IBM System z Technical University

October 4–8, 2010 — Boston, MA

z/VSE Applications accessing DB2 on Linux on System z

Session ID: zDL02

Wilhelm Mild, IBM

Authorized IBM Training

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

Data-consolidation – more important than ever

Decisions for a future oriented Data store

Experiences from last projects / Redbook

A good solution is not standard in detail

Bad Data Can be Costly

83% of data integration projects either overrun or fail

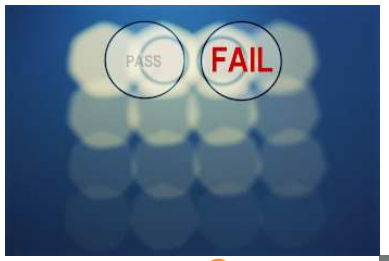

Inaccurate or incomplete data is a leading cause of failure in business-intelligence and CRM projects

Scrap and rework
Increased costs

Lack of consumer confidence

Lost opportunities

Low data quality costs companies $611 billion annually

25% of time is spent clarifying bad data

Undetected defects will cost 10 to 100 times as much to fix upstream

The big Data store

Data Integration – the Base for the future and BI

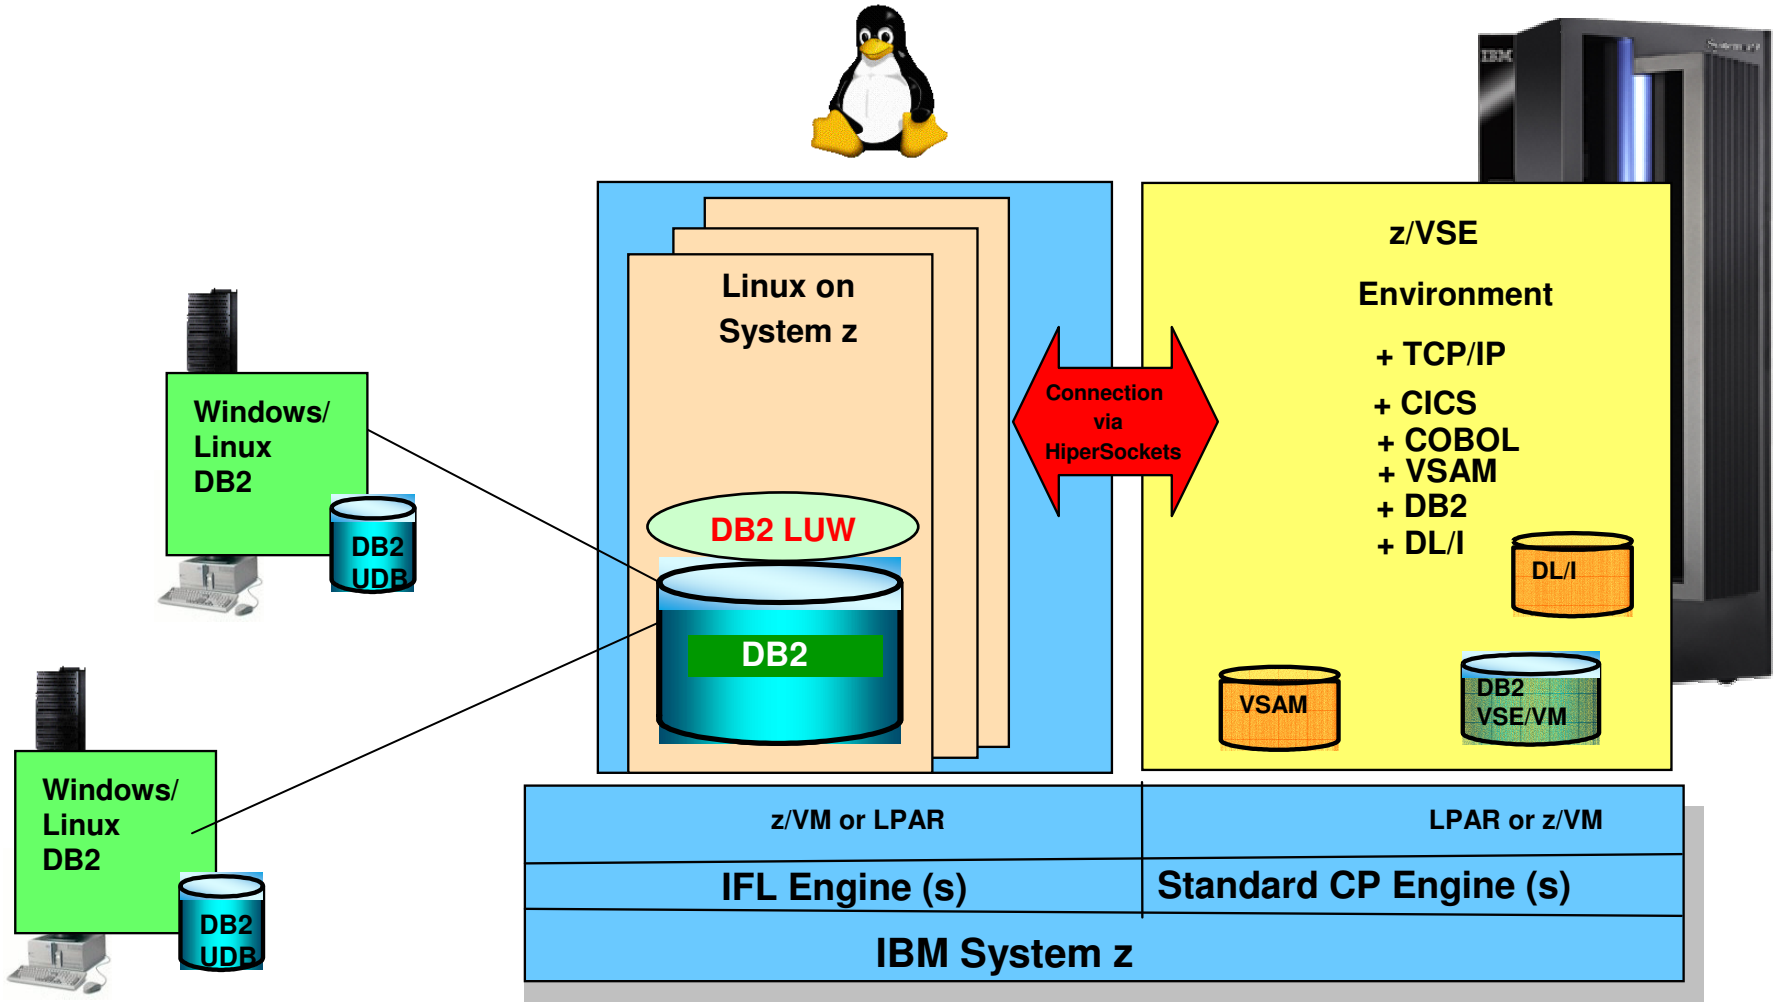

The road to information availability is filled with challenges

- What are the top business challenges?
  - Streamline and improve efficiency of business processes
  - Better understand and meet customer expectations
  - Increase employee productivity

