



IBM Information Protection Capabilities on z/OS

Ernie Mancill
Executive IT Specialist
mancill@us.ibm.com



Agenda

- The need to protect data
- Information Protection entry point to Information Governance
- **IBM's Information Protection capabilities**
- Summary

Organizations facing many of the following challenges

- How to secure your data - DB2 V10 for z/OS
- Audit and separation of roles – privileged user conundrum
- Encryption and data obfuscation
- Discovering what data needs to be secured
- Data in a test environment
- Data life-cycle management and data growth

IBM Information Management Solutions for System z – End to end Solution

Security, Audit, and Encryption for z/OS



Satisfy Your Auditor: Plan, Protect and Audit

■ Data Access


- Minimize the use of a superuser authorities such as SYSADM
- A different group should manage access to restricted data than the owner of the data

■ Data Auditing

- Any dynamic access or use of a privileged authority needs to be included in your audit trail
- Maintain historical versions of data for years or during a business period

■ Data Privacy

- All dynamic access to tables containing restricted data needs to be protected



Today's Mainframe:
*The power of industry-leading security,
the simplicity of centralised management*

Streamline and simplify compliance processes

- Alerts of suspicious activity
- Audit reporting and sign-offs
- Separation of duties – creation of policies vs. reporting on application of policies
- Trace users between applications, databases
- Fine grained-policies
- Sign-off and escalation procedures
- Integration with enterprise security systems (SIEM)



DB2 10 for z/OS Security Enhancements

Help Satisfy Your Auditors using new features

- ✓ New granular authorities to reduce data exposure for administrators
- ✓ New auditing features using new audit policies comply with new laws
- ✓ New row and column access table controls to safe guard your data
- ✓ New temporal data to comply with regulations to maintain historical data




Reduce your risk by minimizing use of SYSADM

New granular system authorities and install security parameters

New in DB2 10

▪ Prior to DB2 10

- SYSADM
- DBADM
- DBCTRL
- DBMAINT
- SYSCTRL
- PACKADM
- SYSOPR



- System DBADM
- ACCESSCTRL
- DATAACCESS
- SECADM
- SQLADM
- EXPLAIN



Prevents SYSADM and SYSCTRL from granting or revoking privileges

- New separate security install zparm parameter
- New install **SECADM** authority manages subsystem security
- SYSADM and SYSCTRL can no longer implicitly grant or revoke privileges

Control cascading effect of revokes

- New revoke dependent privileges install parameter
- New revoke dependent privileges SQL clause

New authority for performing security tasks without ability to change or access data

- **SECADM** authority
 - Allows the user to
 - Issue SQL GRANT, REVOKE statements on all grantable privileges and administrative authorities
 - Manage DB2 9 roles and trusted contexts
 - Manage DB2 10 row permissions and column masks
 - Manage DB2 10 Audit policies
 - Access catalog tables
 - Issue START, STOP, and DISPLAY TRACE commands



New authority for managing objects without ability to access data or control access to data

- **System DBADM** authority

- Allows the user to
 - Issue SQL CREATE, ALTER, DROP statements to manage most objects in the DB2 subsystem
 - Issue most DB2 commands
 - Execute system defined stored procedures and functions
 - Access catalog tables



Satisfy Your Auditor:

New **audit policies** provide needed flexibility and functionality

- New auditing capability allows you to comply without the need of external data collectors
 - New audit policies managed in catalog
 - Audit privileged users
 - Audit SQL activity against a table
 - Audit distributed identities



New Audit Policies Feature

- Your security administrator using the new SECADM authority maintains DB2 audit policies in a new catalog table
 - **SYSIBM.SYSAUDITPOLICIES**
- Audit policies enabled using **-STA TRACE** command
- Audit policies disabled using **-STO TRACE** command
- Up to 8 audit policies can be specified to auto start or auto start as secure during DB2 start up
 - Only user with SECADM authority can stop a secure audit policy trace

DB2 for z/OS is still limited to a total number of active traces (of all types) at 32



New Audit Policies Feature

- Auditor can audit access to specific tables for specific programs during day
 - Audit policy does not require AUDIT clause to be specified using DDL to enable auditing (no more DBA involvement and no package invalidation)
 - Audit policy generate records for all read and update access not just first access
 - Audit policy includes additional records identifying the specific SQL statements
 - Audit policy provides wildcarding of based on schema and table names
- Auditor can identify any unusual use of a privileged authority
 - Records each use of a system authority
 - Audit records written only when authority is used for access
 - External collectors only report users with a system authority



New improved security features provide more effective controls and accurate audit trail for remote access

- Support distributed identities introduced in z/OS V1R11
 - A distributed identity is a mapping between a RACF user ID and one or more distributed user identities, as they are known to application servers
- Support client certificates and password phrases in z/OS V1R10
 - AT-TLS secure handshake accomplishes identification and authentication when the client presents its certificate as identification and its proof-of-possession as authentication
 - A RACF password phrase is a character string made up of mixed-case letters, numbers, special characters, and is between 9 to 100 characters long
- Support connection level security enforcement
 - Enforces connections must use strong authentication to access DB2
 - All userids and passwords encrypted using AES, or connections accepted on a port (Secure Port) which ensures AT-TLS policy protection or protected by an IPSec encrypted tunnel

Satisfy Your Auditor:

New table controls to protect against unplanned SQL access

- Define additional data controls at the row and column level
 - Security policies are defined using SQL
 - Separate security logic from application logic
- Security policies based on real time session attributes
 - Protects against SQL injection attacks
 - Determines how column values are returned
 - Determines which rows are returned
- No need to remember various view or application names
 - No need to manage many views; no view update or audit issues
- All access via SQL including privileged users, adhoc query tools, report generation tools is protected
- Policies can be added, modified, or removed to meet current company rules without change to applications



Table controls to protect SQL access to individual row level

Establish a row policy for a table

- Filter rows out of answer set
- Policy can use session information, e.g. the SQL ID is in what group or user is using what role, to control which row is returned in result set
- Applicable to SELECT, INSERT, UPDATE, DELETE, & MERGE
- Defined as a row permission:

***CREATE PERMISSION policy-name ON table-name
FOR ROWS WHERE search-condition
ENFORCED FOR ALL ACCESS ENABLE;***

Optimizer inserts search condition in all SQL statements accessing table. If row satisfies search-condition, row is returned in answer set

Table controls to protect SQL access to individual column level

Establish a column policy for a table

- Mask column values in answer set
- Policy can use session information, e.g. the SQL ID is in what group or user is using what role, to control what masked value is returned in result set
- Applicable to the output of outermost subselect
- Defined as column masks :

```
CREATE MASK mask-name ON table-name  
FOR COLUMN column-name RETURN CASE-expression  
ENABLE;
```

Optimizer inserts CASE expression in all SQL statements accessing table to determine mask value to return in answer set

RACF and Data Servers on z/OS

- RACF and DB2
 - DB2 Subsystem Access Control (outside of DB2)
 - Control connections to the DB2 subsystem
 - CICS
 - IMS
 - CAF
 - BATCH
 - Assign identities
 - Protect the underlying DB2 data store (underlying data sets of DB2 can be protected by RACF dataset services)
 - In addition to database server-provided security, RACF can be used to control access to database objects, authorities, commands and utilities by using the RACF access control module of the database server.

- RACF and IMS
 - The IBM Information Management System (IMS™) has been enhanced to make use of RACF for controlling access to IMS resources. It is possible to use the original IMS security features, the new RACF features, and combinations of these. RACF provides more flexibility than the older security features. The normal features of RACF can be used to protect both system and database IMS data sets

Tools from Tivoli to enhance RACF

- Tivoli zSecure Admin
 - User friendly layer over the native RACF administration panels
 - Automatically generated RACF commands
 - Reduce complexity
 - Increased RACF administration productivity
 - Fewer errors
 - Less risk of inadvertent data exposure due to inappropriate/insufficient security
- Tivoli zSecure Visual
 - GUI/Windows based UI
 - Insulates security administrators from TSO/ISPF
 - Increased productivity requiring less sophistication in administration skills
- Tivoli Identity Management software
 - Tivoli Directory Server
 - Tivoli Identity Manager

End User Identity Mapping - Why is this Important?

- Proper end user assignment of rights and privileges on the data server is important, but equally important:
 - In many multi-tier implementations, to ease administration, and to influence performance through mechanisms such as connection pooling, thread reuse, etc. Shared (common) authorization IDs are used for connecting to the Data Server
 - In these types of implementations, this leads to loss of end user identification, and any associated ability to completely audit activity on the data server from these types of connections.
 - Various mechanisms can be use to preserve these credentials:
 - SQL Language Extension vis SQLESETI
 - Extended identity propagation using JDBC drivers
 - Enterprise Identity Mapping
- *Support distributed identities introduced in z/OS V1R11*
 - *A distributed identity is a mapping between a RACF user ID and one or more distributed user identities, as they are known to application servers*

The DB2 client information fields

DB2 allows applications to send information about them to the database with each SQL operation.

- The database externalizes this information then in its monitoring data
- The performance impact of setting them is negligible (but for DB2 on LUW V9.1 FP6 is recommended)
- The data can be set by the application itself, or via database driver properties (see next slides)
- The following information can be set:

Field	Description	Length (LUW, z/OS)
Client user ID	This user ID is for identification purposes only, and is not used for any authorization. It typically identifies the user of an application.	255, 16
Client workstation name	The workstation name of the client system. Some applications also use this field to identify the business transaction executed within an application.	255, 18
Client application name	It can be used to identify the application hosted in an application server, or to identify the business transaction within an application.	255, 32
Program name	Identifies the application running on the client. It is only supported for a connected DB2 on z/OS database.	-, 80
Accounting string	It can be used to specify charge-back information, or to add additional monitoring details about the database workload.	200, 200

Ways to instrument your application

JDBC offers methods of class `com.ibm.db2.jcc.DB2BaseDataSource`¹

```
public static void main(String[] args) {
    String url = "jdbc:db2://lap1.boeblingen.de.ibm.com:50000/DEMO";
    Class.forName("com.ibm.db2.jcc.DB2Driver");
    Connection conn = DriverManager.getConnection(url, user, password);

    conn.setClientInfo("ClientUser", "xyz");
    conn.setClientInfo("ClientHostname", "my laptop");

    conn.prepareStatement("SELECT * FROM SYSIBM.SYSDUMMY1" + "WHERE 0 = 1").executeQuery();
}
```

CLI offers the `setSqli()`² interface

```
SQL_API_RC SQL_API_FN sqleseti (
    unsigned short DbAliasLen,
    char * pDbAlias,
    unsigned short NumItems,
    struct sqle_client_info* pClient_Info,
    struct sqlca * pSqlca);

SQL_STRUCTURE sqle_client_info {
    unsigned short type;
    unsigned short length;
    char *pValue; };
```

1) see

<http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.apdv.java.doc/doc/r0021822.html>

2) see

<http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.apdv.api.doc/doc/r0001709.html>


Why auditing is important in a RACF controlled environment

- RACF provides significant controls to protect access to resources, but does little in the way of meaningful access reporting
 - RACF does two things:
 - Prevents people from accessing a resource that is not essential or appropriate for their jobs
 - Allows people access to the necessary data to do their jobs
 - But RACF does NOT:
 - prevent a malicious update if the user has authority to the data.
 - prevent an authorized user from accessing sensitive data that is **NOT** within the scope of their job. E.g. a bank teller looks up the CEO's bank balance or personal customer information
 - provide meaningful information about access to protected DB2 resources (authorized or not).
- DB2 Audit trace will do nothing to protect data, but provides data to help understand what type of access has occurred.
 - Auditing is about ensuring that the appropriate controls are in place to identify inappropriate access and use of production data
 - You need some form of audit facility to watch your privileged users who have RACF and/or DB2 authority and users that have access to sensitive data within the scope of their job
 - Understanding how trusted (privileged) users access sensitive information is essential to ensuring that data is indeed protected

Auditing the Privileged Database User

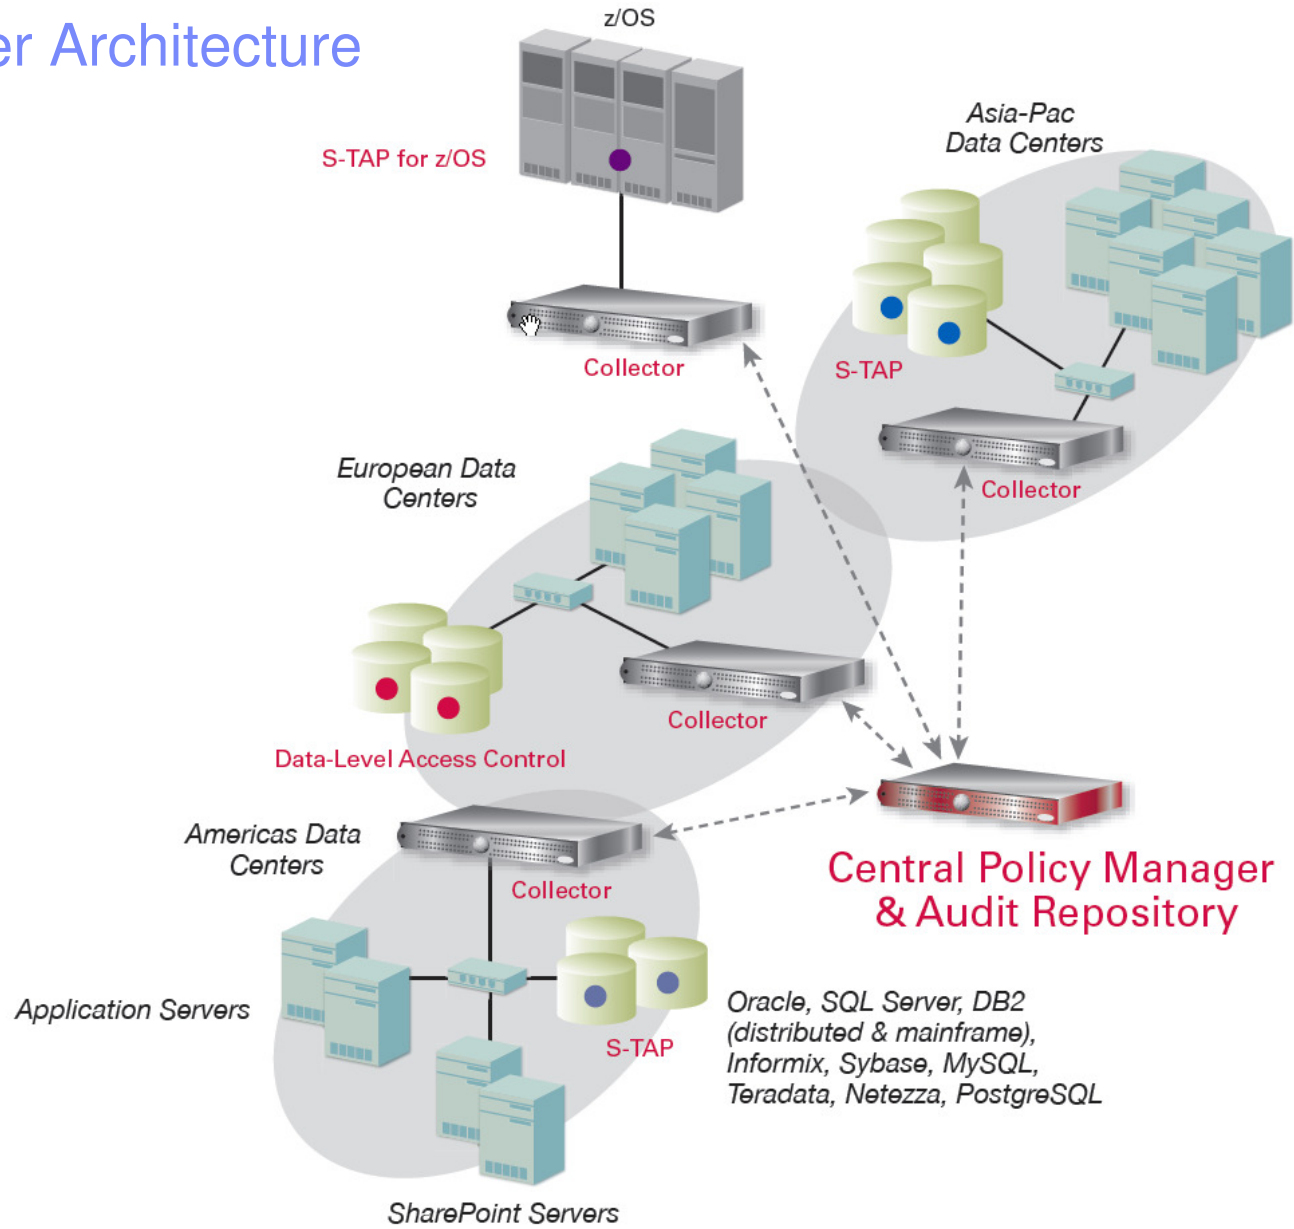
- DB2 trace based processes are managed by DBA's
 - The DBA's are responsible for generating audit data with which they are in turn audited, this constitutes a significant security risk and exposure
 - Trace data collection can be interfered with or turned off completely
 - DBA can issue –DSN Stop Trace
 - Use IFASMFDMF to selectively filter SMF data based on timestamp
 - Use DB2PM (Or Equivalent) filter such as DATE/TIME/EXCLUDE to filter selected records
 - **Having the DBA involved in the collection of audit data is viewed as weak from a compliance and control perspective**
- Security and Auditors with system privileges
 - Also viewed as problematic from a compliance perspective
 - Requires additional technical skills not within their core competencies
 - Misuse of privileges without coordination can result in performance and availability issues
 - Turning on traces without proper filtering to reduce overhead or quantity of trace data collected
 - Altering objects to AUDIT without ensuring that plan/package invalidation is not an issue

Real-Time Database Monitoring with InfoSphere Guardium

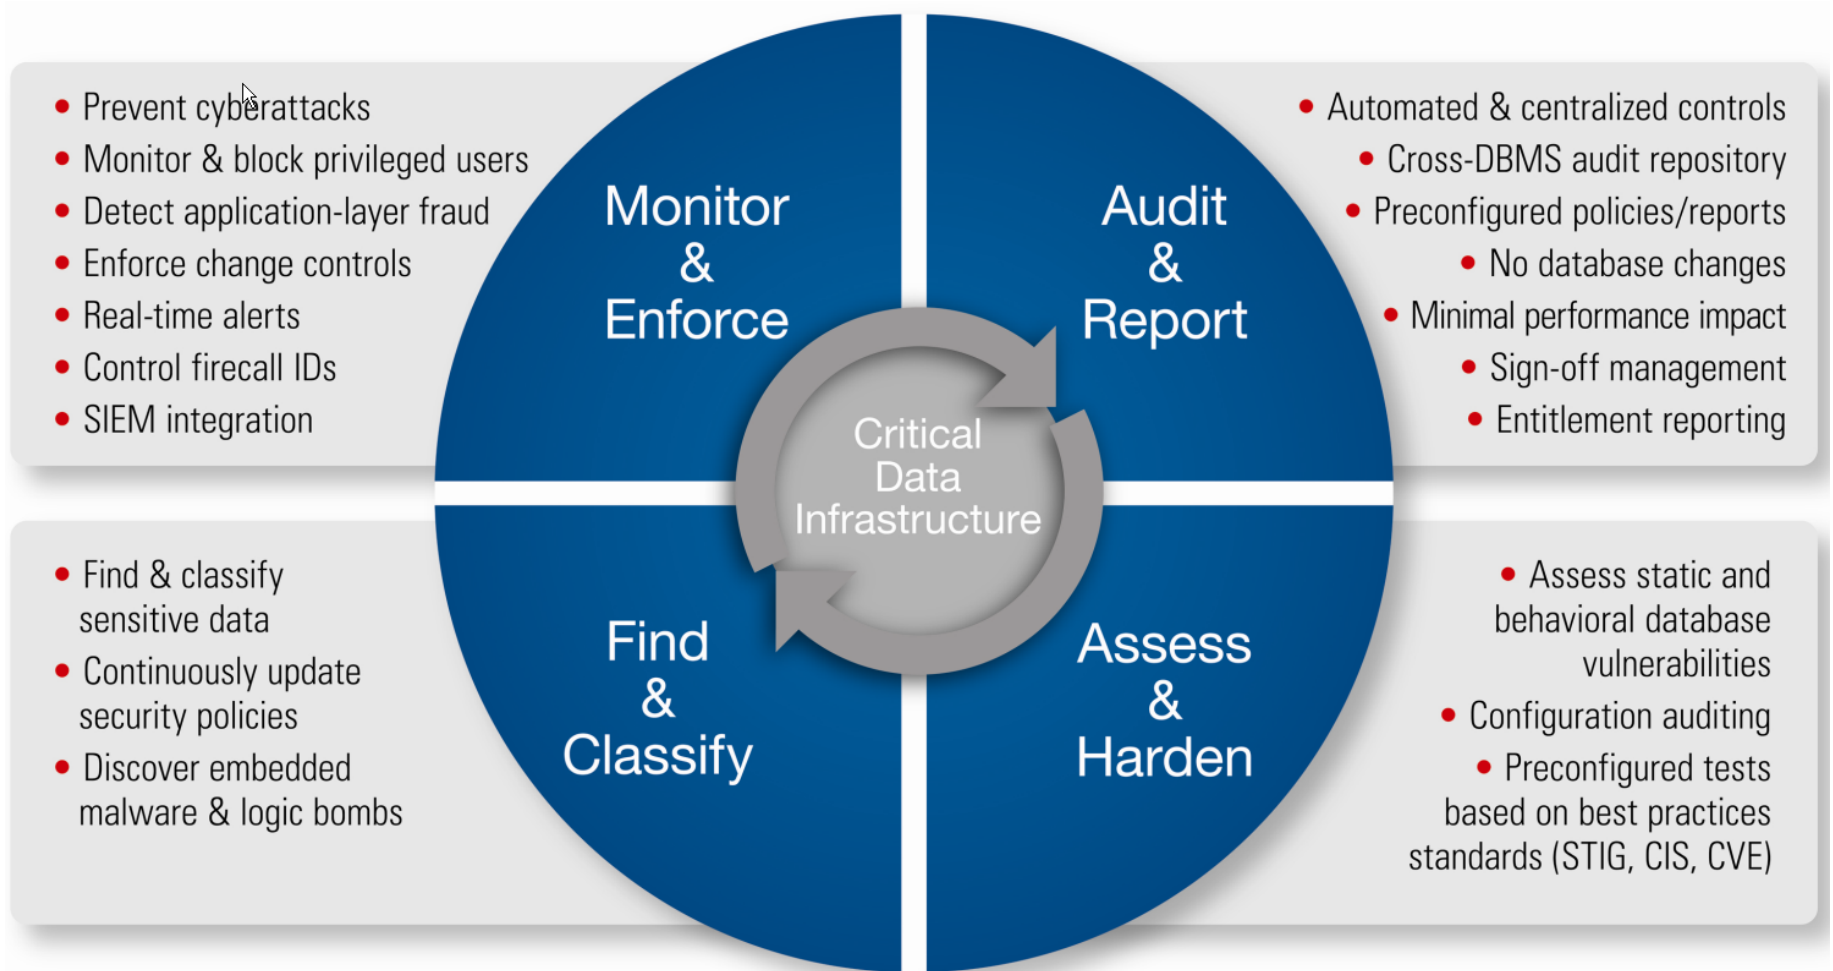


- Non-invasive architecture
 - Outside database
 - Minimal performance impact
 - No DBMS or application changes
- Cross-DBMS solution
- 100% visibility including local DBA access
- Enforces separation of duties
- Does not rely on DBMS-resident logs that can easily be erased by attackers, rogue insiders
- Granular, real-time policies & auditing
 - *Who, what, when, how*
- Automated compliance reporting, sign-offs & escalations (SOX, PCI, NIST, etc.)

