance for Dummies* by Clark Scheffy, Randy Brasche, & David Greene (2005, Wiley Publishing, Inc., ISBN: 0-471-75280-0)
- *Cryptography in the Database: The Last Line of Defense* by Kevin Kenan (2006, Addison-Wesley, ISBN: 0321-32073-5)
- *Electronic Evidence and Discovery: What Every Lawyer Should Know* by Michele C.S. Lange and Kristin M. Nimsger (2004, ABA Publishing, ISBN: 1-59031-334-8)

Session: I01

The Impact of Regulatory Compliance on DBA

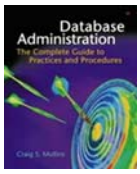

Craig S. Mullins

NEON Enterprise Software, Inc.

craig.mullins@neonesoft.com

www.DB2portal.com

www.CraigSMullins.com

www.neonesoft.com

GoFurther