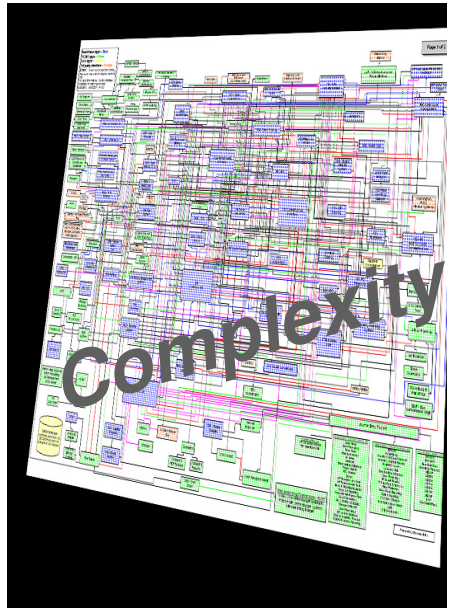

- Key challenges to making information available:
  - **Volume**: Data & content are doubling each year
  - **Variety**: It's not just the transaction data, it's e-mails, document libraries, etc.
  - **Velocity**: The pace of business and business users who need information *now*, in real time
  - **Complexity**: The average $1B company has 40 financial systems; 78% of all companies have 2 or more repositories, 25% have more than 15 repositories.

Information as a Service

From a project based approach to a Service Oriented Architecture based on business needs

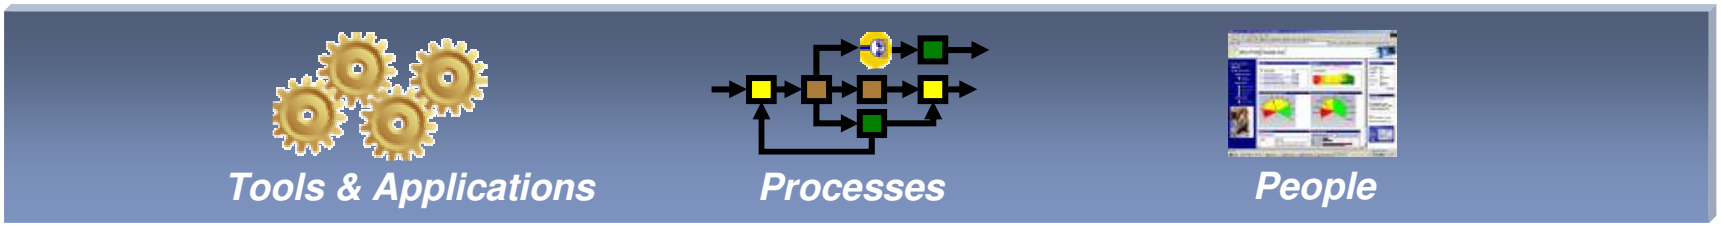

Standards-based

- *SQL*
- *XQuery*
- *JCR*
- *JDBC*
- *Web Services...*

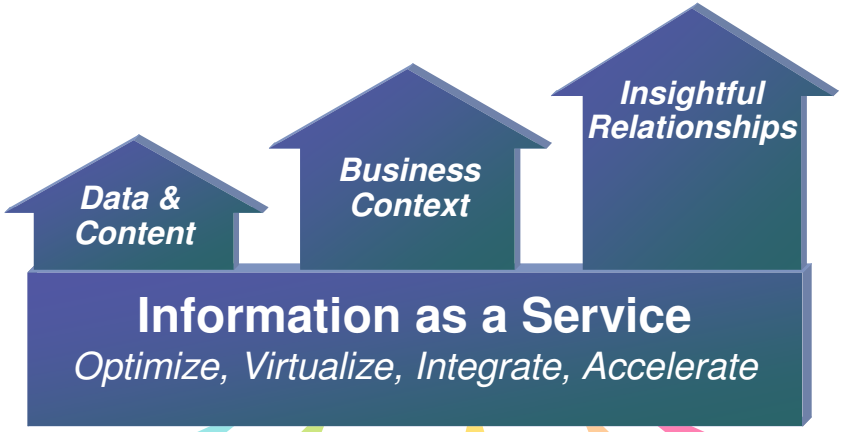

Extracted or Real-time Insight

- *Master Data*
- *Entity Analytics*
- *Information Warehouses*
- *Industry Data Models*

Heterogeneous Applications & Information

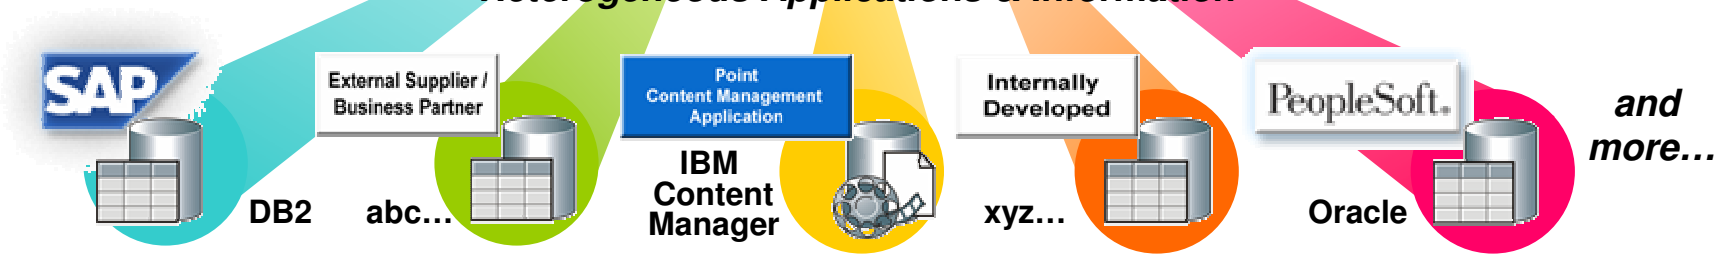

A New Generation Data Server for A New Generation of Applications

Data Server
Services that manage, secure and provide access to the database.

Hybrid Data Server
DB2 supports both relational and pureXML® storage and provides all the necessary services to support both data structures.