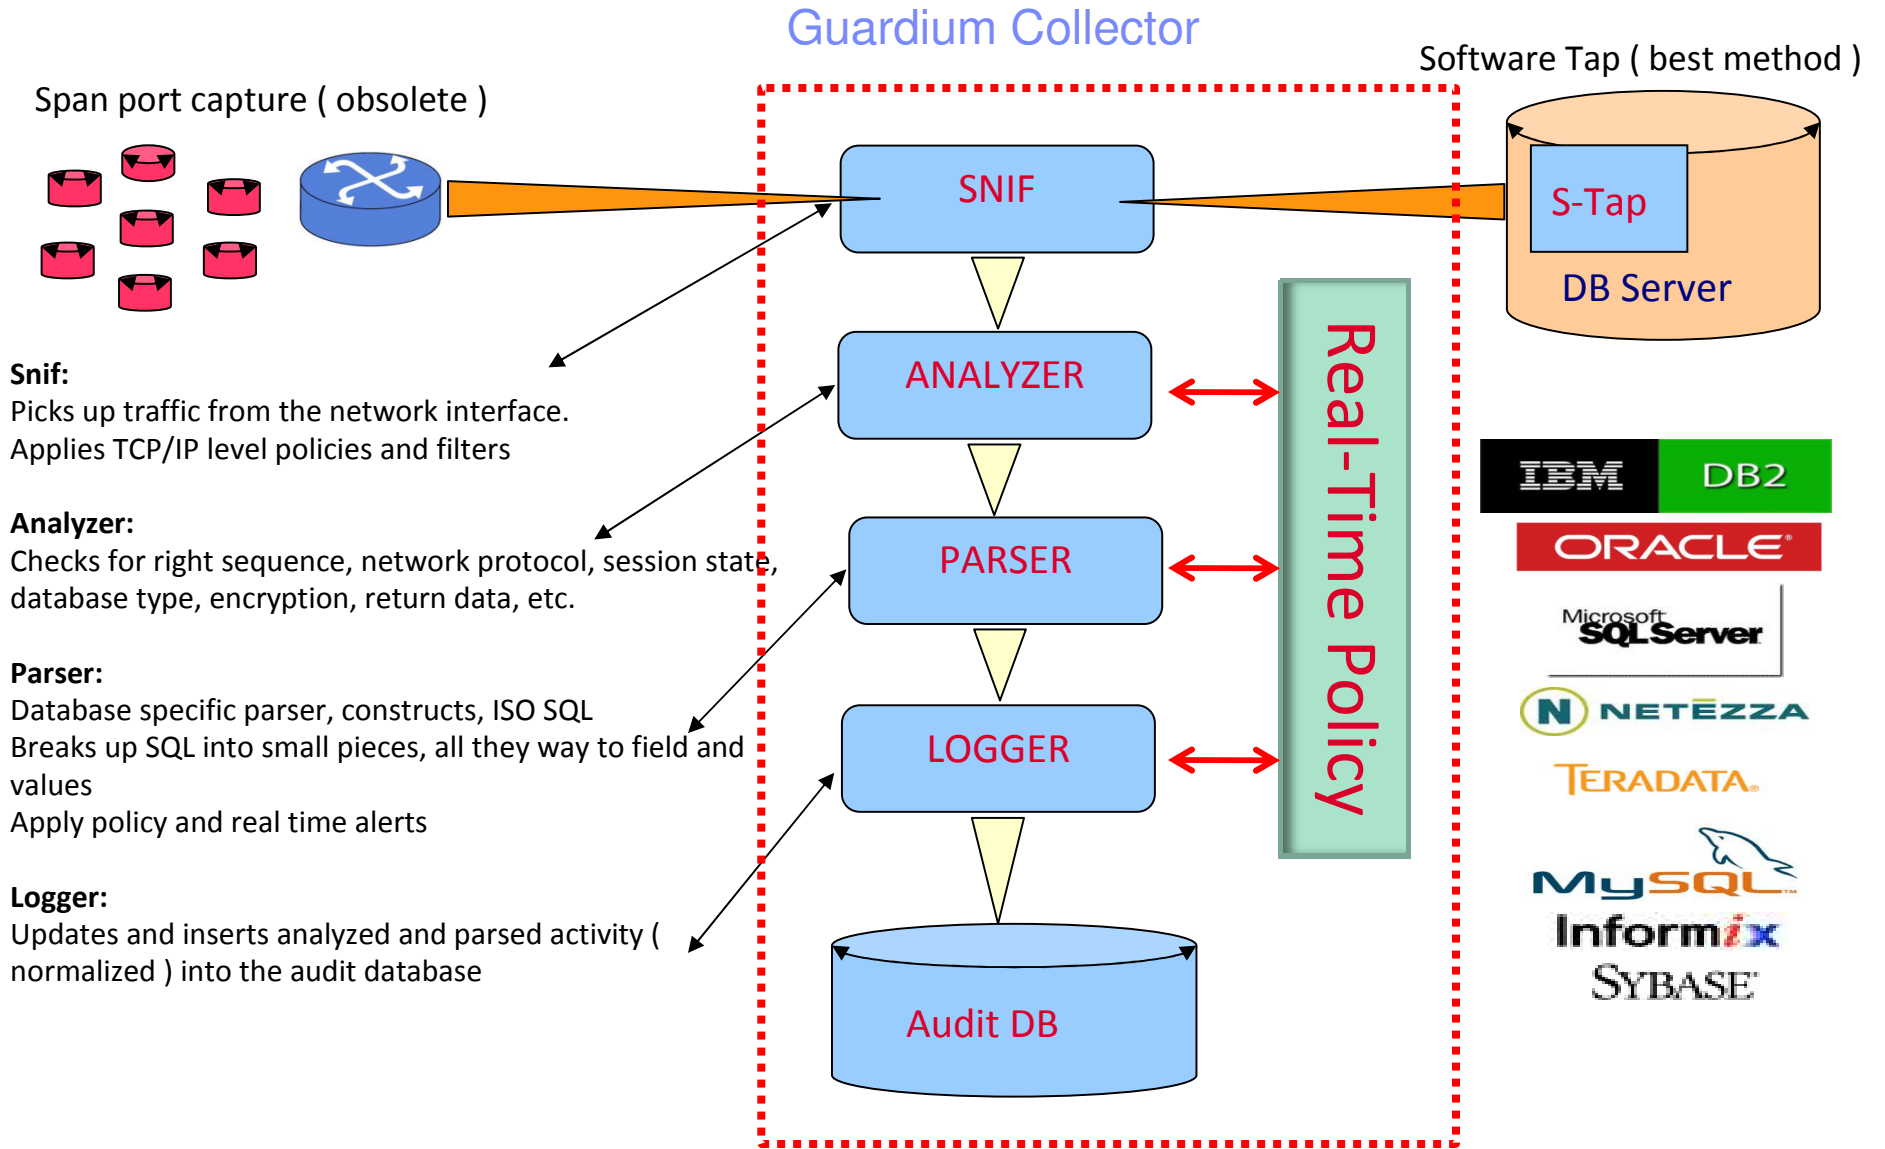
Scalable Multi-Tier Architecture



Addressing the Complete Database Security and Compliance Lifecycle



Architecture



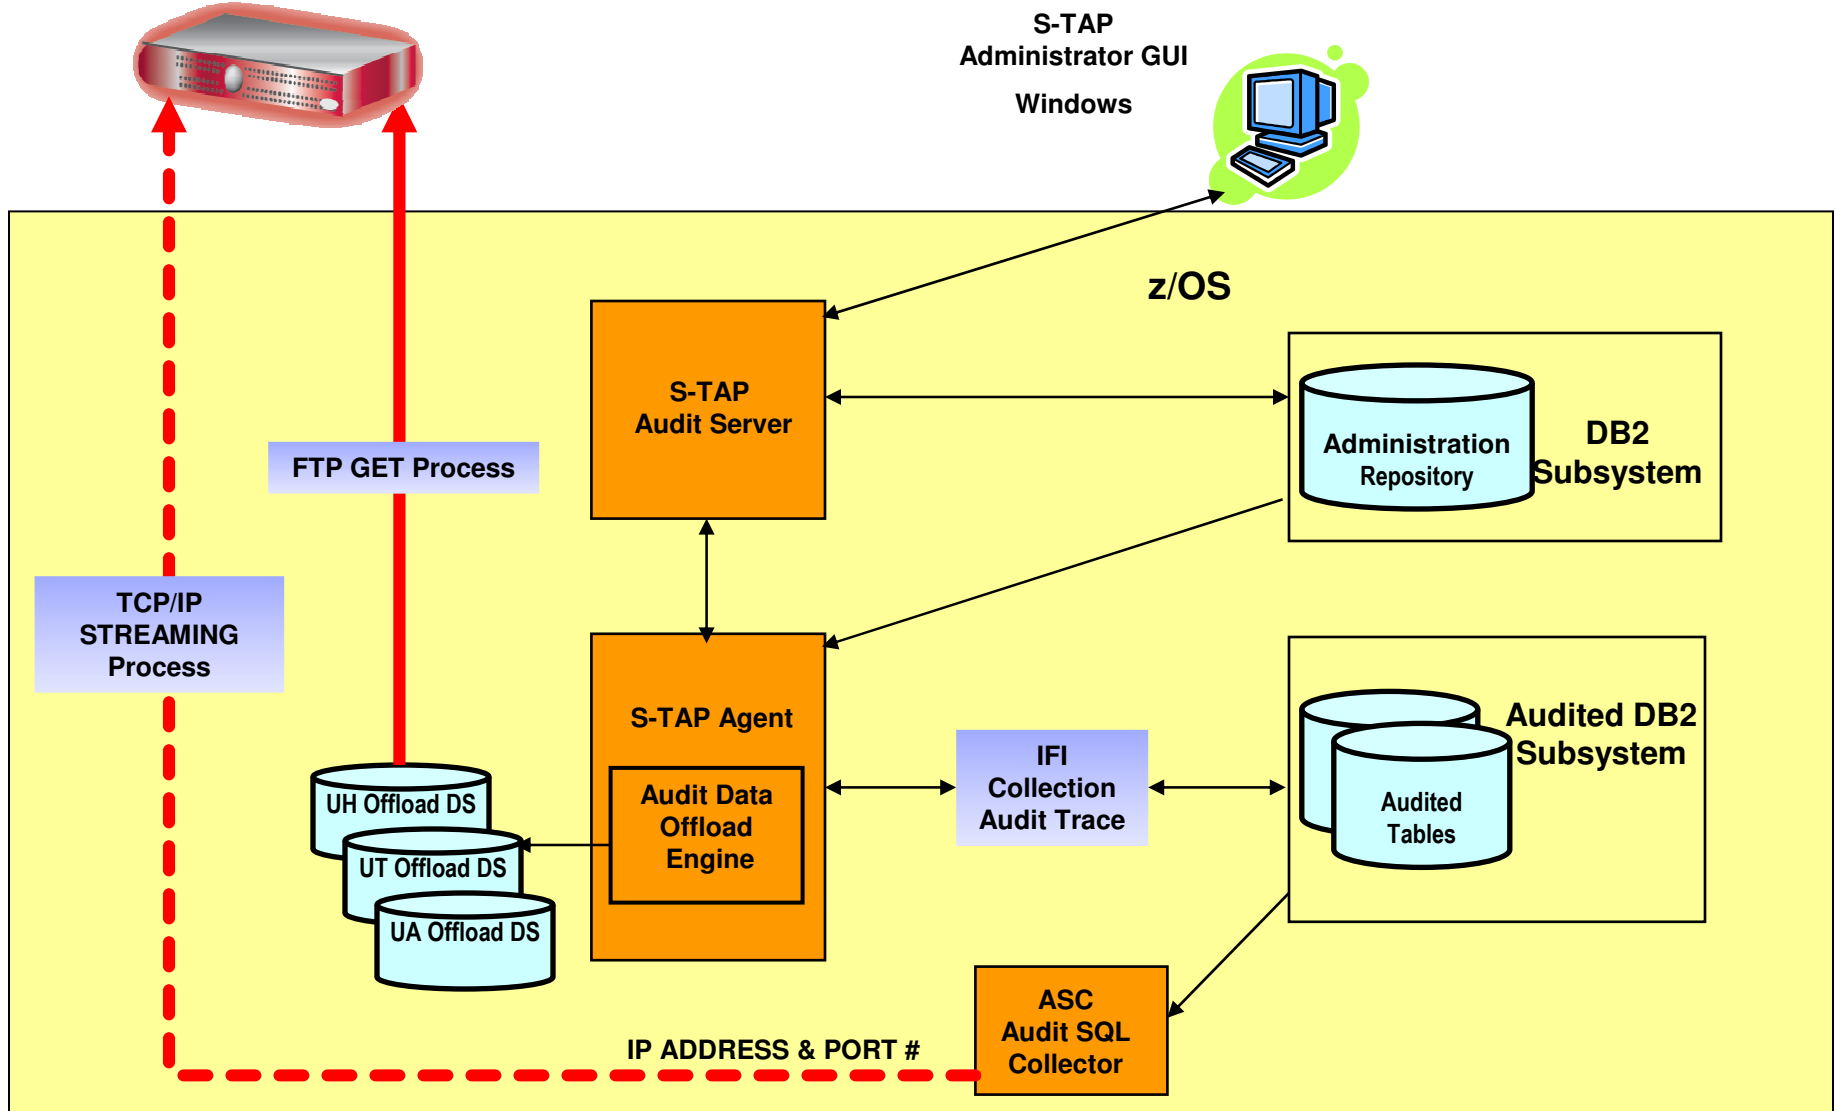
Snif:
Picks up traffic from the network interface.
Applies TCP/IP level policies and filters

Analyzer:
Checks for right sequence, network protocol, session state, database type, encryption, return data, etc.

Parser:
Database specific parser, constructs, ISO SQL
Breaks up SQL into small pieces, all they way to field and values
Apply policy and real time alerts

Logger:
Updates and inserts analyzed and parsed activity (normalized) into the audit database

Guardium S-TAP for DB2 on z/OS Architecture



Guardium S-TAP for DB2/z - Components

- One “Master” started task per LPAR
 - Used to anchor control blocks, no CPU consumption associated with this, the address space is started when the AME
- One S-TAP Server for SYSPLEX
 - Used to manage collection policies and to push collection profile information to each ASC (Audit SQL Collector) started task
- One S-TAP Agent started task per DB2 datasharing member or standalone DB2 subsystem.
 - Externalizing via interval processing the audit data (IFI) to the audit offload datasets
- One ASC started task per DB2 datasharing member or standalone DB2 Subsystem
 - Responsible for hook management, inspection of SQL
 - Started by the S-TAP Agent started task
- S-TAP Administration Client
 - Windows “Fat” client used to administer collection profiles
 - Communicates via SQL with S-TAP server
 - Does not requires DB2 Connect

STAP for z/OS - IP Streaming (SQL Event Collection)

Information Management



System View Administration Console Tools Daily Monitor Guardium Monitor Tap Monitor Incident Management v.8.0 Custom Reports **G2000 - Standalone Unit**

S-TAP Status Monitor

Aliases: ON

S-TAP Host	S-TAP Version	DB Server Type	Status	Last Response Received	Primary Host Name	KTAP	TEE	MSS Shm	Win DB2 Shm	Win Local TCP	Pipes	Encrypted?	Firewall Installed
10.10.9.60	STAP-guard-8.0.xx_r20992_1-20100818_0518	INFORMIX	Inactive	2010-09-10 13:57:19	10.10.9.248	Yes	No	No	N/A	N/A	No	Unencrypted	Yes
DEMOMVS	IBM_5655-STP - IBM Infosphere Guardium S-TAP for DB2 on z/OS_Version 8.1		Active	2011-02-21 09:37:08	187.119.17.10	No	No	No	No	Yes	No	0	No

Records 1 to 2 of 2

```
Display Filter View Print Options Help
-----
SDSF OUTPUT DISPLAY ADHCDSNC STC01572 DSID 102 LINE 0 COLUMNS 02- 133
COMMAND INPUT ==> _ SCROLL ==> CSR
***** TOP OF DATA *****
ADHK101I Compiling filter. flags1 0xc0 trace=0 runtimeTraceFlag 0 runtimeTrace 0, zipRequested 0, supportingStage1 0
2011-01-25-10.46.27.232954: ADHG000I--Attempting connection to server 9.39.66.103 port=16016
2011-01-25-10.46.27.235216: ADHG001I--Establishing connection
2011-01-25-10.46.27.451165: ADHG002I--Connection established
***** BOTTOM OF DATA *****
```

Green is good, red is bad. Look for status of Active

STAP inactive can trigger alert processing in the z2010

STAP for z/OS - IFI (FTP) Collection

Information Management



Agent Editor

Status

Data storage mode The IFI audit data is written to either standard data sets or HFS files, depending on the selected Offload file type.

General settings

Offload dataset

Summary

Offload Files Type

Write Files to Data Sets

Data set HLQ:

Write Files to HFS

HFS directory:

Finish Cancel

Guardium for z/OS Interface Definition Builder ?