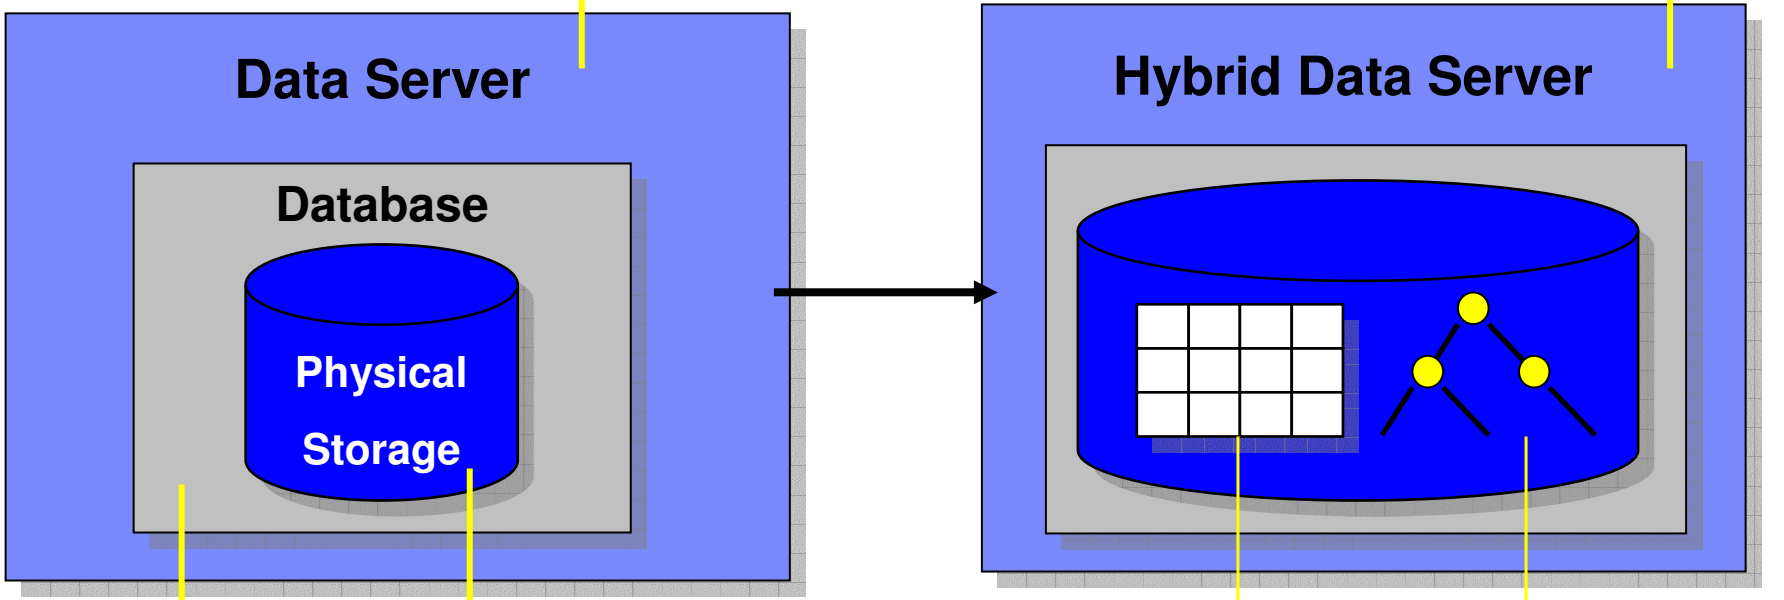

Database
Logical View of storage
- *Tables*
- *Views*

Physical storage
Database Files

Relational Storage
Data stored in a row and column format

pureXML® Storage
Data stored in a pre-parsed hierarchical format, not as a single text object (CLOB)

DB2 9 XML integration is seamless

Offers the Best to Both SQL and XML Worlds

SQL Person "I see a world class RDBMS that also supports XML"

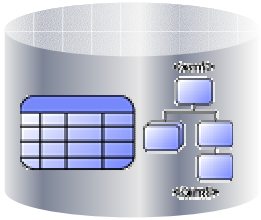

DB2 with XML Support

XML Person… "I see a world class XML repository that also supports SQL"

New XML applications benefit from:

- Ability to seamlessly leverage relational investment
- Proven Infrastructure that provides enterprise-class capabilities