Server Ip

Server Name

Directory ←

User

Password

SSID

File Suffixes

Transfer Method ▼

Use END_USER_ID For Application User Name

Active

Remove Files

Last File Name

Back Add Comments Apply

AME Files

Start Date: 2011-01-22 09:55:19 End Date: 2011-03-02 09:55:19

Aliases: ON FileName: LIKE %

FileStatus: LIKE %

Interface ID	UA File Name	UT File Name	UH File Name	File Status	Total Number Of Events Processed	Number Of Events Failed	Timestamp
3	DSNC.UA.D110201.H11.M04.Z4	DSNC.UT.D110201.H11.M04.Z4		File Removed	8	0	2011-02-01 11:35:03.0
3	DSNC.UA.D110201.H15.M51.Z1	DSNC.UT.D110201.H15.M51.Z1		File Removed	3	0	2011-02-01 16:35:05.0
2	DSNB.UA.D110201.H16.M12.Z1	DSNB.UT.D110201.H16.M12.Z1		File Removed	3	0	2011-02-01 16:45:03.0
2	DSNB.UA.D110201.H16.M37.Z1	DSNB.UT.D110201.H16.M37.Z1		File Removed	3	0	2011-02-01 17:10:03.0
2	DSNB.UA.D110201.H16.M44.Z1	DSNB.UT.D110201.H16.M44.Z1		File Removed	3	0	2011-02-01 17:15:03.0
2	DSNB.UA.D110201.H16.M49.Z1	DSNB.UT.D110201.H16.M49.Z1		File Removed	3	0	2011-02-01 17:20:04.0

Guardium for z/OS - Audit Reporting

Information Management



Query Builder - Mozilla Firefox: IBM Edition

9.39.66.103 https://9.39.66.103:8443/queryBuilderDirectOpen.do?cmd=querySelected&selectedQuery=DSNB+Audit+Report+2&selectedQueryIndex=20022

Entity List

- Client/Server
- Server
- IP/Server Port
- Session
- Access Period
- SQL
- Full SQL

DSNB Audit Report 2

Main Entity: **Command** Add Count Add Distinct Sort by count

Seq.	Entity	Attribute	Field Mode	Order-by	Sort Rank	Descend
<input type="checkbox"/> 1	Session	Timestamp	Value	<input type="checkbox"/>		
<input type="checkbox"/> 2	Client/Server	Client IP	Value	<input type="checkbox"/>		
<input type="checkbox"/> 3	Client/Server	Server IP	Value	<input type="checkbox"/>		
<input type="checkbox"/> 4	Client/Server	DB User Name	Value	<input type="checkbox"/>		

DSNB Audit Report 2

Start Date: **2011-02-09 10:00:59** End Date: **2011-02-22 10:00:59**
 Aliases: **OFF** BindInfo: **LIKE %**
 OSUser: **LIKE SYS248** SQL_Verb: **LIKE %**
 server_IP: **LIKE %**

Timestamp	Client IP	Server IP	DB User Name	OS User	SQL Verb	Full Sql	Bind Variables Values	Host Variables	Bind Info	Plan Name	Timestamp	Application User
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	N/A		0			2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	SELECT T.*,HEX(T.HIGH2KEY),HEX(T.LOW2KEY),HEX(HIGH2KEY) AS HIGH2KEX, HEX(LOW2KEY) AS LOW2KEYX FROM SYSIBM.SYSCOLUMNS T WHERE T.NAME LIKE 'SSN%' ORDER BY TBCREATOR, TBNAME, COLNO FOR FETCH ONLY		52	PLAN=ADB	2011-02-16 06:40:36.0	2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	N/A		0			2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	N/A		0			2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	SELECT * FROM "DB2ADMIN"."CARDHOLDER" FOR FETCH ONLY		6	PLAN=ADB	2011-02-16 06:41:02.0	2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	SELECT * FROM "DB2ADMIN"."CREDIT_CARD" FOR FETCH ONLY		5	PLAN=ADB	2011-02-16 06:41:35.0	2011-02-16 12:19:29.0	
2011-02-16 12:19:29.0	127.0.0.19.39.68.147		SYS248	SYS248	SELECT	SELECT * FROM "DB2ADMIN"."CARDHOLDER_CREDIT_CARD" FOR FETCH ONLY		6	PLAN=ADB	2011-02-16 06:43:02.0	2011-02-16 12:19:29.0	

Vulnerability Assessment - DB2 on z/OS

Information Management



IBM® InfoSphere™ Guardium®



Results for Security Assessment: **z/OS DSNB VA**

Assessment executed **2010-12-14 16:02:03.0**

From: **2010-12-13 16:02:03.0**

To: **2010-12-14 16:02:03.0**

Client IP or IP subnet: **Any**

Server IP or IP subnet: **Any**

-- Select a result --

Download PDF

Tests passing: **0%***

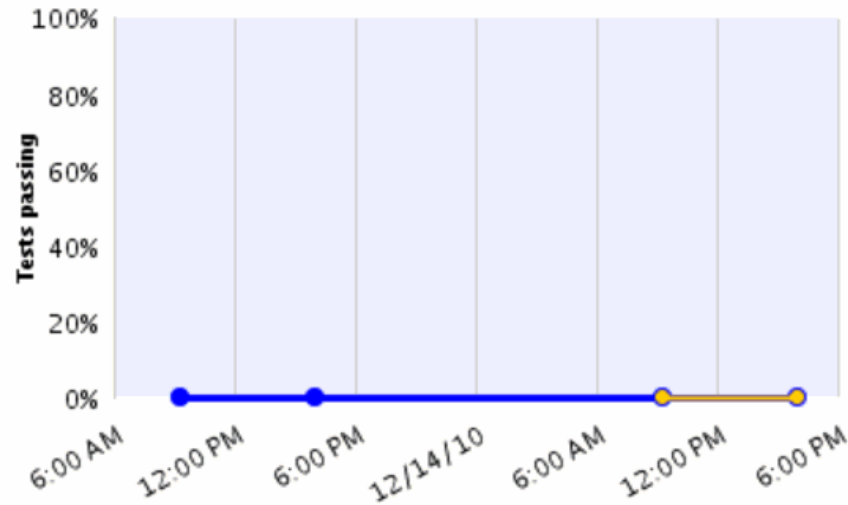
*Percentage does not take into account any current filtering

Based on the tests performed under this assessment, data access of the defined database environments requires significant improvement across a number of areas. Refer to the recommendations of the individual tests to learn how you can address problems within your environment, focusing on severe issues first. Continue running repeats of this assessment with every issue you address to track improvement.

[View log](#)

[Jump to Datasource list](#)

Assessment Result History



Result Summary *Showing 1 of 1 results (0 filtered)*

	Critical	Major	Minor	Caution	Info
Privilege	1f				
Authentication					
Configuration					

Current filtering applied:

Test Severities: - [Show All](#) -


Datasource Severities: - [Show All](#) -

Scores: - [Show All](#) -

Types: - [Show All](#) -

Why is Guardium Needed With TSIEM and zSecure?

- Traditional approaches to database security and compliance rely upon native logging, which:
 - Does not meet auditors requirements for separation of duties
 - Can easily be circumvented by DBAs
 - Imposes a higher performance overhead on database servers
 - Does not provide real-time alerting or blocking
 - Does not enable enforcement of consistent policies in heterogeneous environments



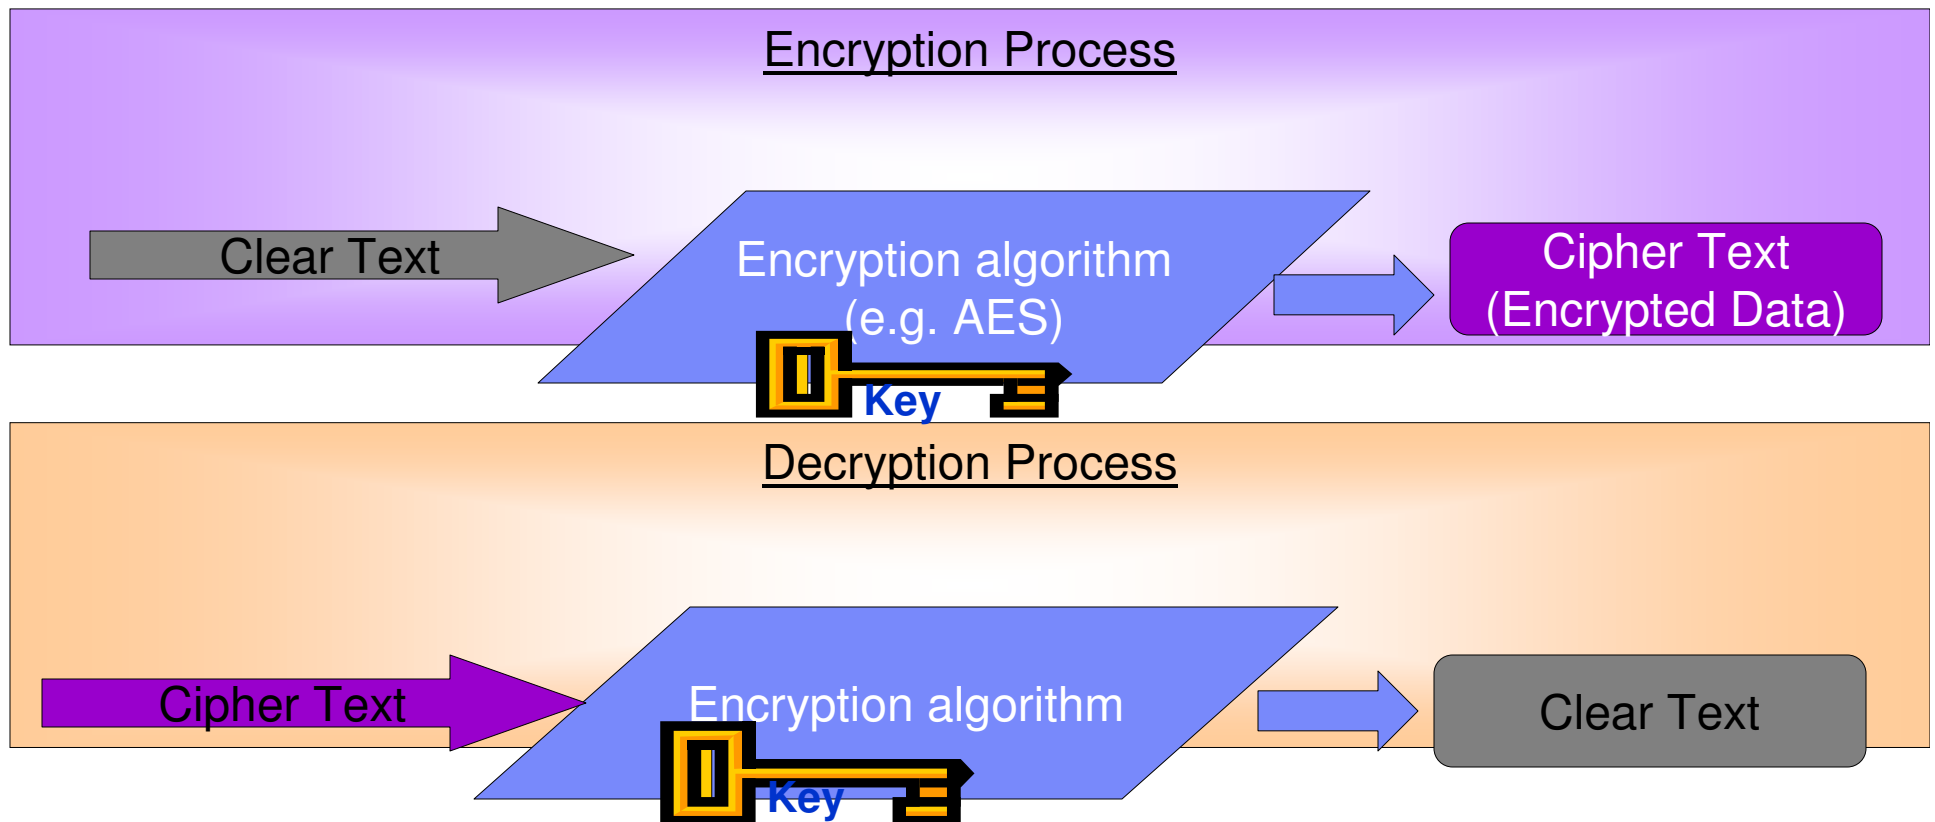
Encryption and “Data at Rest” Protection

- Key requirement for most of the “popular” data protection initiatives
- Main requirement is to protect “data at rest” to ensure that only access if for business need-to-know, and through mechanisms which can be controlled by the native security mechanisms (such as RACF)
- Consider the following scenario:
 - DB2 Linear VSAM datasets are controlled via RACF from direct access outside of DB2 via dataset access rules
 - DBA or Storage Administrator has RACF authority to read VSAM datasets in order to perform legitimate storage administration activities.
 - Administration privileges can be abused to read the linear VSAM datasets directly and access clear-text data outside of DB2/RACF protections.
- Now consider the above scenario, but with the underlying Linear VSAM datasets encrypted
 - When DBA or Storage Administrator uses their RACF dataset authorities in a manner which is outside of business need-to-know, the data retrieved is cybertext and thus remains encrypted and protected.
 - Only way to access and obtain clear-text data will be via SQL which can be protected via DB2/RACF interface



Encryption and DB2 for z/OS

- IBM Data Server Drivers starting in V9.5 support SSL protocol and AES encryption.
- Starting with Fix Pack 2, non-Java clients supports the Secure Sockets Layer (SSL) protocol. All DB2 Version 9.5 clients now support SSL. In addition, Java and CLI clients now support 256-bit AES encryption.
- SSL connectivity and AES user ID and password encryption requires Communication's AT-TLS configured and ICSF started.
- Starting with DB2 for z/OS V8, column level encryption implemented via SQL primitives is supported. TDES 128 bit support only.
- Row level encryption implemented for all supported releases of DB2 for z/OS using the IBM Infosphere Guardium Encryption Tool for IMS and DB2 databases
- DS8000 family DASD Based Encryption
- TS1120/TS1130 Tape Based Encryption
 - TKLM (Tivoli Key Lifecycle Manager) Required for DS8000 and recommended for TS1120/TS1130

Encryption is a technique used to help protect data from unauthorized access



- Data that is not encrypted is referred to as “clear text”
- Clear text is encrypted by processing with a “key” and an encryption algorithm
 - Several standard algorithms exist, include DES, TDES and AES
- Keys are bit streams that vary in length
 - For example AES supports 128, 192 and 256 bit key lengths



CP Assist for Cryptographic Functions

- Problem State Instructions
- Clear Keys Only
- DES/TDES Encryption
- AES (128 Bit)
- SHA-1 (256 on z9)

Crypto Express 2 Coprocessor

- ICSF Access Only (Key 0)
- Master Key Stored Within Boundary of Crypto Express 2 Feature
- Secure Key DES/TDES Encryption
- SSL Accelerator
- Tamper Resistant

System z9/z10 Cryptographic Support Summary

CP Assist for Cryptographic Function (CPACF) “free”

- Supports DES, TDES and SHA-1
- Standard on System z9/z10 (feature code 3863)
- Standard on every CP and IFL
- Advanced Encryption Standard (AES)
- Secure Hash Algorithm – 256 (SHA-256)
- Pseudo Random Number Generation (PRNG)




Crypto Express2 (feature code 0863) “fee”

Crypto Express3 (feature code 4863) “fee”

- Two configuration modes
- Coprocessor (default)
- Federal Information Processing Standard (FIPS) 140-2 Level 4 certified
- “Tamper Resistant”
- (Secure Key) – “Exclusive”
- SSL Accelerator (Handshake offload)

Three configuration options

- Default set to Coprocessor (1)
- SSL Acceleration (3)
- Mixture of configuration (2)



Encryption - key forms on z/OS

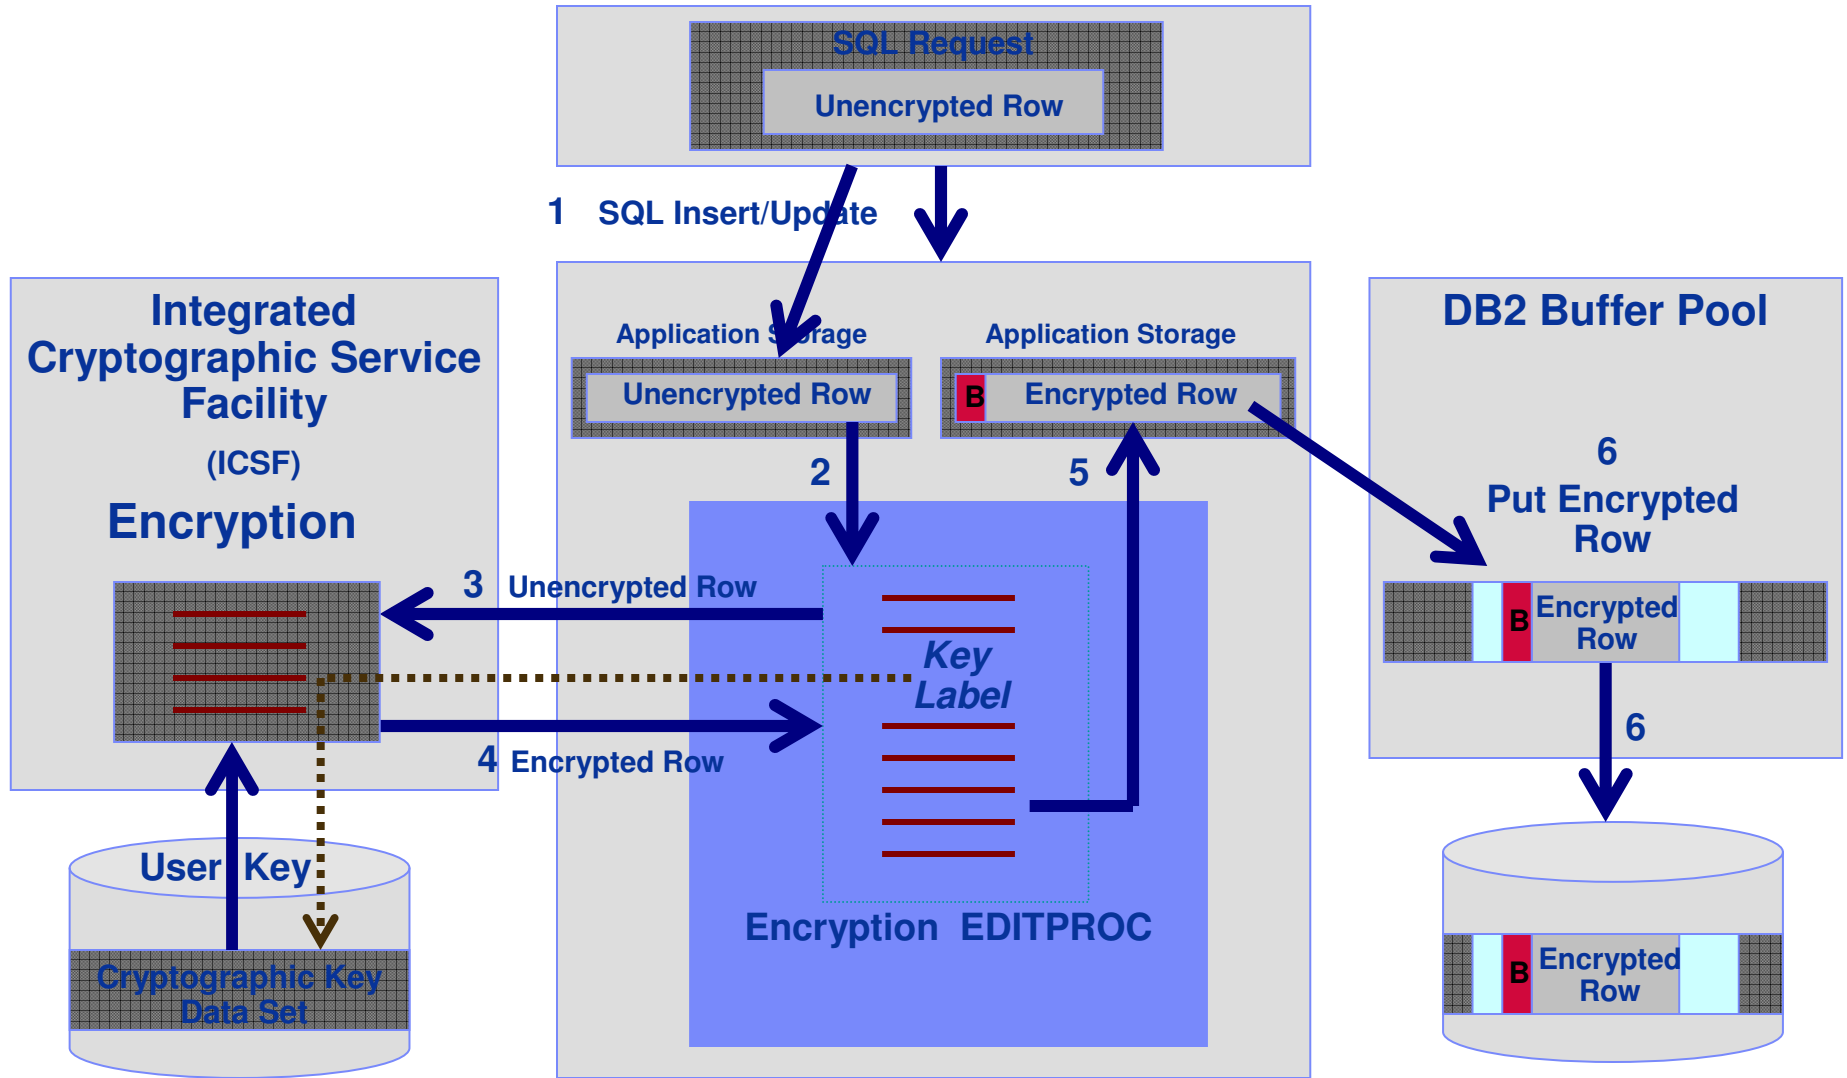
	zSeries 900	System z9 & z10		System z10
	CCA Secure Key	Clear Key	CCA Secure Key	CPACF Protected Key
Key Wrapping: Host Storage	CCA Master Key – Key material is never visible in the clear outside the tamper resistant hardware boundary	None – Key material is visible in the clear in system and application storage .	CCA Master Key – Key material is never visible in the clear outside the tamper resistant hardware boundary	CPACF Wrapping Key – Key material is not visible in the clear in <i>operating system or application storage.</i>
Key Wrapping: Key Store	CCA Master Key – Key material is never visible in the clear outside the tamper resistant hardware boundary	None – Key material is visible in the clear key store.	CCA Master Key – Key material is never visible in the clear outside the tamper resistant hardware boundary	CCA Master Key – Key material is never visible in the clear outside the tamper resistant hardware boundary
Key Store	CKDS or <i>application key file</i>	CKDS or <i>application key file</i>	CKDS or <i>application key file</i>	CKDS only
Encryption Engine	CCF	CPACF or software	CEX2C	CPACF
Symmetric Encryption Algorithms	DES and TDES	DES, TDES and AES	DES, TDES and AES	DES, TDES and AES
Benefits	High Security	High Performance	High Security	High Performance High Security