Powering a Flexible Approach

XML and SOA are the Keys

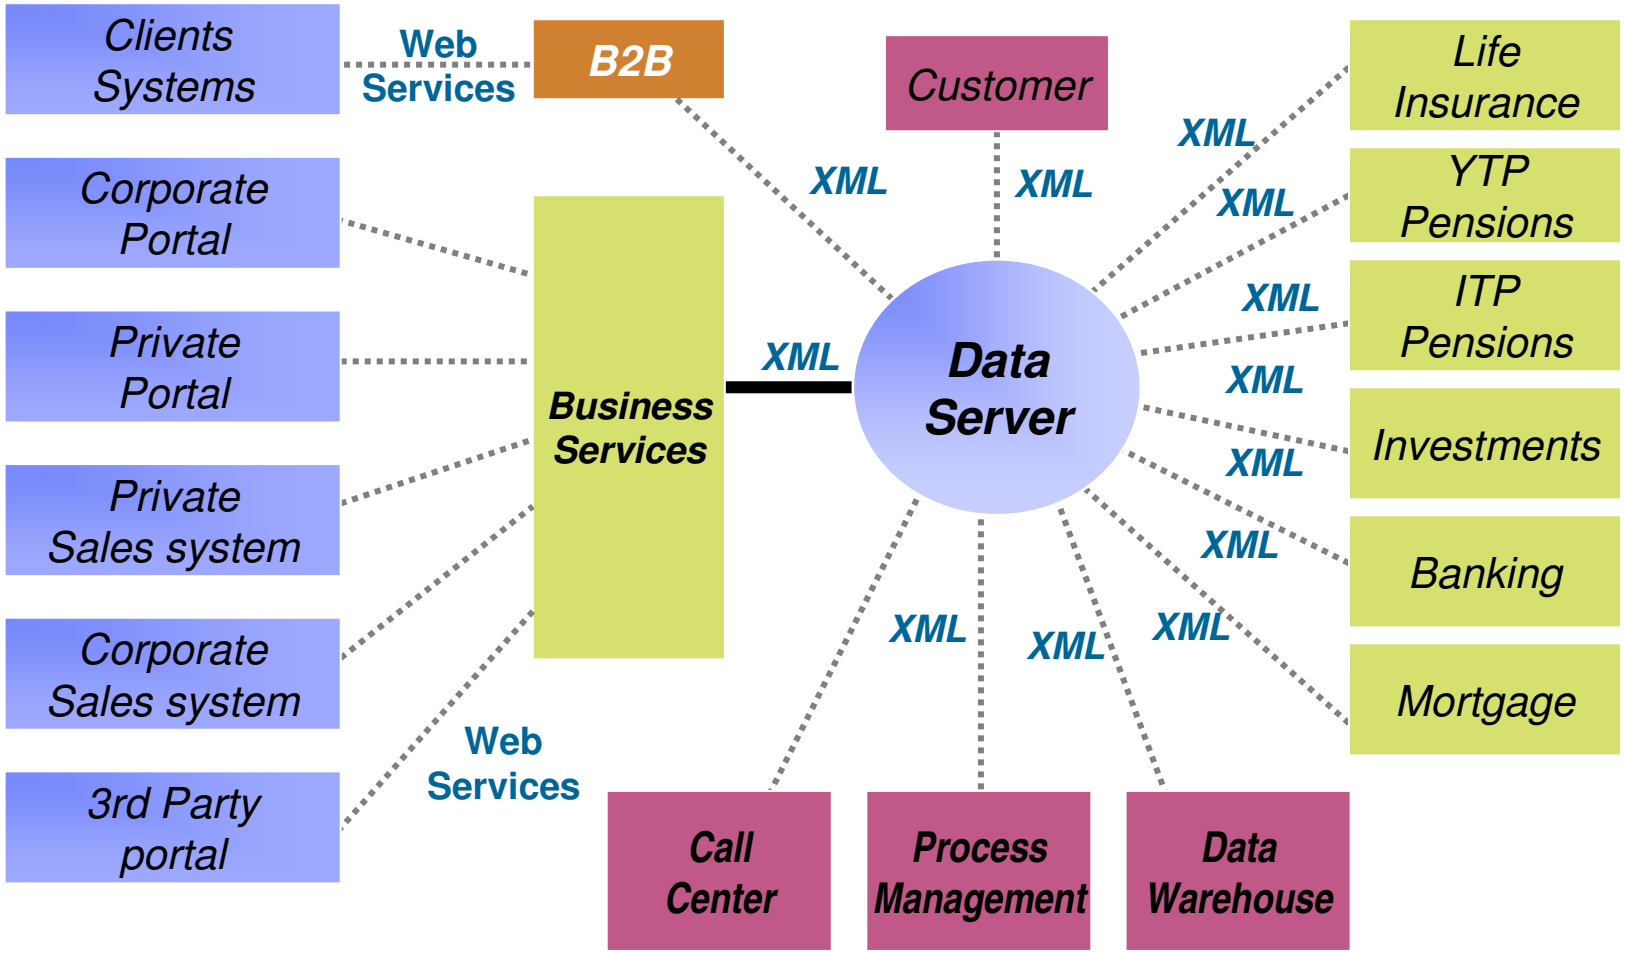

Top Scenario: Linux on System z as data hub

Consolidate, Integrate, Evaluate, Decide, Base for Business Intelligence (BI)

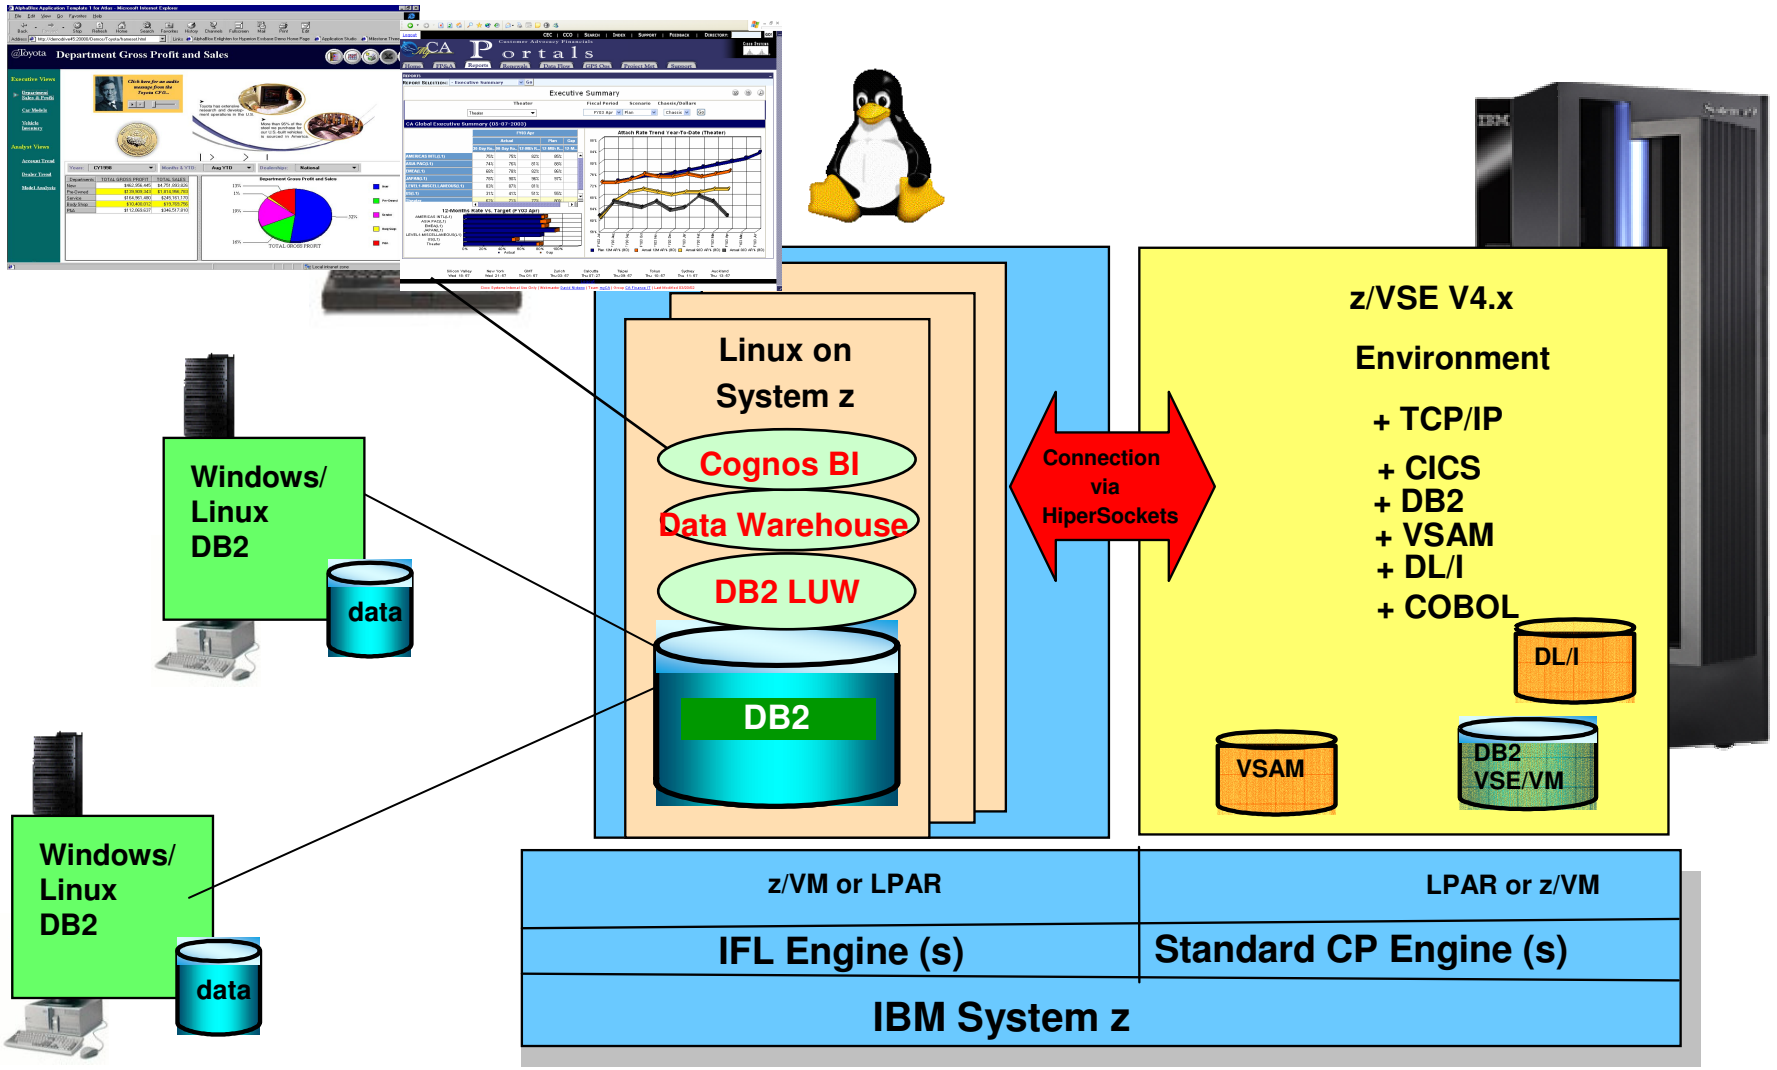

IBM InfoSphere Federation Server

- Integrating at the data layer – Federation of data
  - Read from and write to federated mainframe data sources using SQL
  - Standards-based access via JDBC, ODBC, or Call Level Interface
    - Including for VSAM
  - Multithreaded with native drivers for scalable performance
  - Metadata-driven means...
    - No mainframe programming required
    - Fast installation & configuration
    - Ease of maintenance
  - Works with existing and new...
    - Mainframe infrastructure
    - Application infrastructure
    - Toolsets

SQL

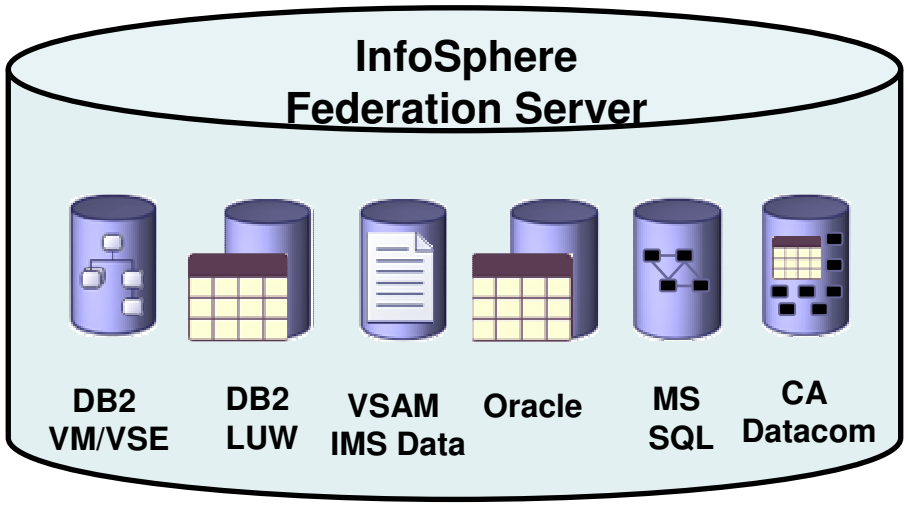

Agenda

Data-consolidation – more important than ever

Decisions for a future oriented Data store

Experiences from last projects / Redbook

A good solution is not standard in detail

From Planning to the Implementation and tuning

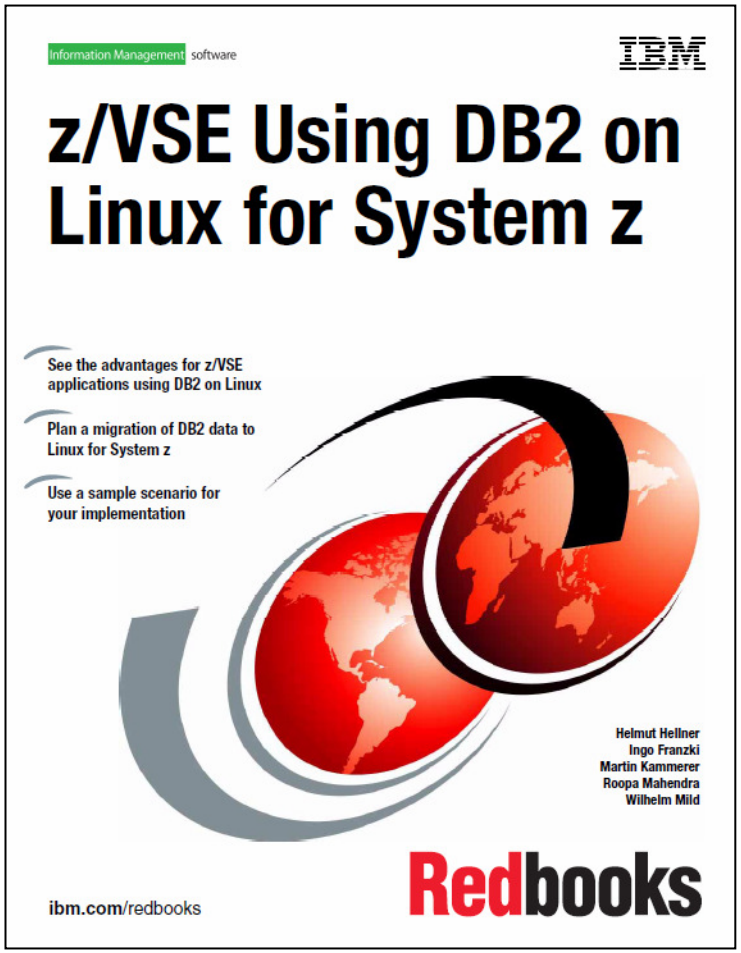

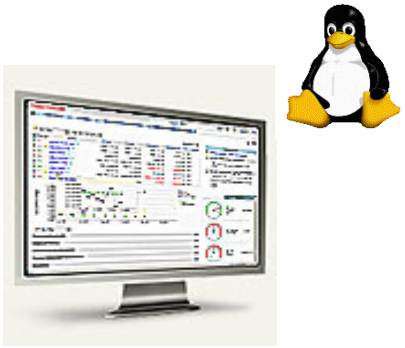

SG24-7690

DB2 Redbook

- Overview-
  - Strategical Decision
  - Advantages (Business Requiremnts)
  - Possible architectures
  - Technical overview(DB2 VM&VSE)
- Planning
  - Capacity Planning
  - Storage planning
  - Network
  - Database- DB2 Linux (LVM)- DB2 VM/VSE
  - The Transition phase
- Setup and Customization
  - DB2 Linux on System z
  - DB2 VSE (AR, AS)
  - DRDA Communication