Infosphere Guardium Encryption Tool for IMS and DB2 Databases

- **Generates standard DB2 EDITPROC for Accessing Cryptographic Functions**
 - **All Supported DB2 Versions**
 - **Member of IBM IMS | DB2 Tools Family of Products**
 - **Pre-coded EDITPROC for encryption of DB2® Data**
 - **Encryption/Decryption occurs at the DB2 Row Level**
 - **Unique EDITPROC can be defined for each DB2 Table**
 - **Exploits z/OS Integrated Cryptographic Service Facility (ICSF)**
 - **Exploits zSeries CPACF Cryptographic Hardware Directly**
 - **Requires no changes to your applications**
 - **Fast implementation**


- **Edit Procedures (EDITPROC) are Programs that:**
 - **Transform Data on INSERT | UPDATE | LOAD**
 - **Restore Data to Original Format on SELECT**
 - **Transformations on Entire ROW**
 - **Supported by Utilities**
 - **Implemented via Create Table specification**
 - **Requires unload/load of data**

DB2 Data Encryption Flow – Insert / Update



zIIP Assisted IPsec (VPN) on z/OS

- **Benefits of having secure channel end-point on z/OS**
 - No clear-text data on any network segments
 - Security regulations compliance
 - End-to-end authentication of secure channel end-points
 - Both end-point authentication and message authentication
 - Key management and storage done on System z by z/OS
 - Compliance with end-to-end security regulations
- **System z CPU cost is a concern**
 - Encryption/decryption CPU cost can be a significant percentage of overall CPU cost for a given application
 - Especially the case for streaming workloads (file transfer type of workload)
- **zIIP processors**
 - Specialty processor on System z9 or later hardware
 - zIIPs priced lower than general purpose processors
 - No IBM software charges on zIIPs
- **zIIP Assisted IPsec**
 - Use zIIP processors for most IPsec encryption/decryption
 - Lower the cost of doing IPsec processing on z/OS

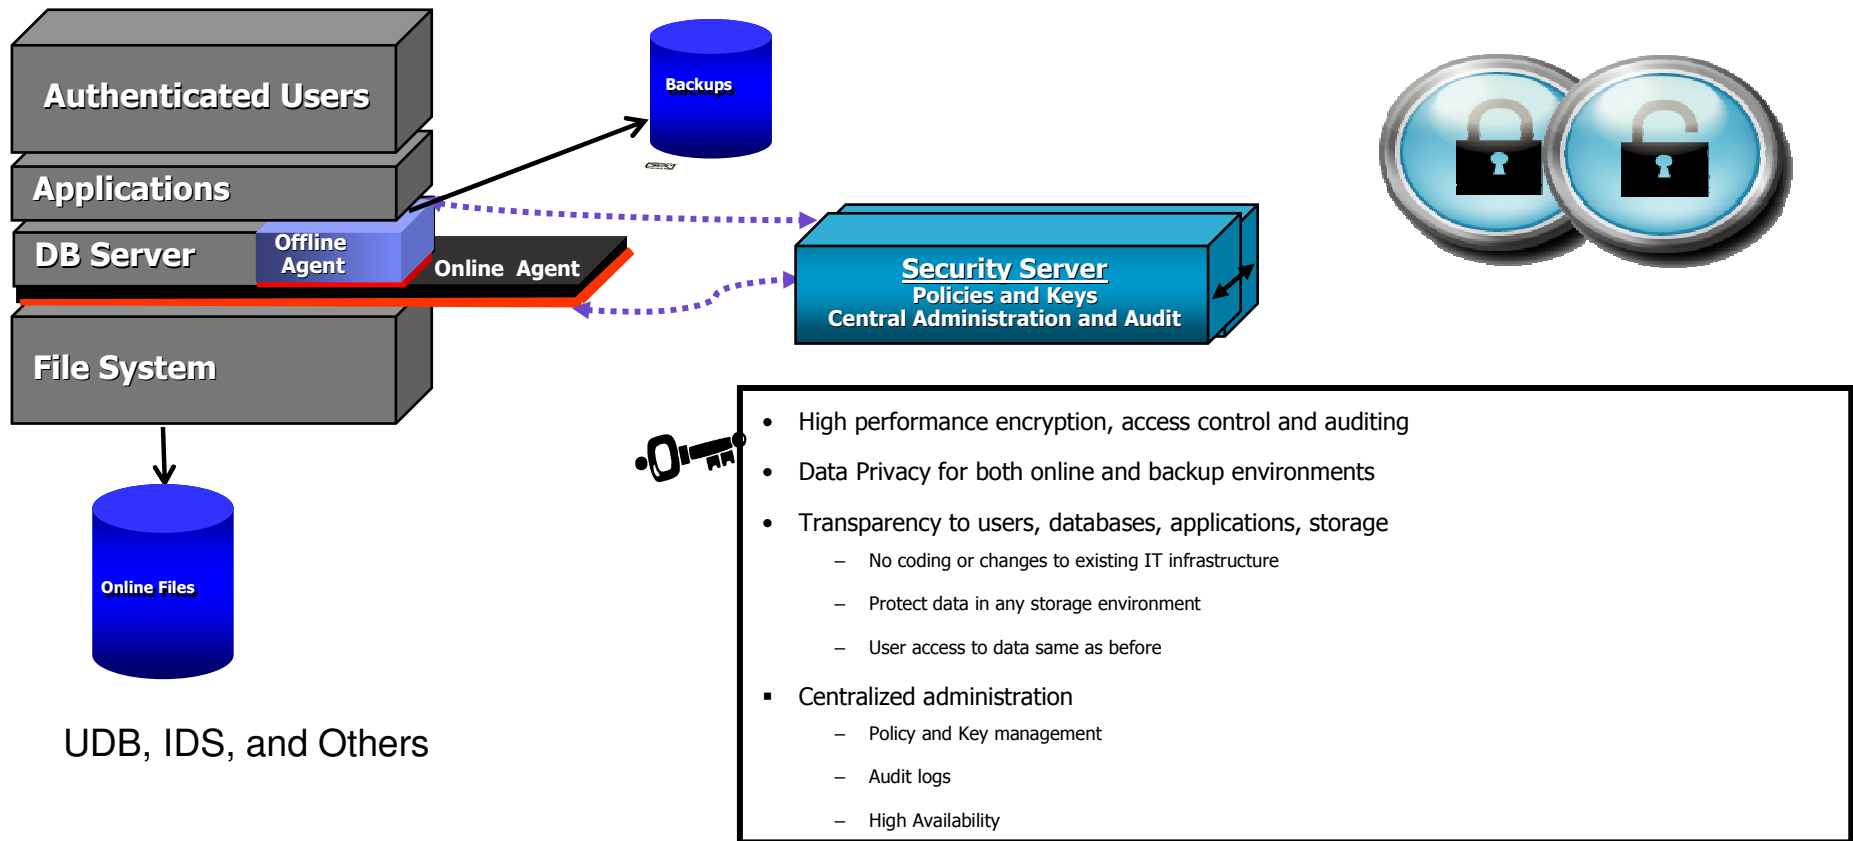


System z9 or later
z/OS CS V1R8 + PTFs
z/OS CS V1R9

IBM DS8000 Disk Encryption - Characteristics

- Customer data at rest is encrypted
 - Data at rest = data on any disk or in any persistent memory
- Customer data in flight is not encrypted
 - Data in flight = on I/O interfaces or in dynamic memories (Cache, NVS)
 - If you can read/write to disk, you get access to clear-text data.
- Uses Encrypting Disk
 - Encryption hardware in disk (AES 128)
 - Runs at full data rate
 - 146/300/450 GBs 15K RPM
 - No measurable performance impact
- Integrated with Tivoli Key Lifecycle Manager (TKLM)
 - DS8000 automatically communicates with TKLM when configuring encryption group or at power on to obtain necessary encryption keys to access customer data
 - Each disk has an encryption key
 - Data is always encrypted on write and decrypted on read
 - Encryption key is wrapped with access credential and maintained within the disk
 - Access credential maintained by TKLM
 - Establishing a new encryption key causes cryptographic erasure
- Key attack vectors prevented:
 - Disk removed (repair, or stolen)
 - Box removed (retire, or stolen)

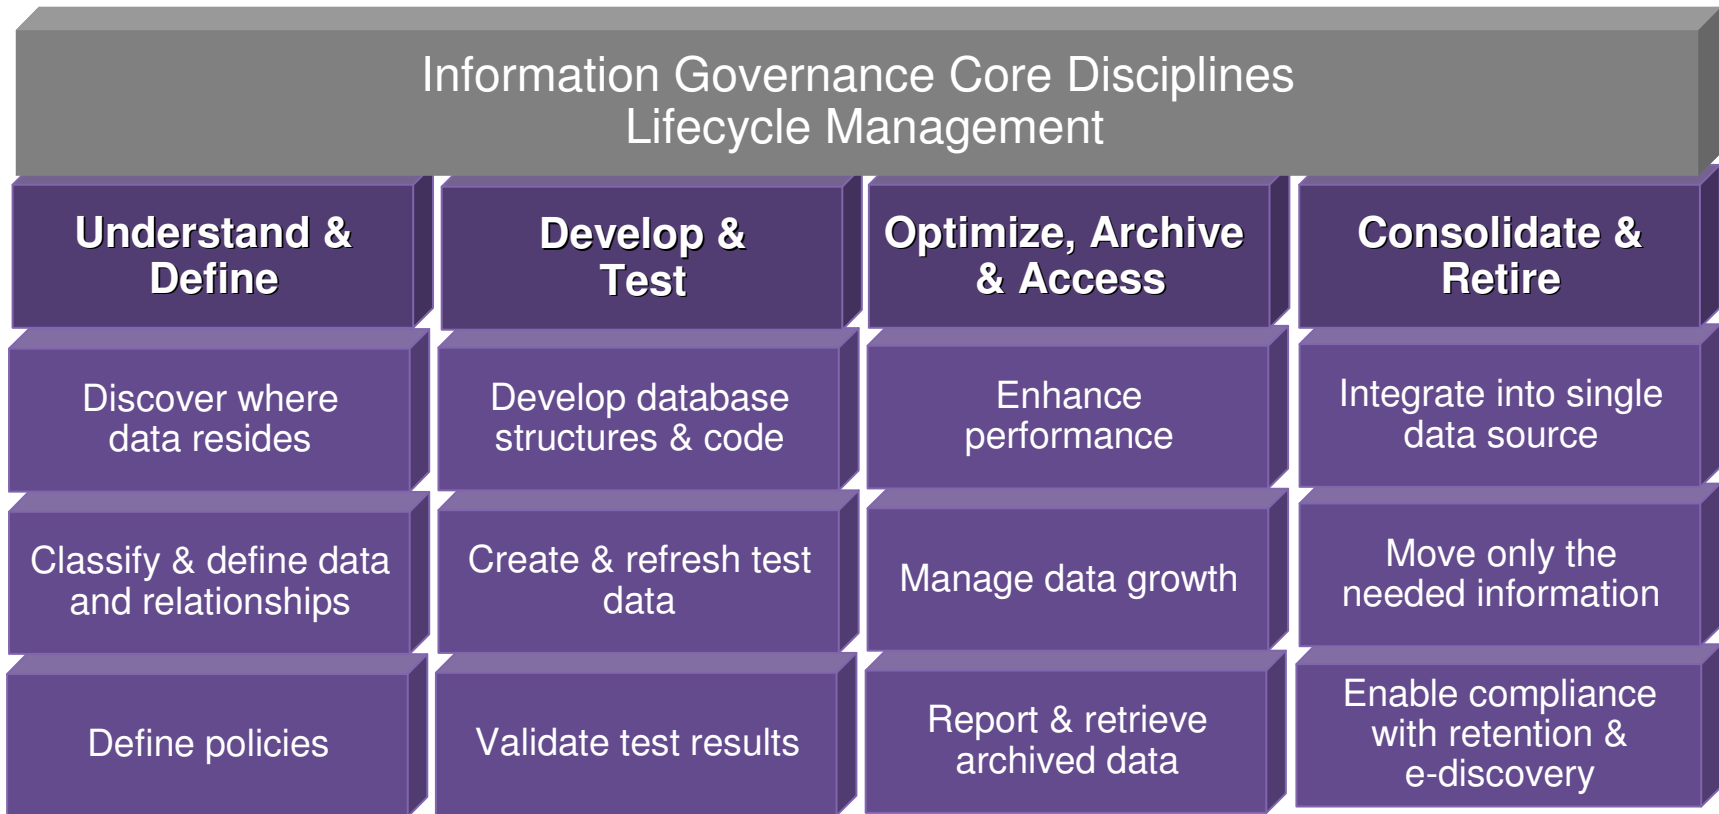
Optim Encryption Expert – Data Encryption



Discovery, Test Data Management/Obfuscation, and Data Growth



Requirements to manage data across its lifecycle



Slide 48

A1

Remove Access

Author, 11/23/2010

Challenges due to lack of information lifecycle management

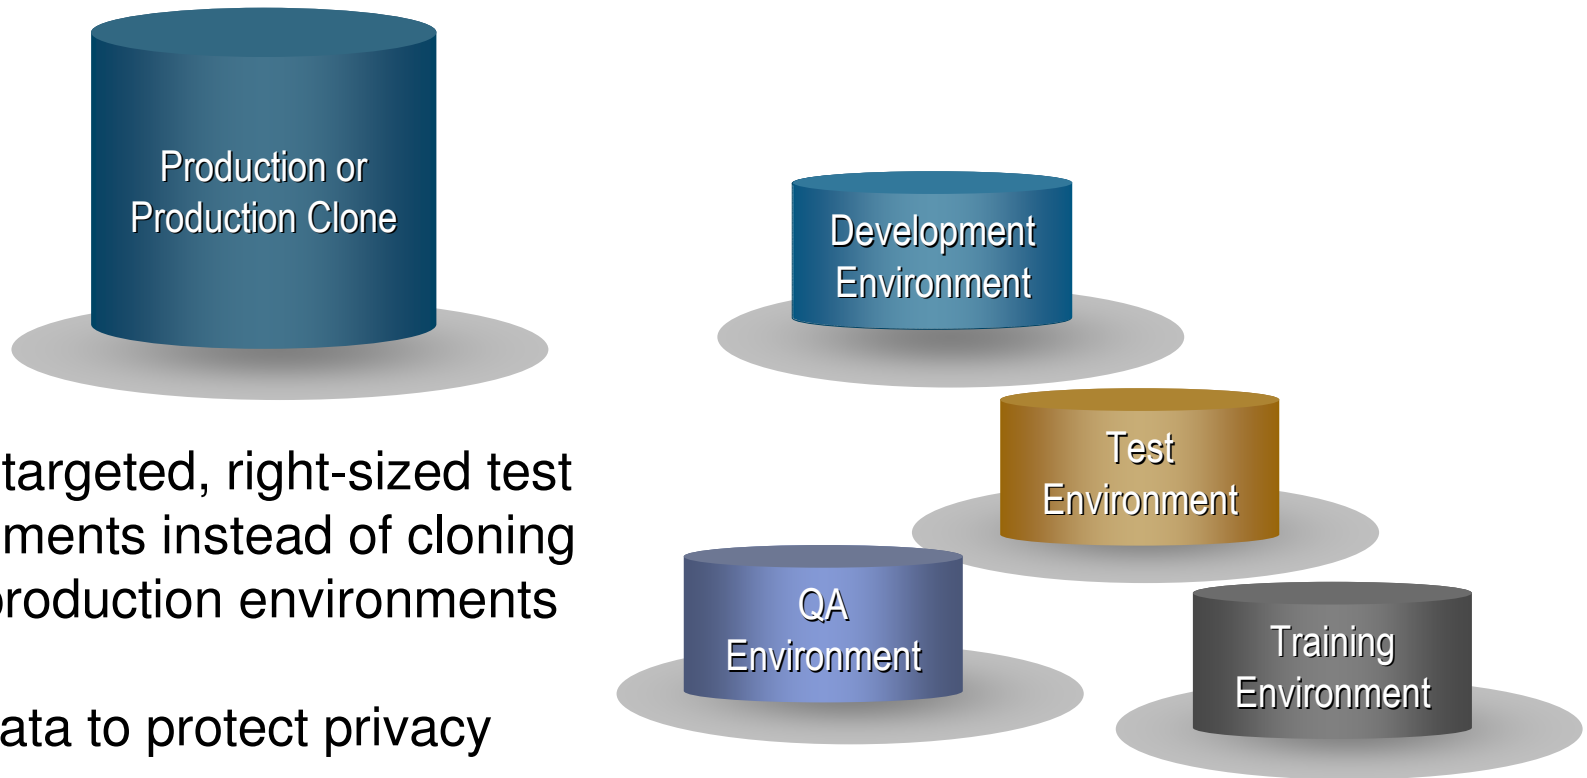
- New application functionality to meet business needs is not deployed on schedule
- Increased operational and infrastructure costs impact IT budget
- Application defects are discovered late in the lifecycle including after deployment
- Unintentional disclosure of confidential data kept in test/development environments

“

*Forrester estimates that 85%
of data stored in databases is inactive*

Source: Noel Yuhanna, Forrester Research, Database Archiving Remains An Important Part Of Enterprise DBMS Strategy, 8/13/07

Implement test data management with masking

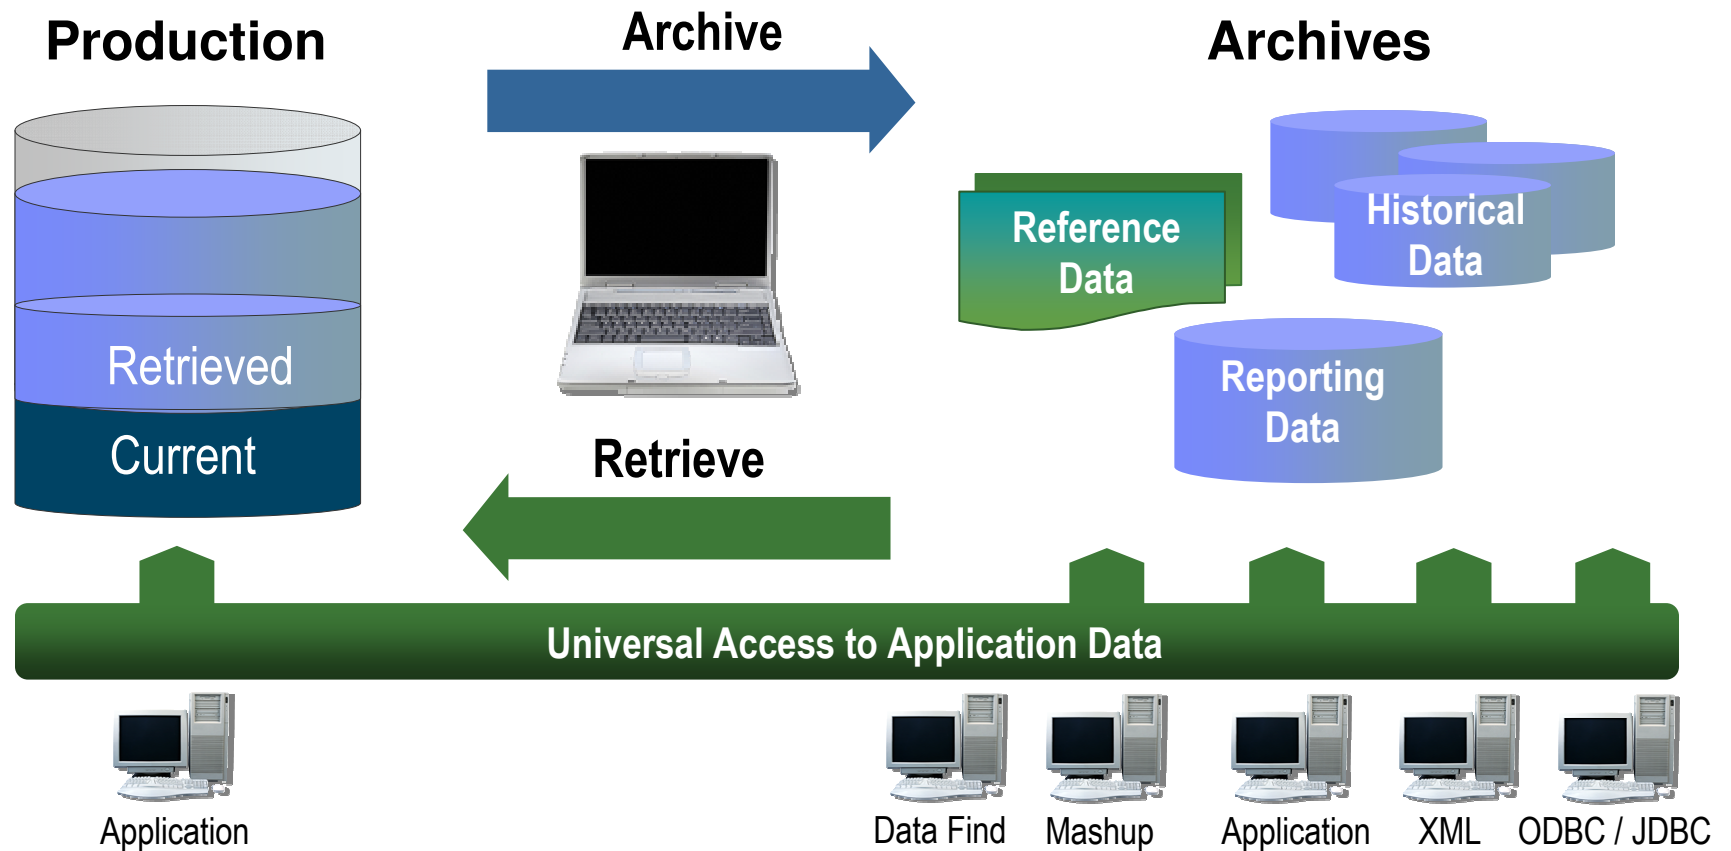


Create targeted, right-sized test environments instead of cloning entire production environments

Mask data to protect privacy

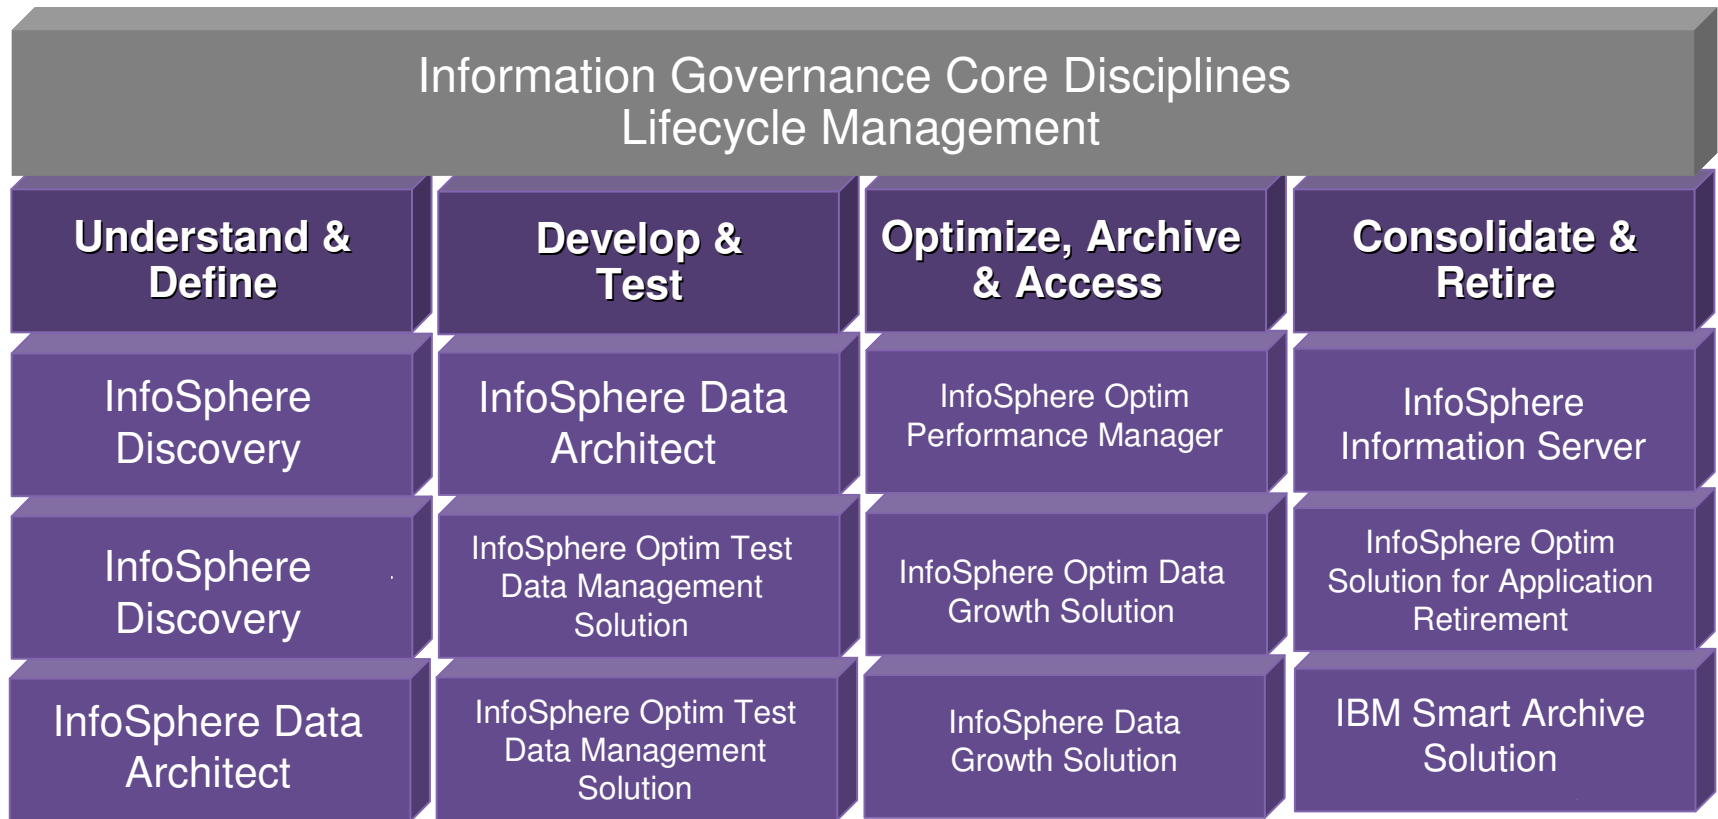
Compare data pre/post test to identify quality issues

Archive to manage data growth




Archiving is an intelligent process for ***moving*** inactive or infrequently accessed data that still has ***value***, while providing the ability to ***search and retrieve*** the data

IBM provides the solutions required to manage information throughout its lifecycle from requirement to retirement




Understand your information



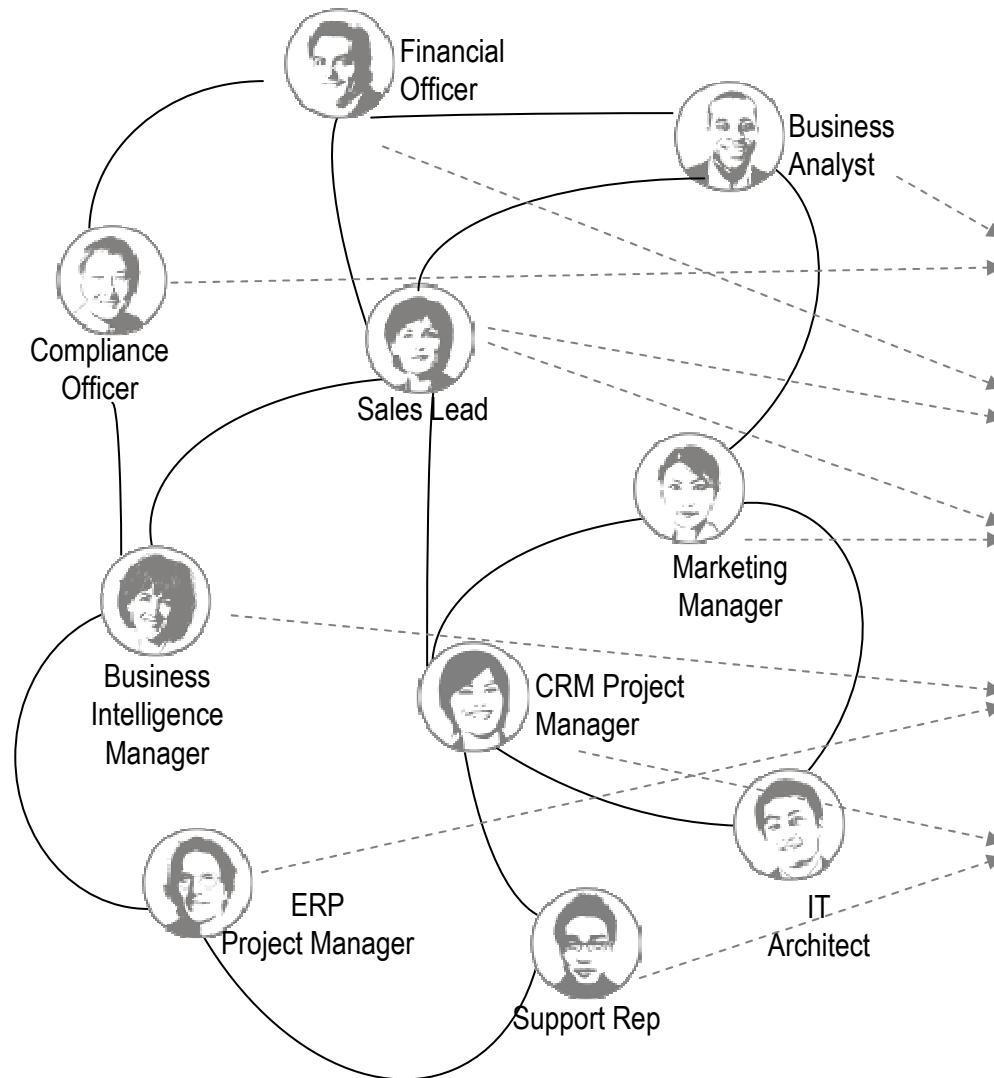
- Data can be distributed over multiple applications, databases and platforms
- Relationships are complex and poorly documented
- Relationships are not well understood

Automate Discovery and Accelerate Information Understanding

- Significant Acceleration of Information Agenda projects
 - Data Growth Management
 - Test Data Management
 - Sensitive Data De-identification
 - Application/Data Consolidation, Migration & Retirement
 - Master Data Management and Data Warehousing
- Why is this Different?
 - Data-based discovery
 - Automate discovery of business entities, cross-source business rules & transformation logic
 - Evaluate multiple data sources simultaneously
 - Identify & remediate cross-system rules and inconsistencies



Define a common vocabulary



For example, define

“Active Subscriber”

- **Mobile user who has used “any” service in the mobile network**
- **User who paid for the service at least 1 time in the past 90 days.**
- **Mobile user who has a phone plan, but not SMS**
- **Only post-paid customers, not pre-paid customers**
- **User who makes at least 1 call over the period of 90 days**

Discover where sensitive data may be hidden

Sensitive Relationship Discovery

System A Table 1		System A Table 15		
Number	Name	Patient	Result	Test
4600986	AlexFulltheim	3802468	N	53
8150000	BernardCole	4182715	N	53
6123913	Karalyn Jones	6123913	Y	47
5567193	EileenKratzman	5567193	N	72
7409934	FredSimpson	6736304	N	34
5061085	JamieSlattery	7409934	N	34
4182715	JimJohnson	8150928	N	47
8966020	MartinAston	8966020	N	34

System Z Table 25

Code	Name
53	Streptococcus pyogenes
72	Pregnancy
32	Alzheimer Disease
47	H1N1
34	Dermatamycoses

Compound sensitive data:
Test results could potentially be revealed.

- Relationships and sensitive data can't always be found just by a simple data scan
 - Sensitive data can be embedded within a field
 - Sensitive data could be revealed through relationships across fields & systems

- When dealing with hundreds of tables and millions of rows, this search is complex – you need the right solution

InfoSphere Discovery Components

Cross-Profiler

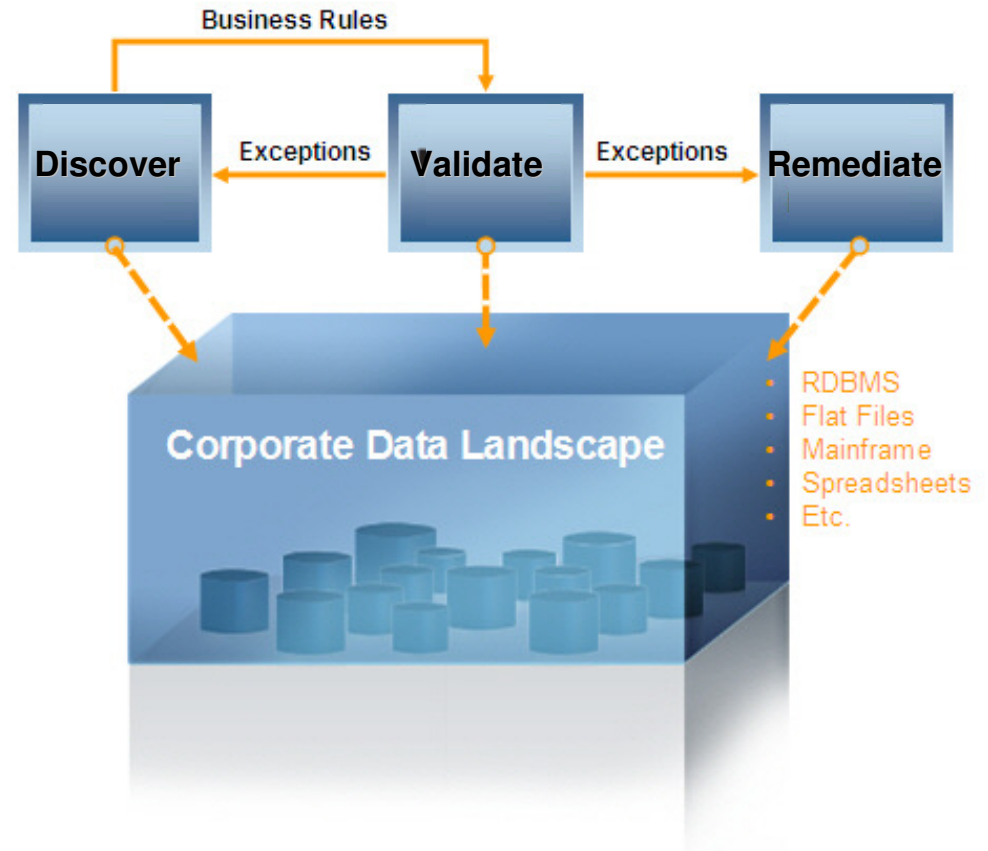
- Basic profiling plus automated primary-foreign key, business entity & cross-source overlaps discovery

Unified Schema Builder:

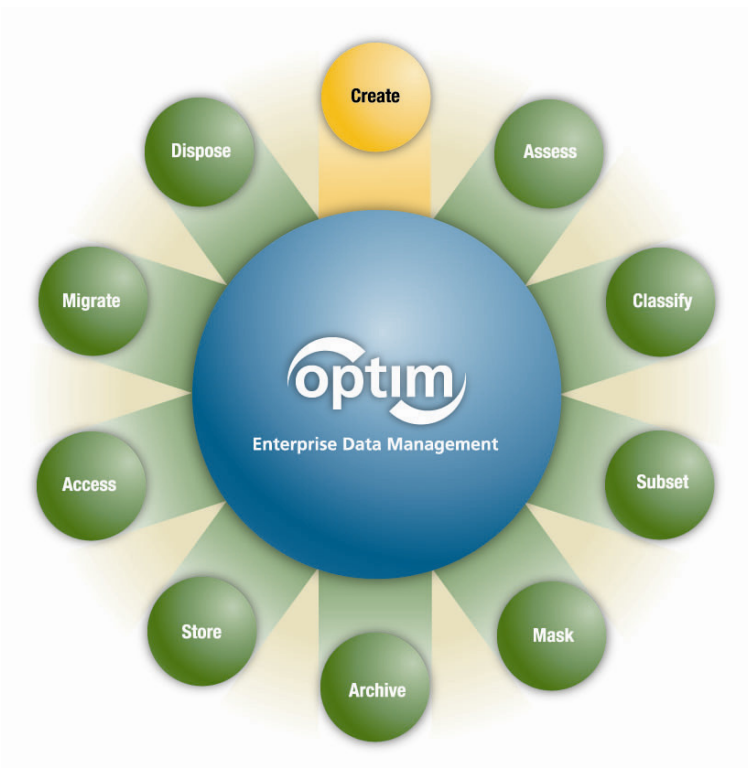
- Prototype empty targets from the combination of many data sources

Transformation Analyzer:

- Discover complex business rules and transformation logic between two data sources



Optim™ Solutions



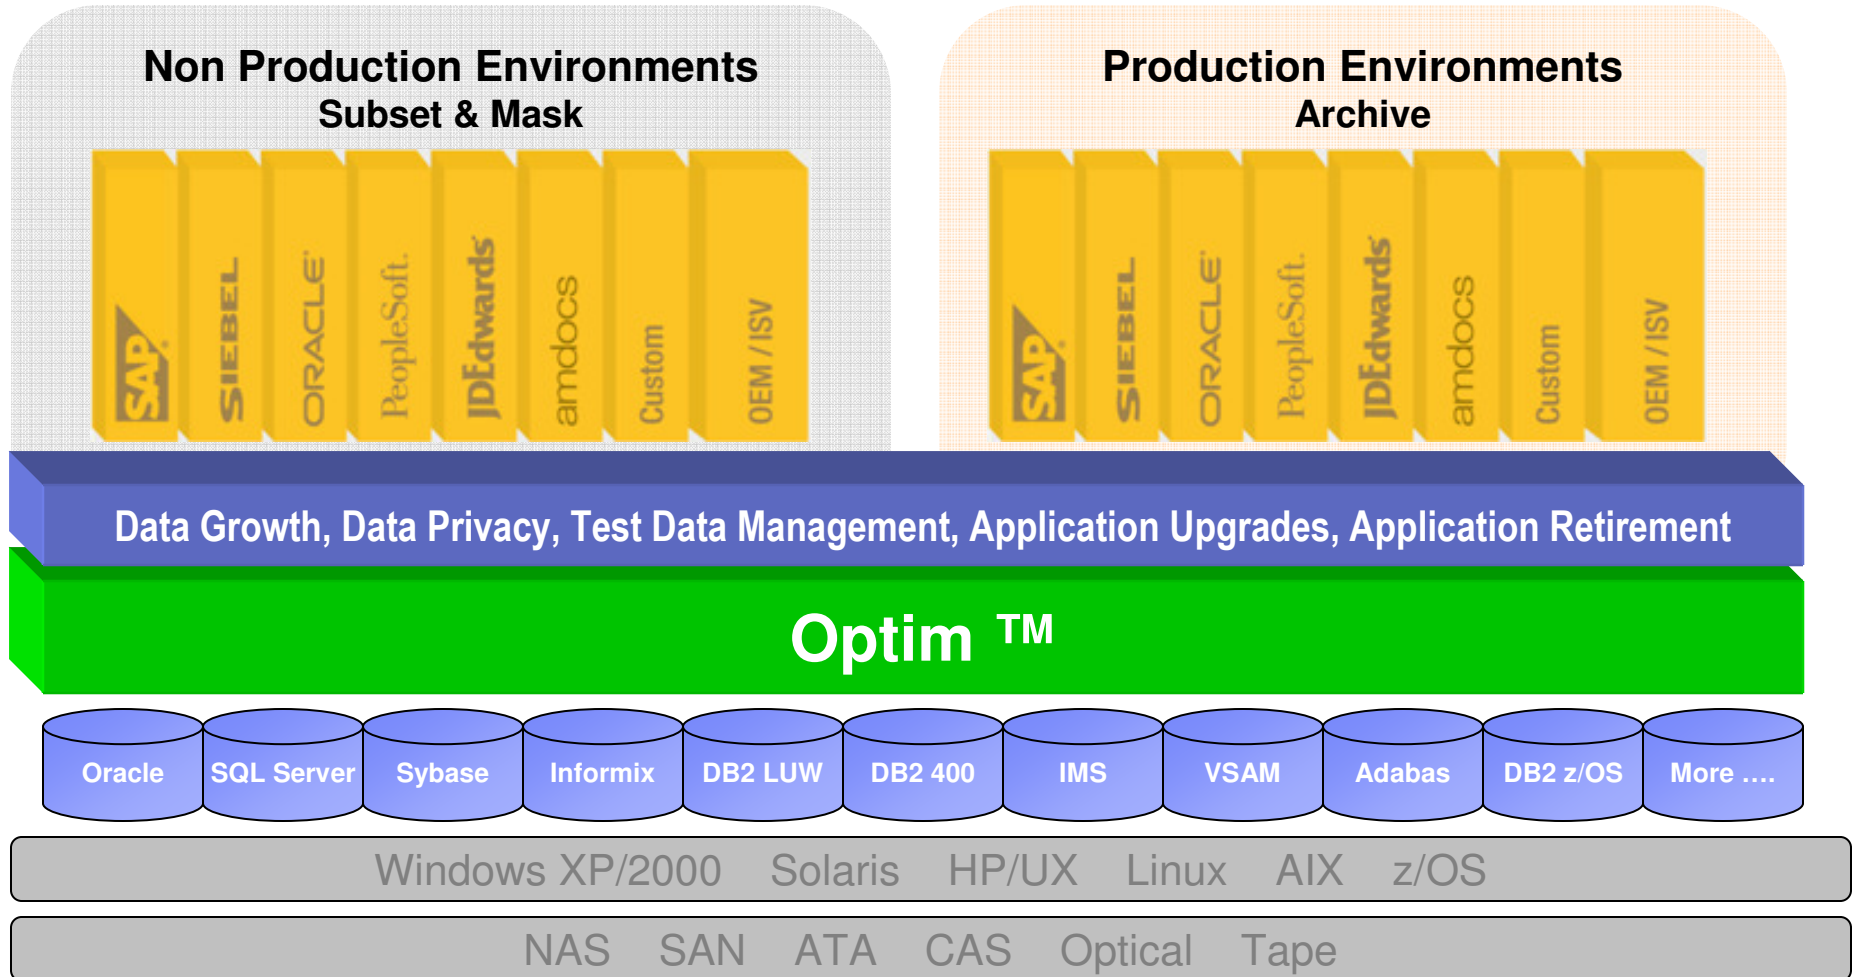
- Optim™ Data Growth Solution (Archiving)
 - Improve performance
 - Control data growth, save storage
 - Support retention compliance
 - Enable application retirement
 - Streamline upgrades