- DBMS Migration
  - Data Migration
  - Packages Migration
  - Application considerations
  - Transition / Coexistence environment
- Monitoring and tuning
  - DB Monitoring
  - AR VSE
  - Appl. Monitoring (DB)
  - Connections / Interfaces
  - Network monitoring
  - System monitoring/tuning
  - Tuning considerations

DB2 Redbook

- **Overview**
  - Strategical Decisions
    - The decision for a modern Data Management System can enhance your business value substantially
  - Advantages (Business Requirements)
    - Business processes can be simplified a lot
  - Possible architectures
    - Data stores can be homogenous or heterogeneous,
  - Technical prerequisites
    - DB2 Server for VM&VSE (Server & Client)
    - DB2 Server for VM and VSE Client Editions

DB2 Szenarios – with DB2 LUW on Linux on System z

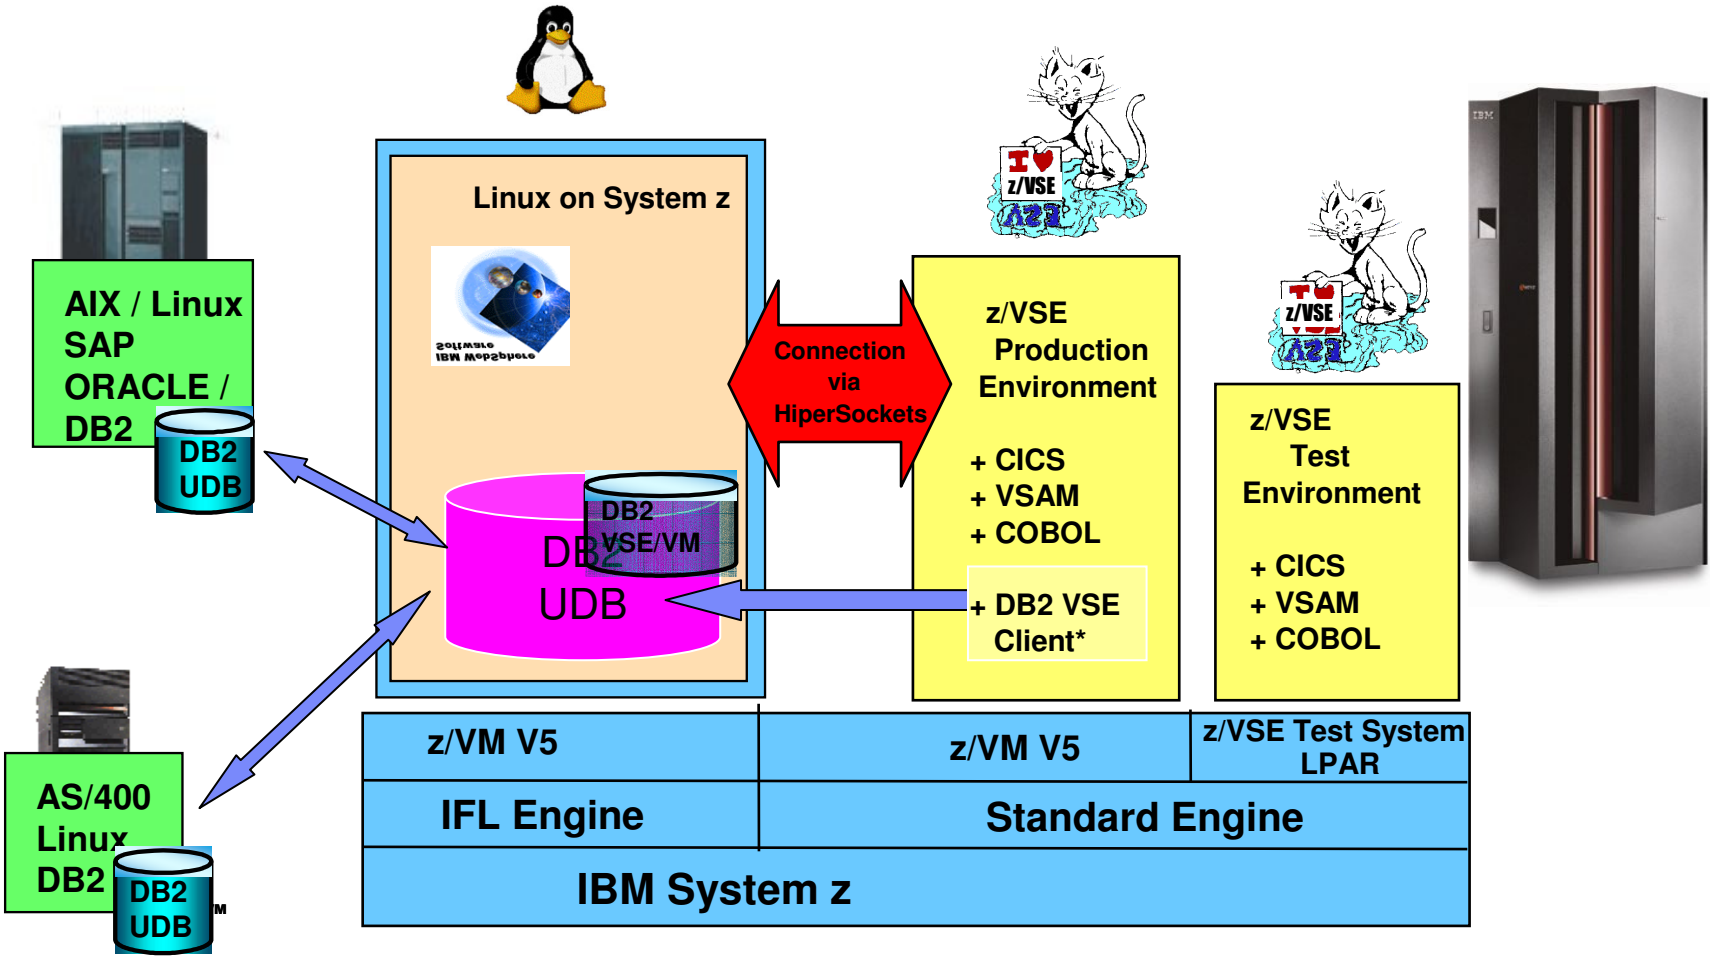

(*) DB2 VSE Client – the client functionality only, can be obtained with DB2 Server for VSE & VM 7.5 **Client Edition**