- Optim™ Test Data Management Solution
 - Create targeted, right sized test environments
 - Improve application quality
 - Speed iterative testing processes

- Optim™ Data Privacy Solution
 - Mask confidential data
 - Comply with privacy policies

- Enterprise Capabilities
 - Single, scalable solution for complex multi-DB application environments

Enterprise Architecture

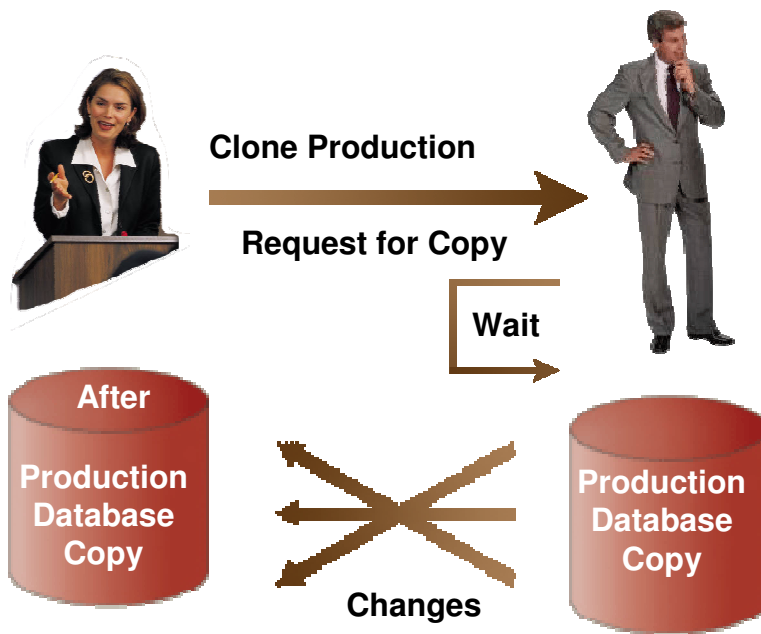


- Single, scalable, interoperable EDM solution provides a central point to deploy policies to extract, store, port, and protect application data records from creation to deletion


Current Practices?

#1 - Clone Production


Repeat ?*%\$!



Manual examination:
 Right data?
 What Changed?
 Correct results?
 Unintended Result?
 Someone else modify?



#2 - Write SQL




• RI Accuracy?
 • Right Data?

Expensive,
 Dedicated Staff,
 Ongoing
 Responsibility.

Share test database
 with everyone else

Data Multiplier Effect

Actual Data Burden = Size of production database + all replicated clones

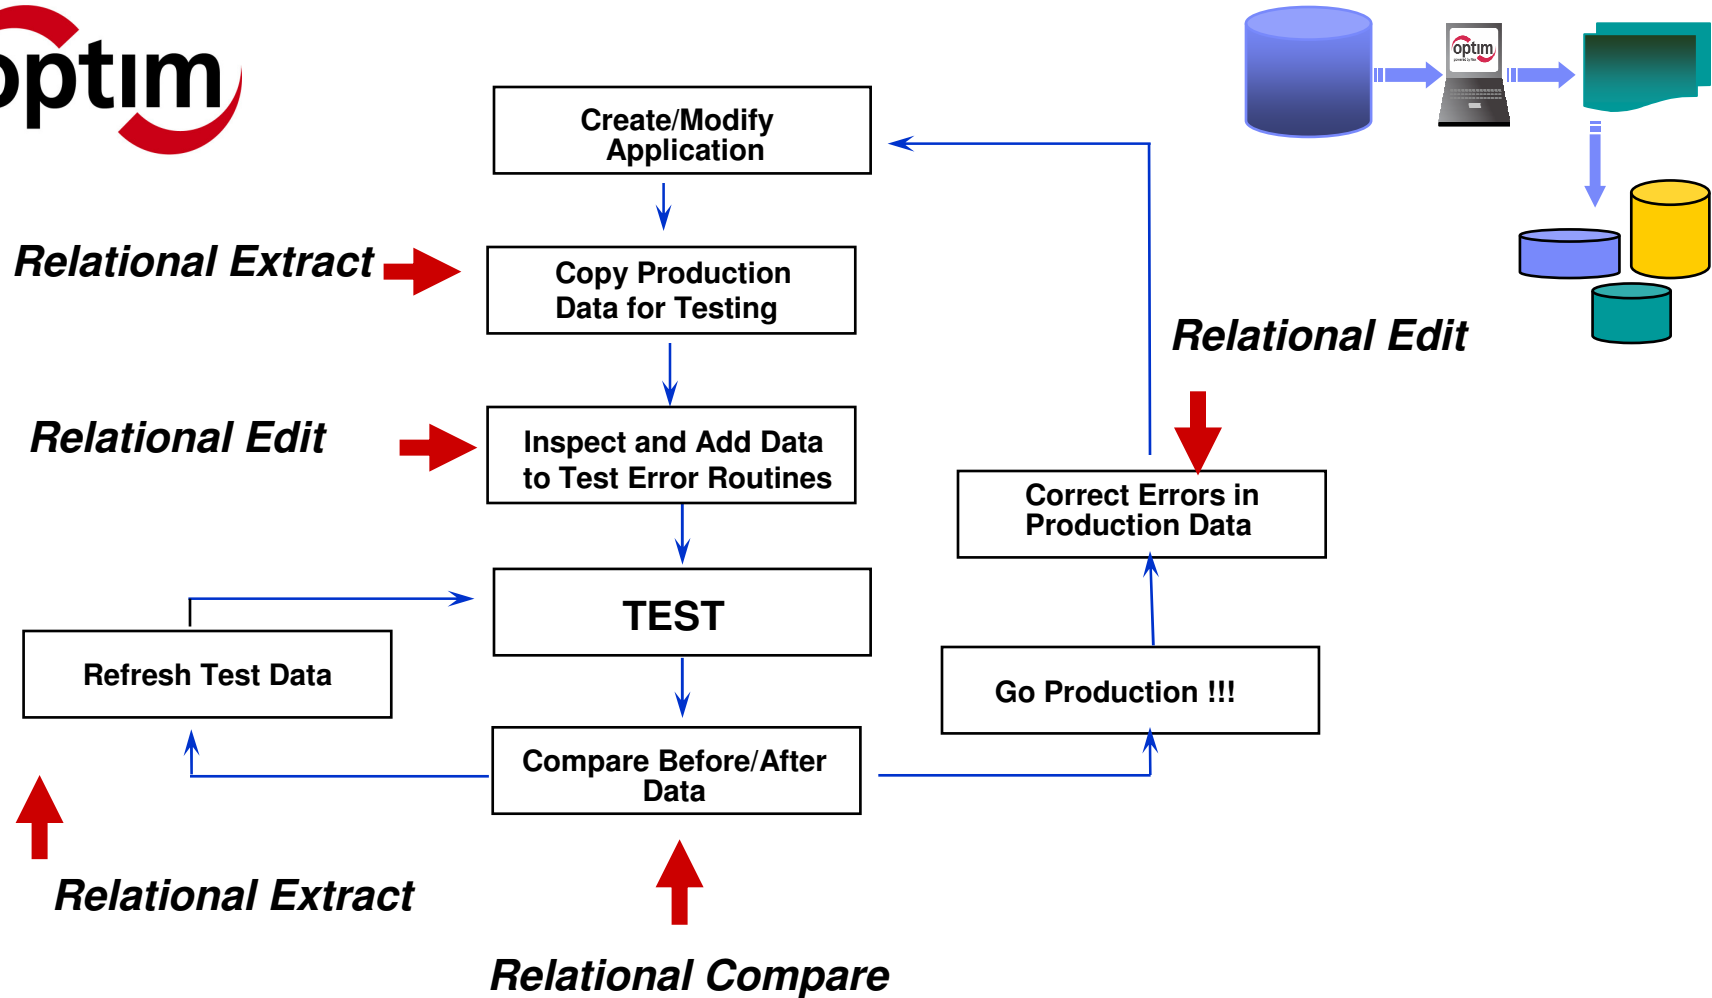


Optim Overview

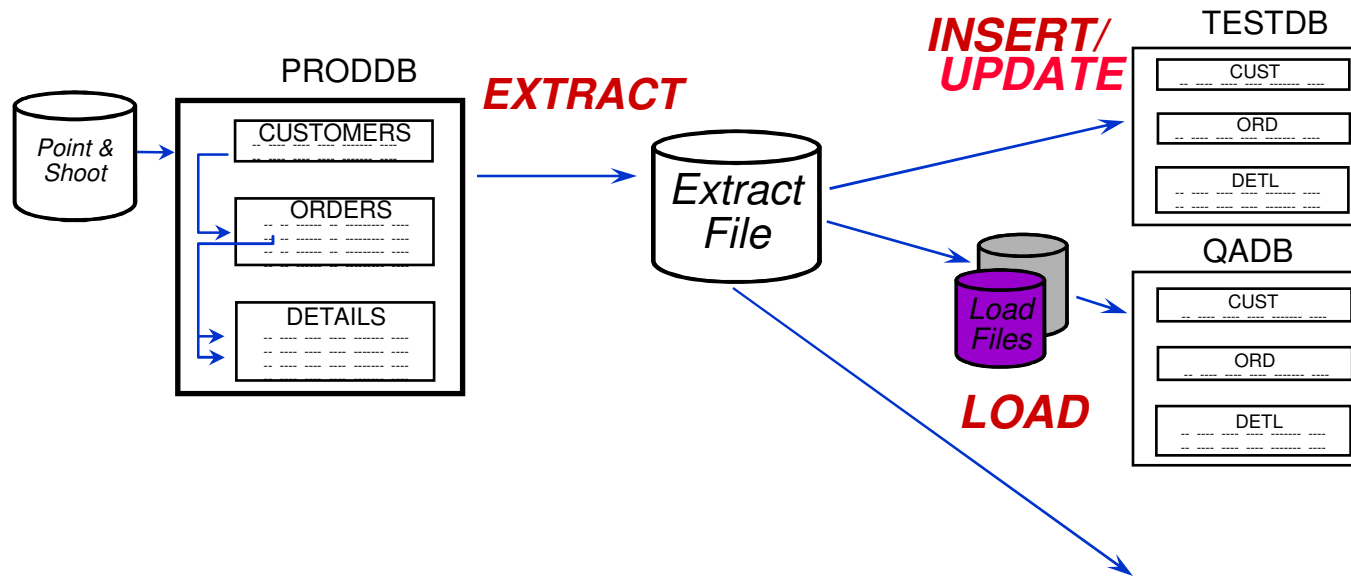
Relational Extract Facility / Test Data Management




Product Overview : Optim Test Data Management



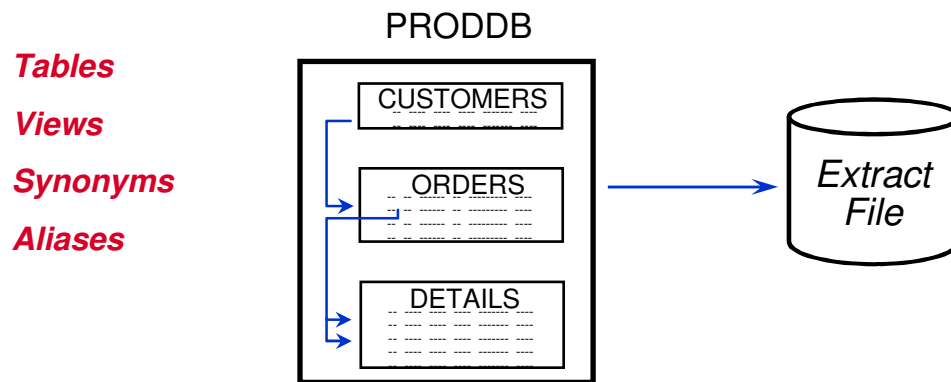
OPTIM Relational Extract Facility



- Creating and maintaining test data bases
- Migrating data
- The data and/or the object metadata can be extracted



Defining the Extract.....



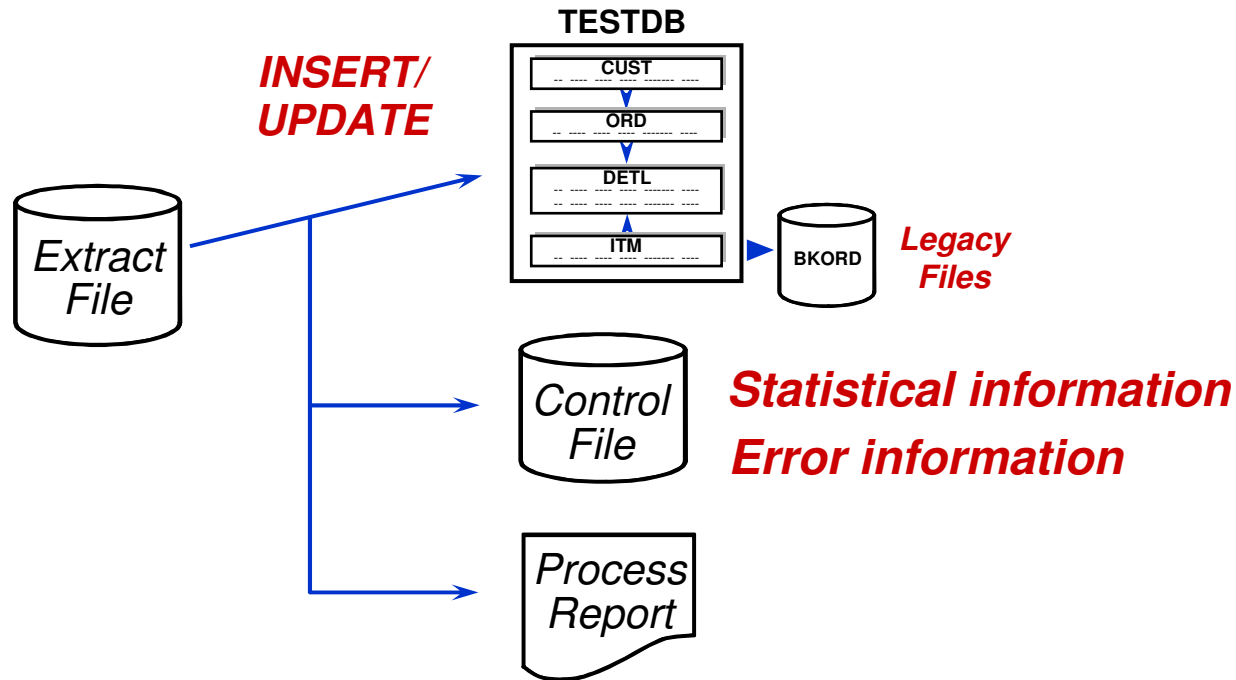
Required:

- **Start Table**
- **Set of Tables**

Optional:

- Selection Criteria
- Data Sampling
- Data Partitioning
- Point and Shoot
- Relationship Usage

Populate Destination Tables Control File



If INSERT/UPDATE errors occur:

1. **BROWSE** the control file for error information
2. **RETRY/RESTART** the INSERT/UPDATE

De-identify data without impacting test and development

- Mask or de-identify sensitive data elements that could be used to identify an individual
- Ensure masked data is contextually appropriate to the data it replaced, so as not to impede testing
- Support referential integrity of the masked data elements to prevent errors in testing



Personal identifiable information is masked with realistic but fictional data for testing & development purposes.

Protect sensitive data values within documents


- Redact (or remove) sensitive unstructured data found in documents and forms, protecting confidential information while supporting the need to share critical business information
- Leverage an automated redaction process for speed, accuracy and efficiency
- Prevent unintentional disclosure by using role-based masking to confidently share data
- Ensure multiple file formats are supported, including PDF, text, TIFF and Microsoft Word documents



Data Privacy in Application Testing

Only Users authorized to see Private data



Extract a relationally intact subset from production database(s)




Transform / mask sensitive data




**INSERT/
UPDATE**



LOAD




Sanitized Data



- Most Secure Approach
 - **Extract data only**
 - Convert during extract
- Extract file already contains masked data
- Can be shared with testers to reuse

Data Privacy in Application Testing


**Extract a relationally intact subset
from production database(s)**



Data transformation functions:

- 📄 Hard-code literals,
- 📄 special registers such as date, time
- 📄 Arithmetic calculations
- 📄 Sequential number generation
- 📄 Random number generation
- 📄 Substring and/or concatenation of values
- 📄 Lookup Table Functions Random, Specific or HASH
- 📄 Intelligent TRANSformation Library – SSN, CCN, email,...
- 📄 Access to client-defined exit routines to apply complex algorithms, encryption, ...
- 📄 Propagation of masked primary keys to dependent foreign keys

Propagating Keys



Without Key Propagation...

Original Data

Customers Table

Cust ID	Name	Street
08054	Alice Bennett	2 Park Blvd
19101	Carl Davis	258 Main
27645	Elliot Flynn	96 Avenue

Orders Table

Cust ID	Item #	Order Date
27645	80-2382	20 June 2004
27645	86-4538	10 October 2005

Without Key Propagation

Customers Table

Cust ID	Name	Street
10000	Auguste Smith	Mars23
10001	Claude Jones	Venus24
10002	Pablo Adams	Saturn25

Orders Table

Cust ID	Item #	Order Date
27645	80-2382	20 June 2004
27645	86-4538	10 October 2005



First Names and Last Names Data Sets

Production Database

First Name	Last Name	GPA	High School
Paul	Smith	3.2	Princeton
Johnson	NJ		

First Name
Lookup
Table

NY

Last Name
Lookup
Table

John
Bob
Danielle
Dave
Stacey

Newton
Nelson
Kline
Howell
Reese

1) Client is a University who wishes to mask the first and last name fields in their admissions database

2) Optim now has a first name lookup table with over 5,000 male/female names and a last name lookup table with over 80,000 names

Test Database

3) Use Lookup Tables to randomly replace table first and last names

First Name	Last Name	GPA	High School	Advisor
Stacey	Nelson	3.2	Princeton	
Johnson	NJ			

Dave

Reese
NY

2.7

Albany

Kline

Street Address/City/State/Zip Code Data Sets

Total Assets	Customers	Street	City	State	Zip Code
\$534,674,233	54,999	12 Buttercup Ln	Cleveland	OH	44101
\$8,777,733,811	105,333	6767 Rte 1 S	Princeton	NJ	08540

Address
Lookup
Table

1) Client is a Bank who wishes to mask its assets by location

288 Elm St	Milwaukee	WI	53201
12 Rodeo Dr	Los Angeles	CA	90001
3526 Diamond Rd	Seattle	WA	98101
12 Street Road	Las Vegas	NV	89101
2 Applegarth Ln	Brunswick	ME	04011

2) Optim provides corresponding Street Address/City/State/Zip Codes for masking

New Table with Masked Data

Total Assets	Customers	Street	City	State	Zip Code
\$534,674,233	54,999	3526 Diamond Rd	Seattle	WA	98101
\$8,777,733,811	105,333	21 Street Rd	Las Vegas	NV	89101

3) Leverage Multiple Column Replacement. Entire address row can be masked with a valid CASS address using enhanced random lookup function

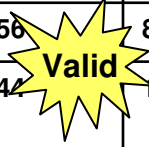

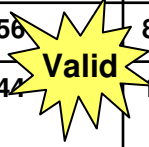

Intelligent Masking Capability

Production Database

F. Name	L. Name	Credit Card#	SSN#
John	Denver	5298774132478855	254-77-6644
Vanessa	Jones	4324115574123654	154-74-7788

Data before Masking

Test Database

F. Name	L. Name	Credit Card#	SSN#
John	Denver	5326458711224956 	854-77-6644 
Vanessa	Jones	4972584612457744 	154-74-7788 

Data after Masking... Masked with Valid CC# and SS#

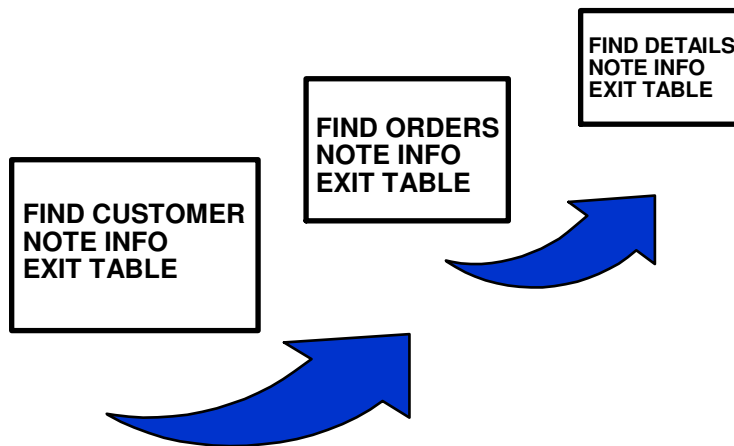
How are these numbers valid?

For Social Security Numbers	For Credit Card Numbers
A Social Security Number (SSN) consists of nine digits. The first three digits is called the "area number". The central, two-digit field is called the "group Number". The final four-digit field is called the "serial Number". All numbers must fit the latest available criteria for each section.	Most credit card numbers are encoded with a "Check Digit". A check digit is a digit added to a number (either at the end or the beginning) that validates the authenticity of the number. A simple algorithm is applied to the other digits of the number which yields the check digit.

Traditional vs. Relational Tools


Single Table Editors

- One table/view at a time
- No edit of related data from multiple tables

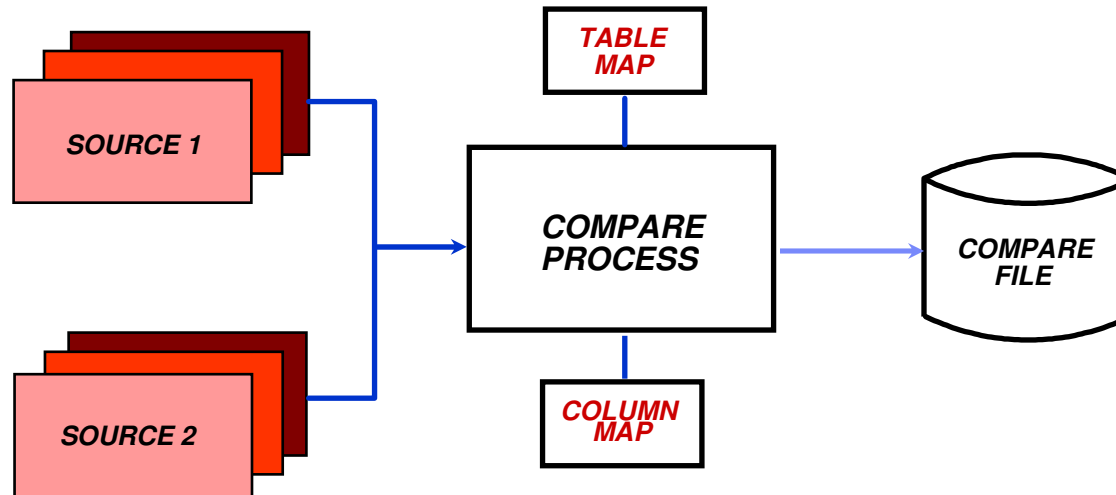


The Relational Editor

- **Simultaneous browse/edit of related data from multiple tables**



OPTIM TDM Compare Facility



- Single-table or multi-table compare
- Creates compare file of results
- Displays results on screen
- For application testing, QA, and to verify database contents
- Enhances productivity by finding unexpected changes in the data


What are the Key Drivers of Data Growth?

- Mergers & acquisitions
- Organic business growth
 - eCommerce
 - ERP/CRM
- The digital revolution
- Records retention
 - Basel II
 - SOX
 - Euro-SOX
- Data multiplier effect
- Forrester estimates that 85% of data stored in databases is inactive



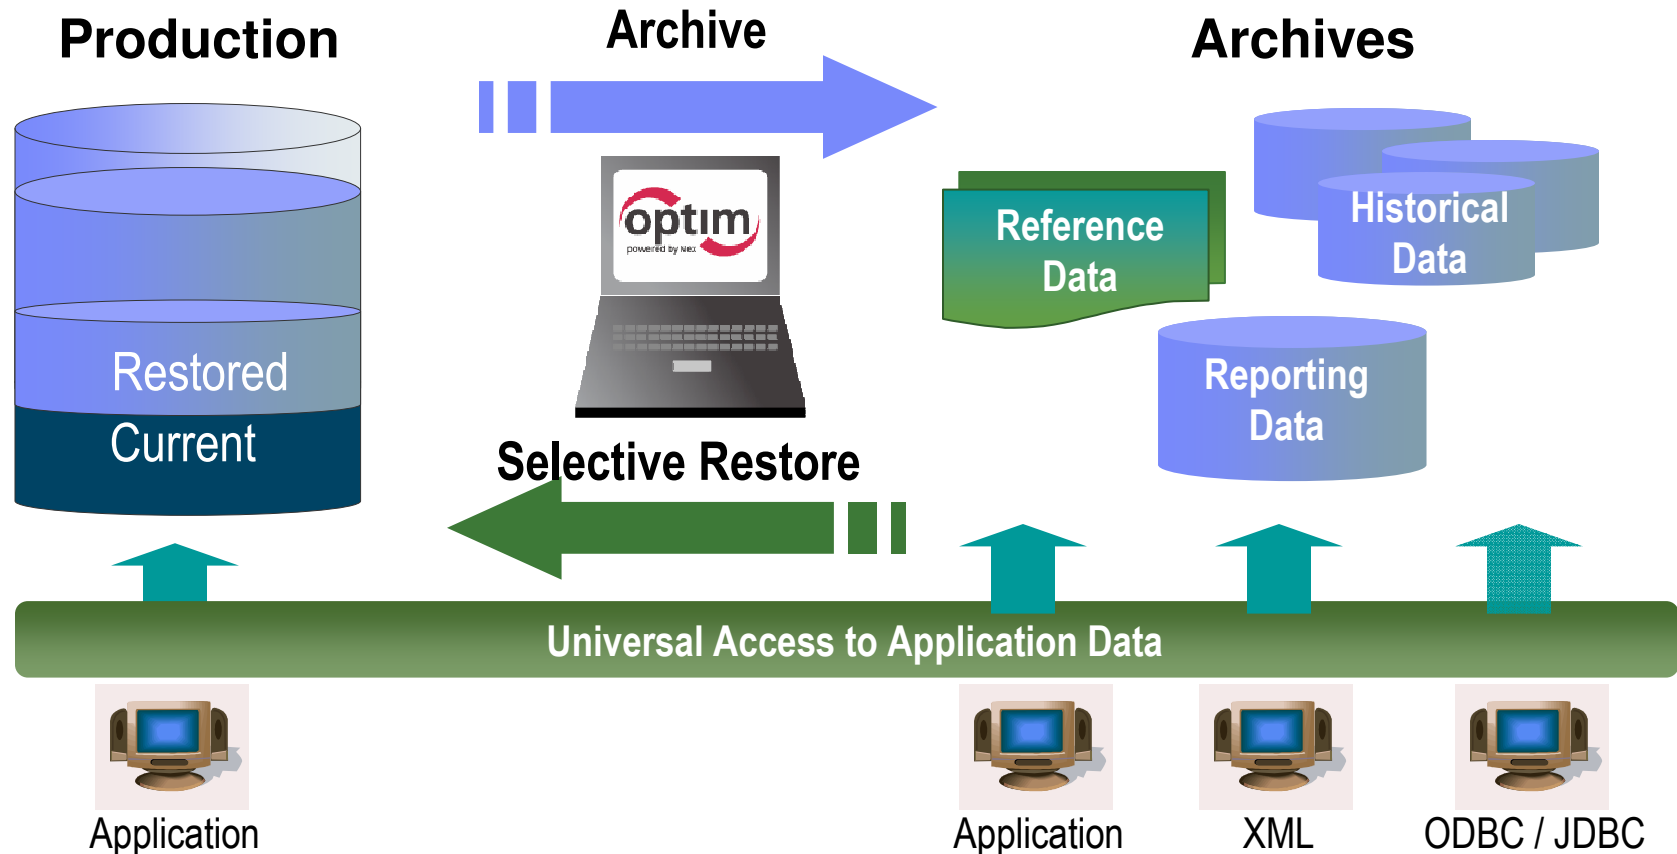
* Source: Noel Yuhanna, Forrester Research, Database Archiving Remains An Important Part Of Enterprise DBMS Strategy, 8/13/07

The data multiplier effect



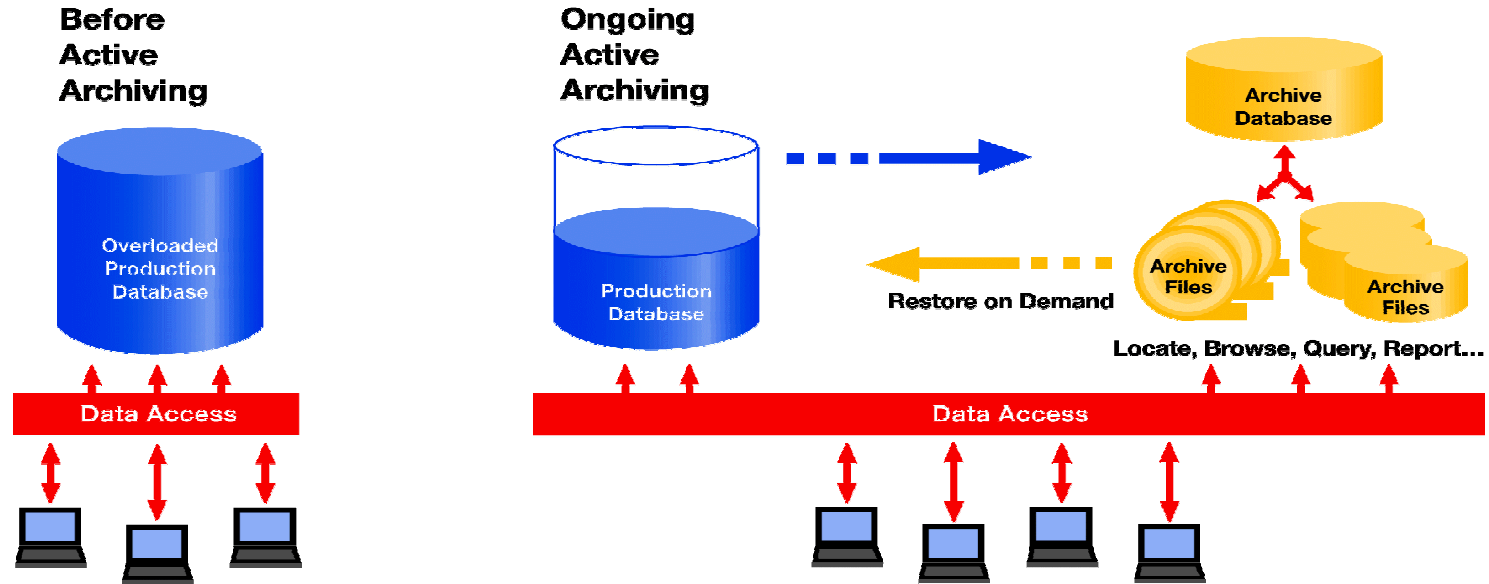
Actual Data Burden = Size of production database + all replicated clones

Optim™ Data Growth Solution: Archiving



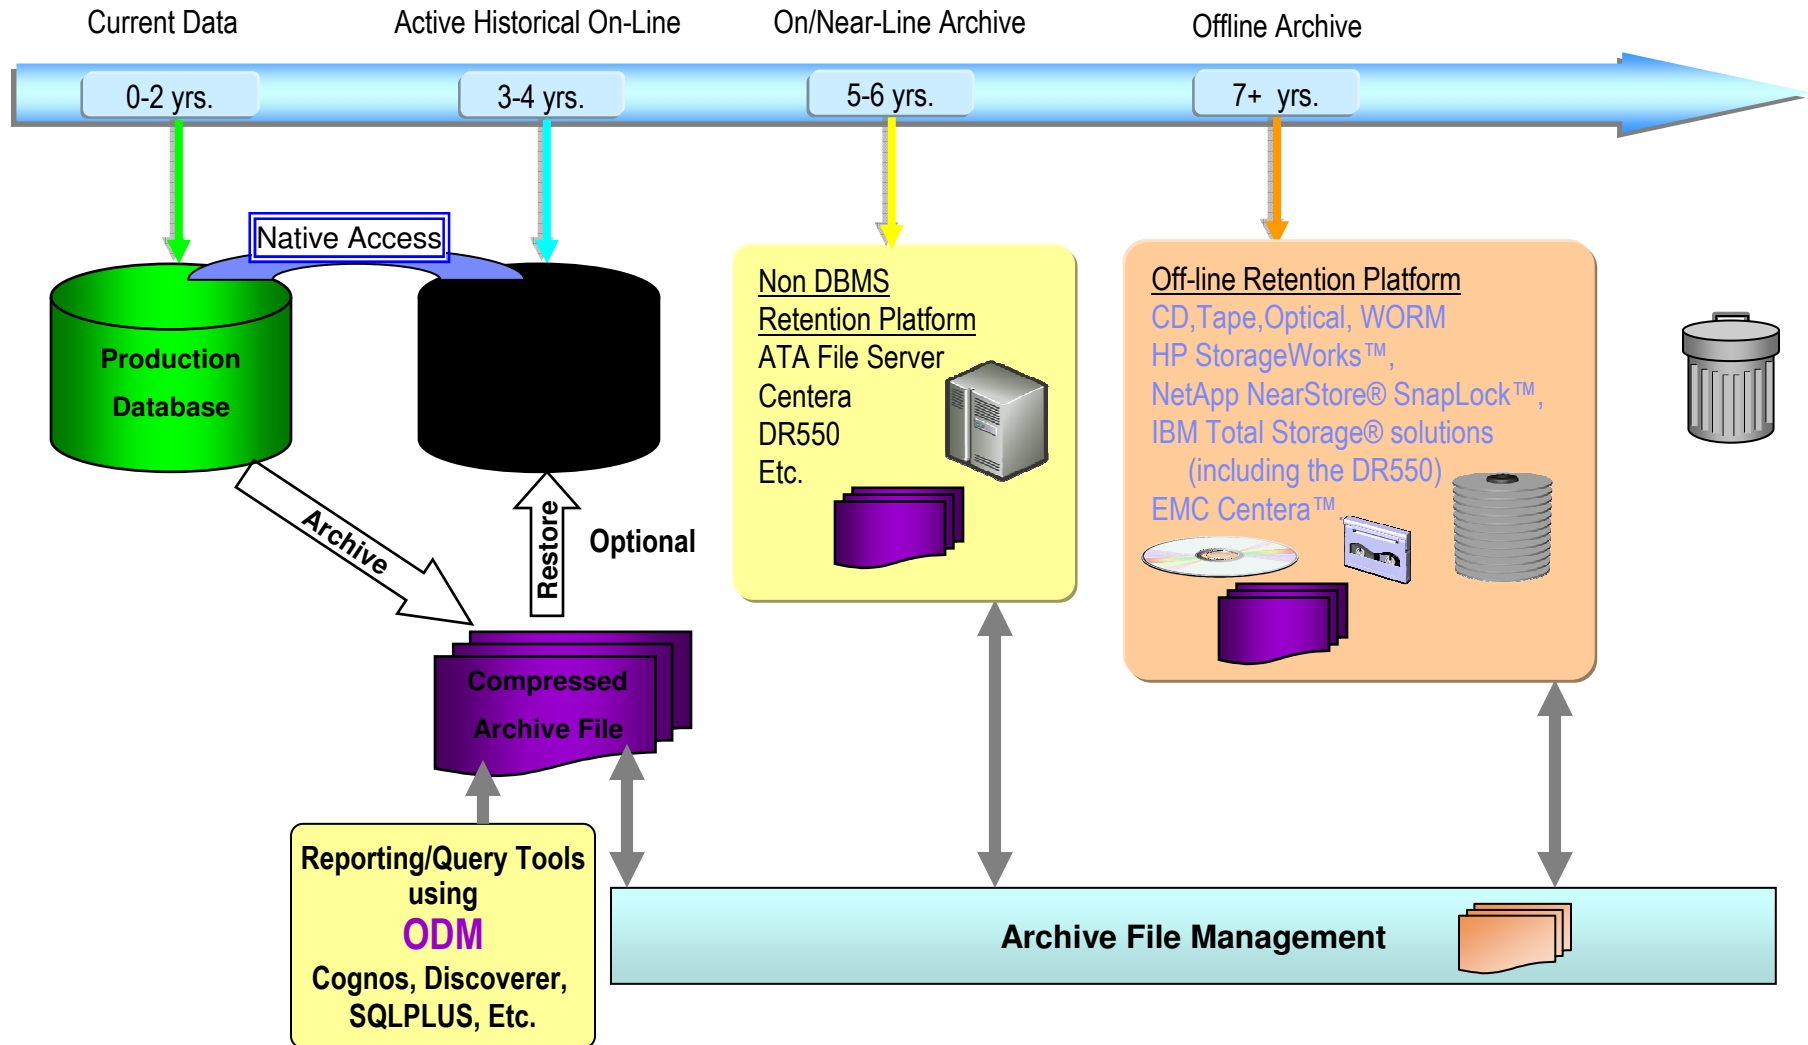
- Complete Business Object provides historical reference snapshot of business activity
- Storage device independence enables ILM
- Immutable file format enables data retention compliance

Active Archiving Defined




- Reduce the amount of data in the application database by:
 - Separating infrequently accessed data from transactional data
 - Preserve metadata and relationships of archived data outside db
 - Archive relational subsets vs. entire files
- Enable easy user access to archived information
 - View, research and restore as needed
- Complementary to Information Lifecycle Management (ILM)