DB2 Redbook

- **Planning**
  - Capacity Planning
    - CPU load depends on many factors (parallel workload, IP traffic, application design)
    - z/VM virtualization increases flexibility and connectivity
  - Storage planning
    - The most advanced possibilities of the System z Architecture
      - use LVM (in Linux) or striped storage function (in DS8000)
      - use ECKD for system and FCP/ SCSI disks for large databases
    - High Availability
      - Mirroring / Redundant Connections
  - Database Planning on Linux
    - use LVM, Container Striping, PAV
  - Network
    - Hipersockets the fast communications
    - Shared OSA and VSWITCH the alternative Communication
  - Transition Phase
    - ‚Step by Step' always better instead of ‚Big Bang'!

Network alternatives

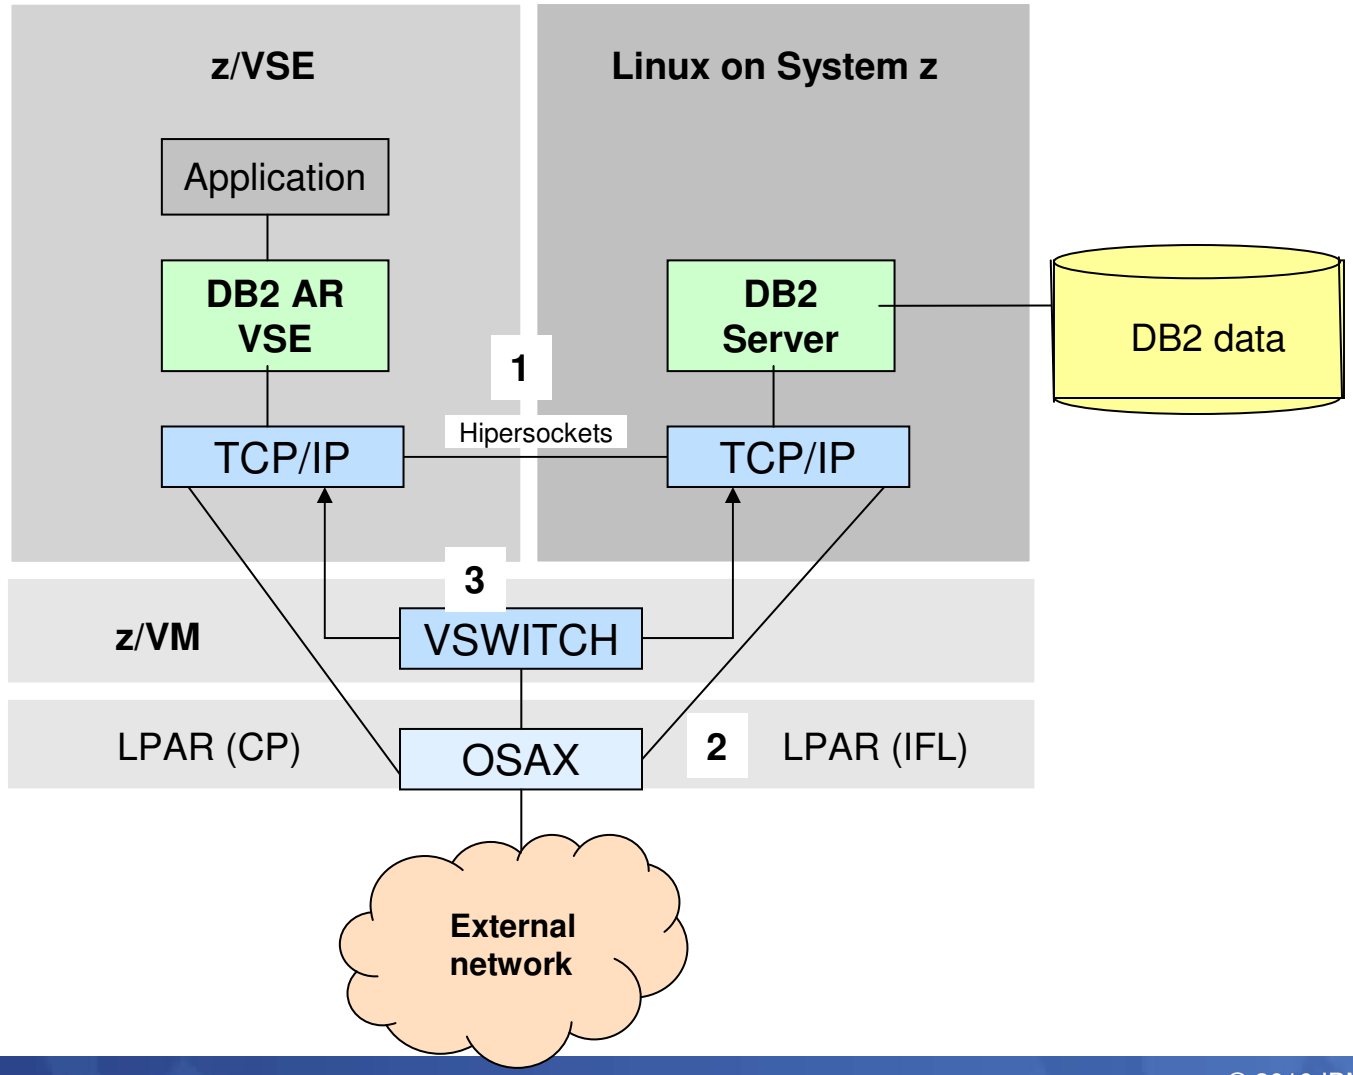

Network

- Network setup

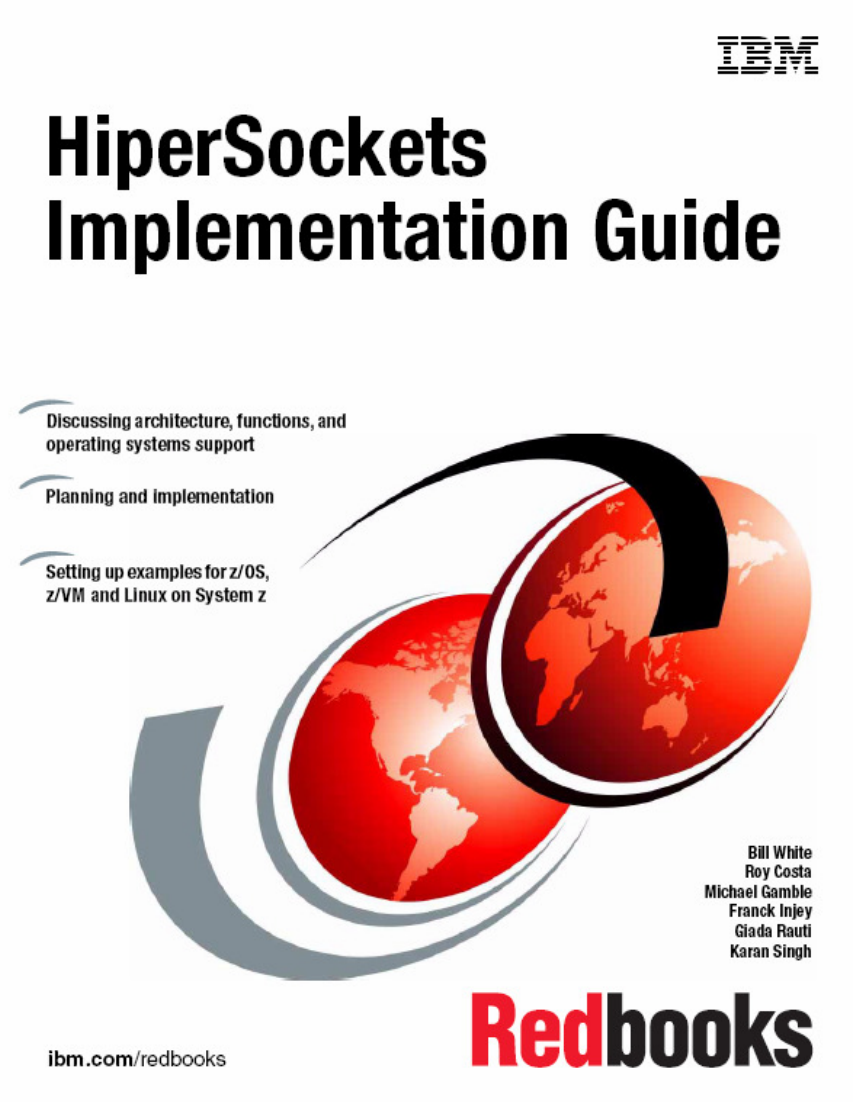

DB2 Redbook

- **Setup and Customization**
  - DB2 Linux on System z
    - Database-Definitions need to be adopted for the workload
      - Codepage (SBCS / Unicode)
      - EBCDIC versus ASCII Sort order ‘Collating Sequence’
      - Federation to implement complex requirements
  - DB2 VSE (Application Requestor)