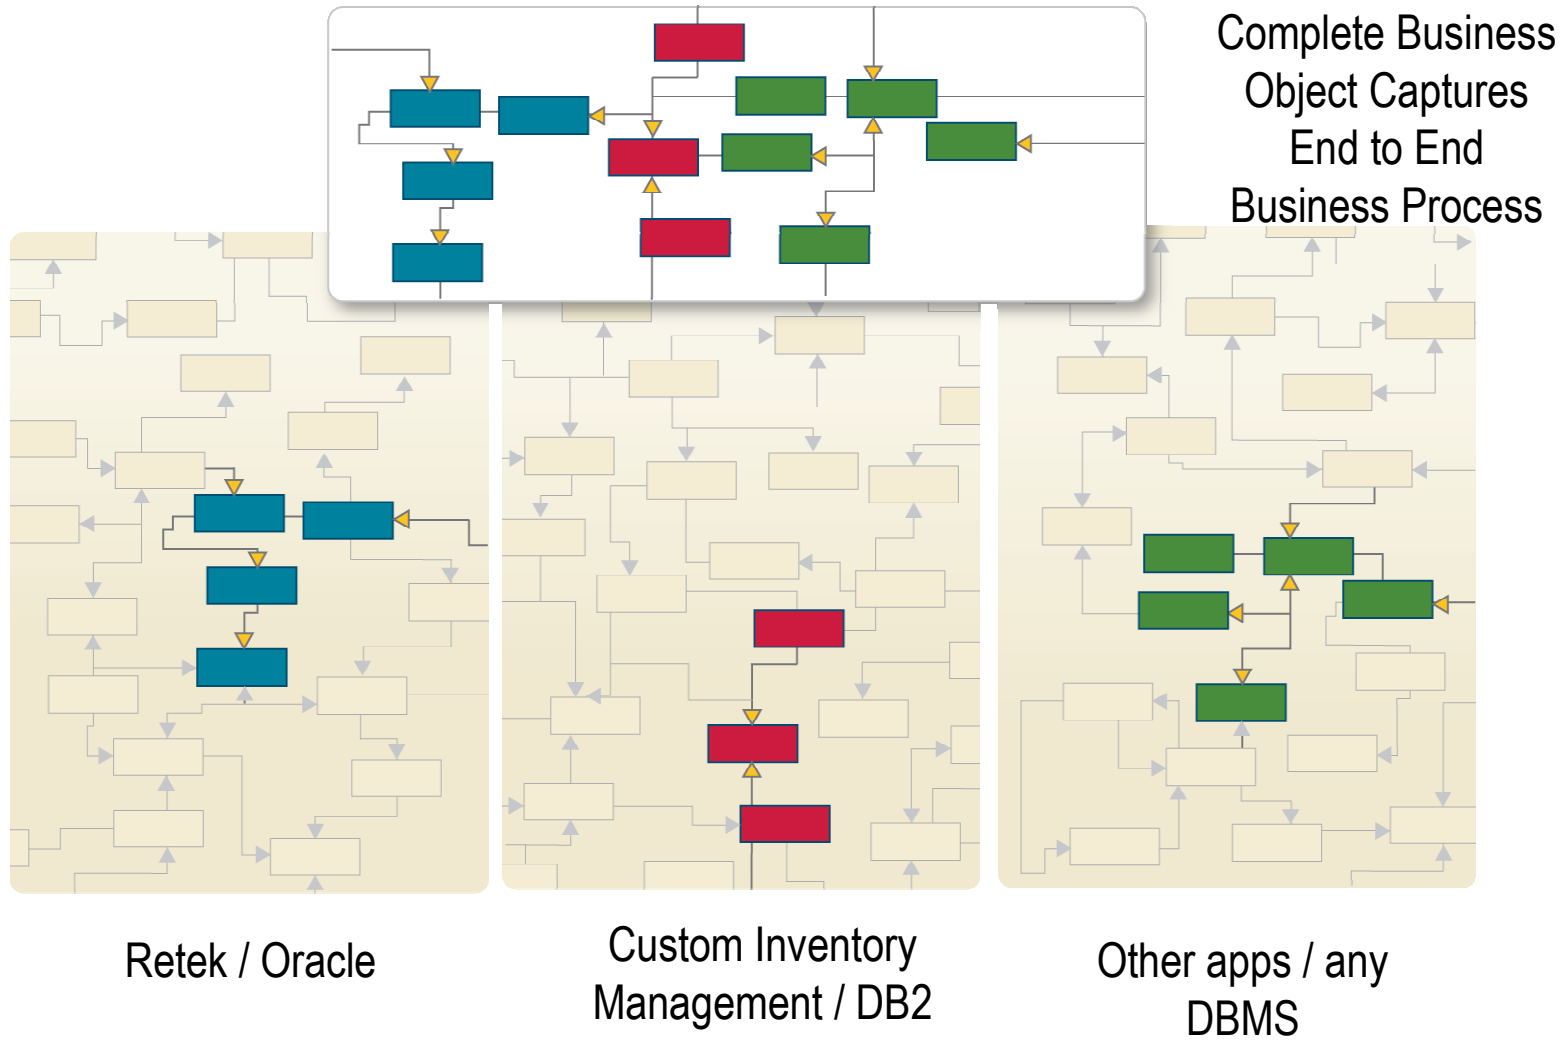
Information Lifecycle



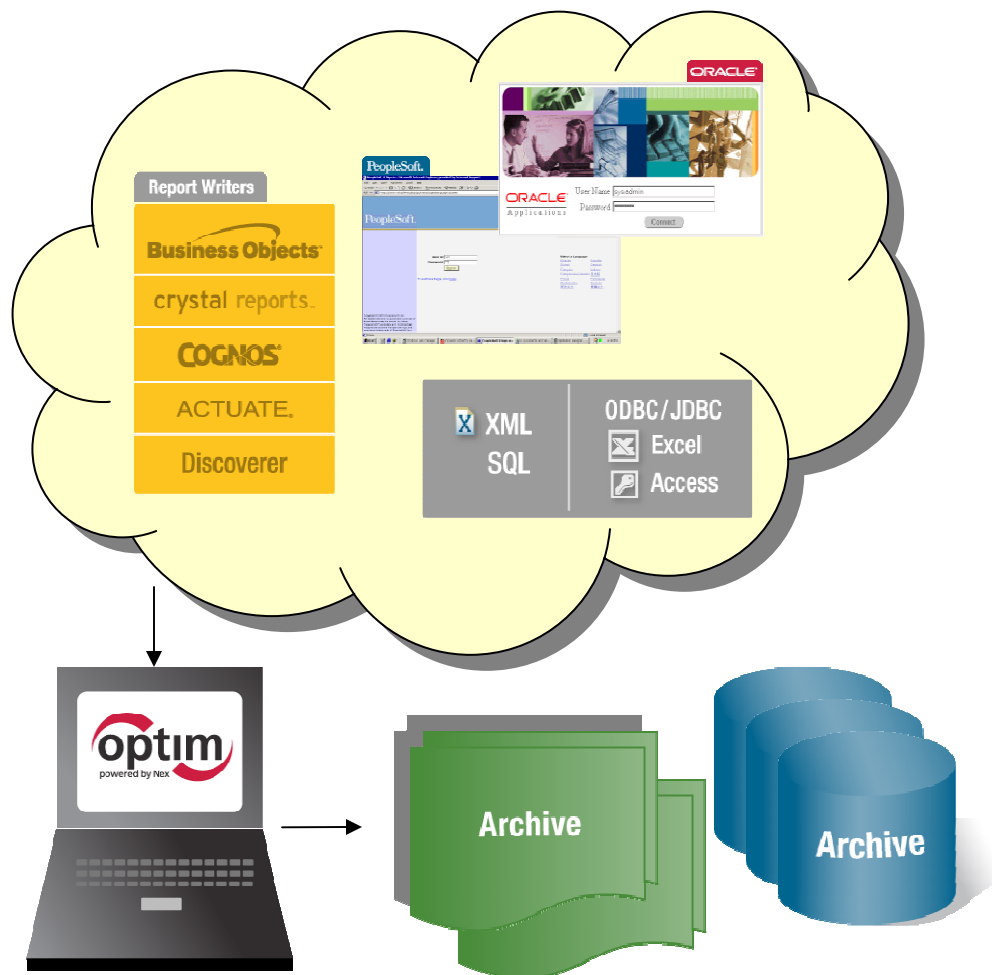
Our Unique Capability: Complete Business Object



Extract - Federated Data Support



Universal Access



- Native application access
 - Familiar screens and processes
- Application independent access
 - Industry standard methods: SQL, ODBC/JDBC, XML
 - Portals
 - Report writers: Crystal Reports, Cognos, Business Objects, Discoverer, Actuate
 - Desktop formats: Excel, CSV, MS Access
 - Database formats

Access Any Record, Anytime, Anywhere!

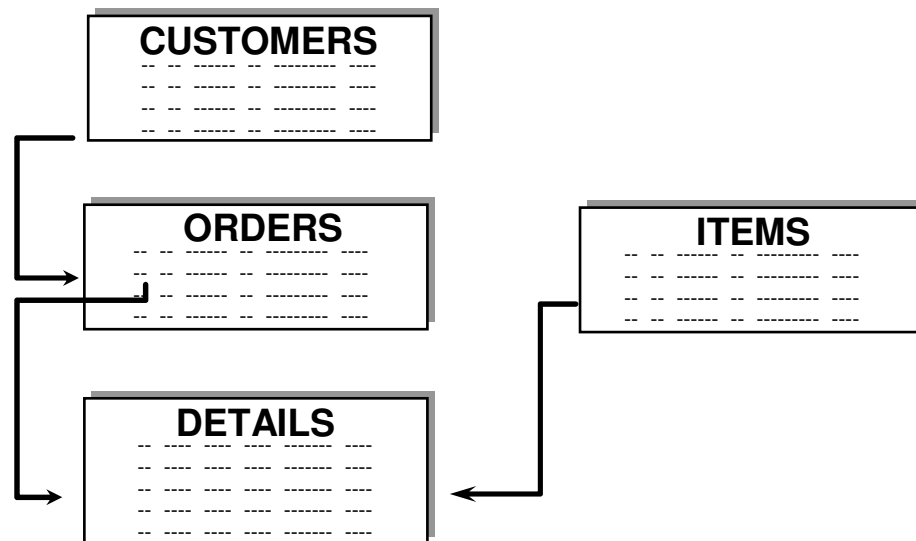
Steps for Archiving Data

- **Identify the data to be archived**
- **Define the data to be deleted**
- **Choose a delete method**
- **Create the archive & Delete the data**
- **Find Data in the Archives**
- **Browse or Restore**

Identify the data to be archived

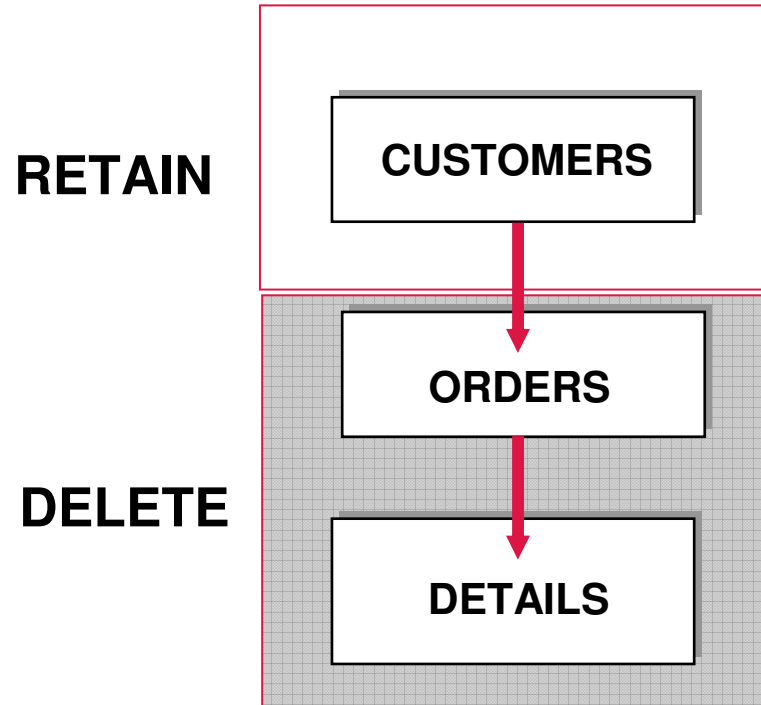
Access Definition

Defines a subset of of relational data



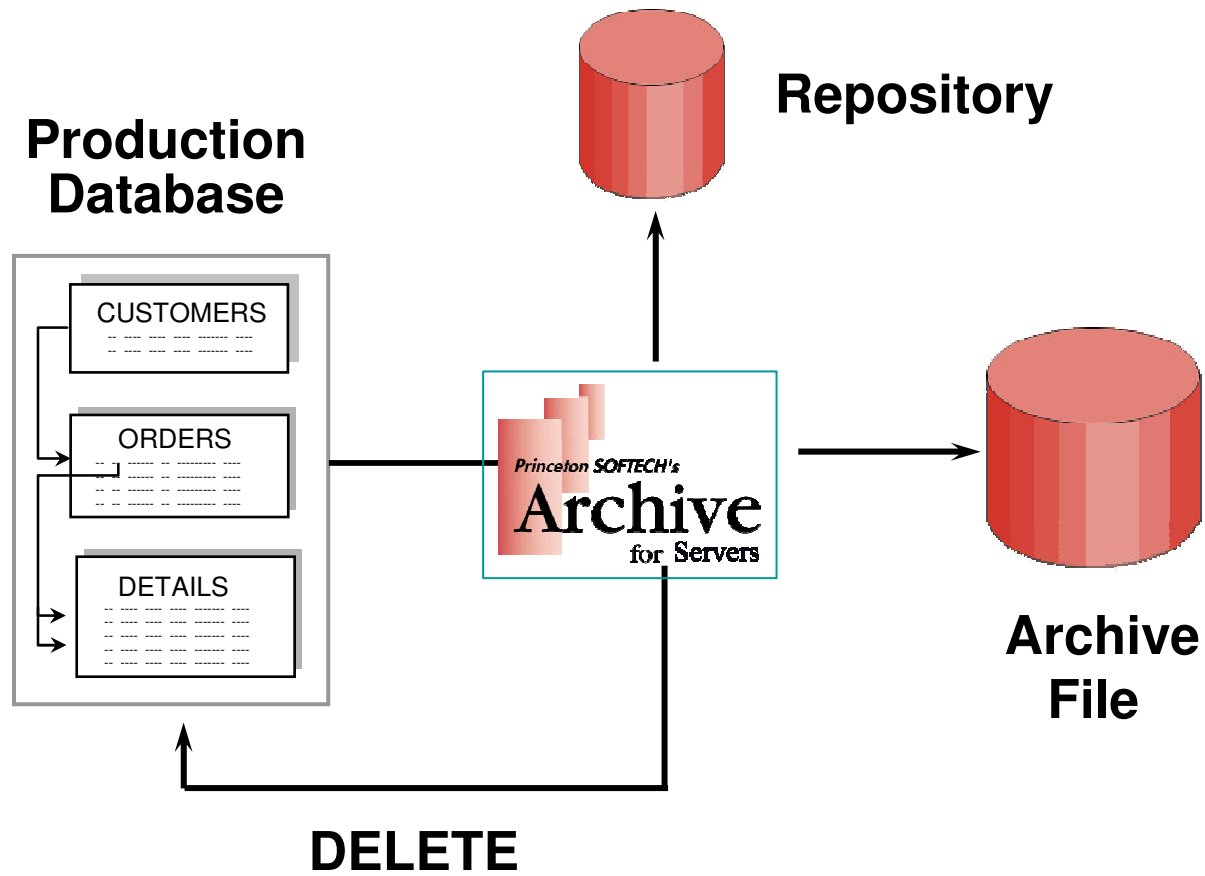
- Start table
- Associated data
- Relationships
- Extraction rules
- Index specifications

Define the data to be deleted



- Archive all data
- Delete orders and details after they are safely archived
- Preserve semantic intelligence

Create the archive



Researching the Archives



Direct access to archived data:

- User maintainable indexes
- Global searches
- Simple or complex criteria
- Intelligent browse
- ODBC Access
- ODM Access
- Save as CSV

Restore archived data only when you need to

Why Restore?



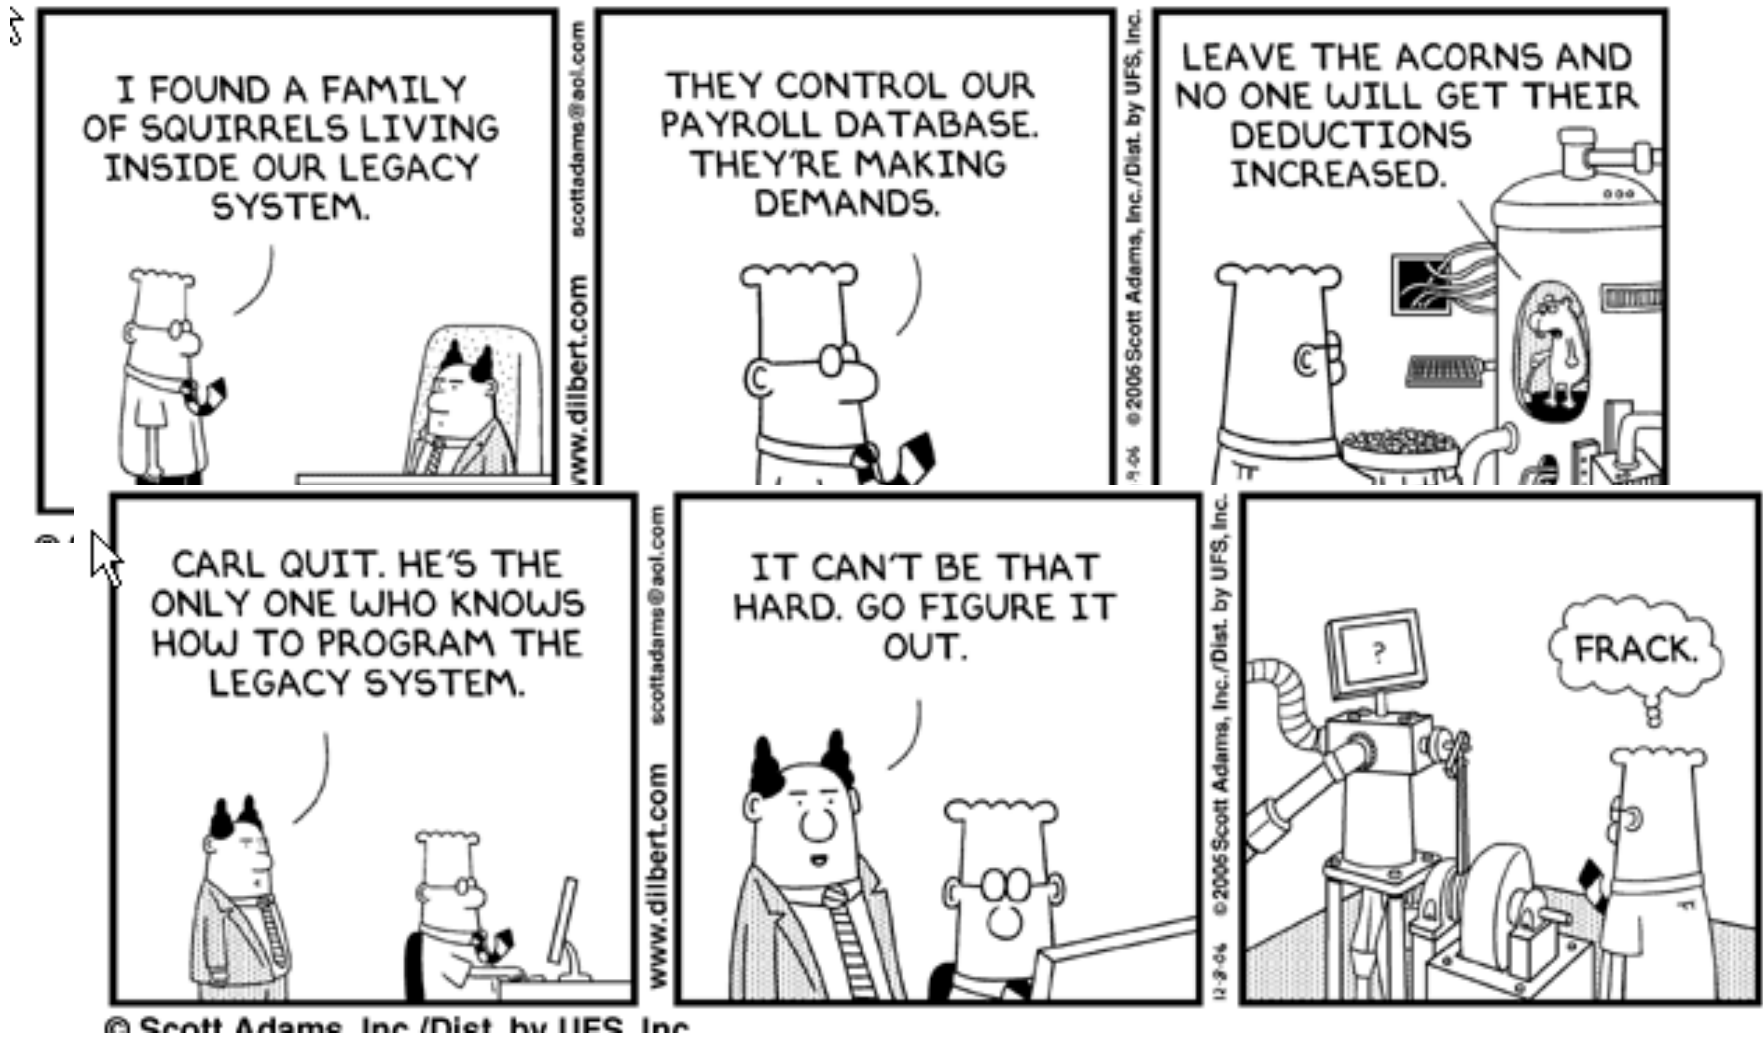
Browse archived data for:

- Customer service
- Answering questions
- Archive research

Restore archived data for:

- Audit situations
- Application-generated reports

Untold stories of legacy applications



When you retire or consolidate applications don't move all of the data

- Application portfolio has redundant systems acquired via mergers and acquisitions
- Line of business divested; application is no longer needed
- Legacy technologies not compatible with current IT direction or supported by vendor
- Required technical skills or application knowledge no longer available
- Budget pressures – do more with less

In almost ALL cases, access to legacy data MUST be retained while the application and database are eliminated



Issue: Retire obsolete applications

CIO:

- Reduce risk and cost by sunsetting obsolete or redundant technologies
- Reduce IT expenses (software, hardware, personnel)
- Preserve access to legacy system data for retention compliance

Business Line Executive:


- Minimize negative budget variances and reduce IT charge-backs to line of business
- Allocate scarce resources for priority business needs
- Maintain access to historical business data for retention compliance

Technical Management:

- Eliminate expenses associated with underperforming assets
- Enable application-independent access to legacy system data for retention compliance
- Reduce risks from dependence on specialized labor and no longer supported vendor products

Enterprise Challenge: Application Retirement

Optim Supports Application Retirement Strategies



Summary

- **Take Back Control with IBM Information Protection solutions on System z:**
 - Transform your information from a Liability into your most strategic, valuable Asset
 - Help manage business risk by enforcing security, audit, privacy and policy controls
 - Lower operational costs by optimising data management, retention and archiving

- **Software, Hardware and Expertise.**
 - Information Management - the most complete end-to-end Information Protection software solutions
 - Information Protection Entry point as part of your wider Information Governance strategy
 - System z - the ultimate platform to for security
 - Clear ROI business cases for each area of Information Protection .

- For more information visit
 - www.ibm.com/software/data/db2imstools/solutions/data-governance.html
 - [Download the Information Protection white paper now.](#)

THANK
YOU