    - Client Edition (AR only!) or Server & Client for VM/VSE
  - DRDA Communication
    - DRDA Performance is dependant on the application
    - Connection Pooling / Buffered Insert helps
    - TCP/IP Setup tuning for the workload (MTU, Window size)

Federated access for EBCDIC considerations

1) Linux applications can access the database as ASCII database
2) z/VSE applications access the database via vselnxdb as EBCDIC collated database

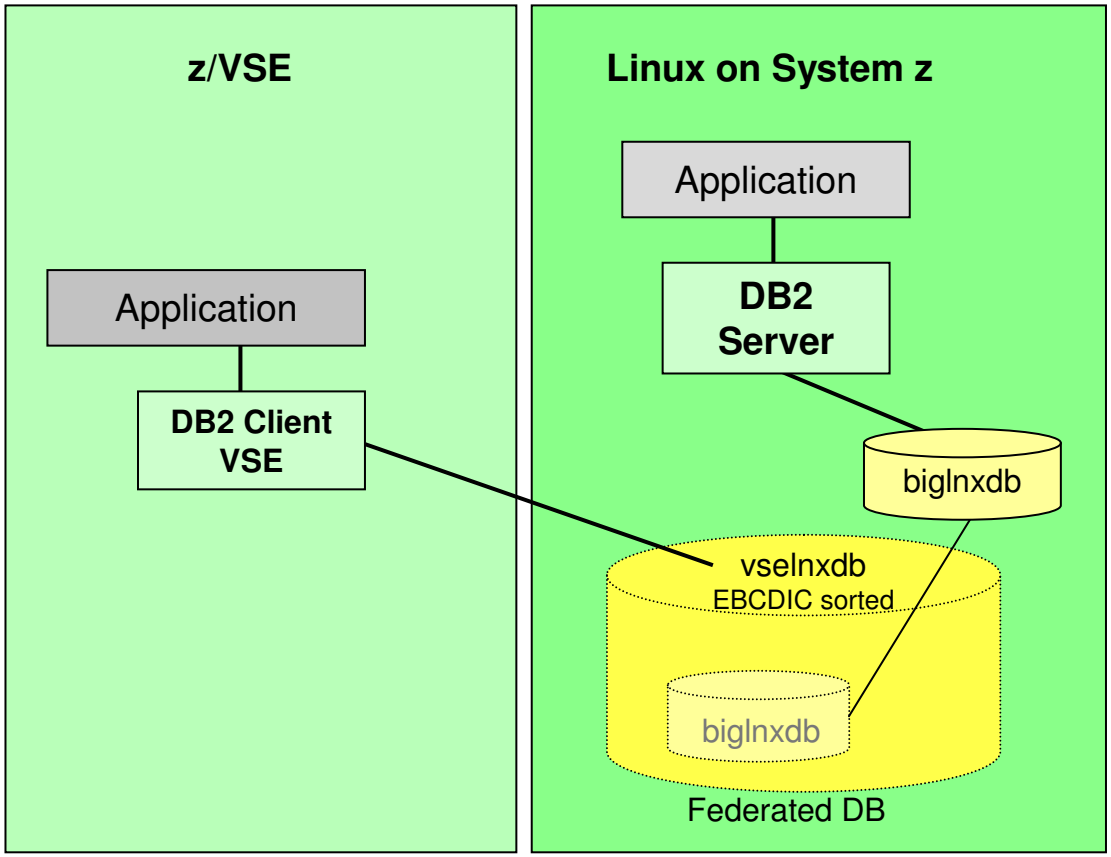

DB2 Redbook

- **DBMS Migration**
  - Data Migration
    - Data Migration: small effort / repeatable solution recommended
    - Federation is very effective
  - Package Migration
    - Bind Files build! (CICS or ‚Batch Binder‘)
    - Export of DB2/VM&VSE Packages and Import in DB2 Linux possible (not recommended)
  - Application Considerations
    - Applications may need adaptions (ASCII-EBCDID, HEX-Sort)
    - Dynamic SQL uses functionality of the server
  - Transition / Coexistence Environment
    - with Replication or ‚Federation‘, a coexistence is possible

Data migration to DB2 Linux with DB2 federation feature

1) Linux applications can access the databases using Federation feature
2) z/VSE applications access the database in z/VM or Z/VSE local

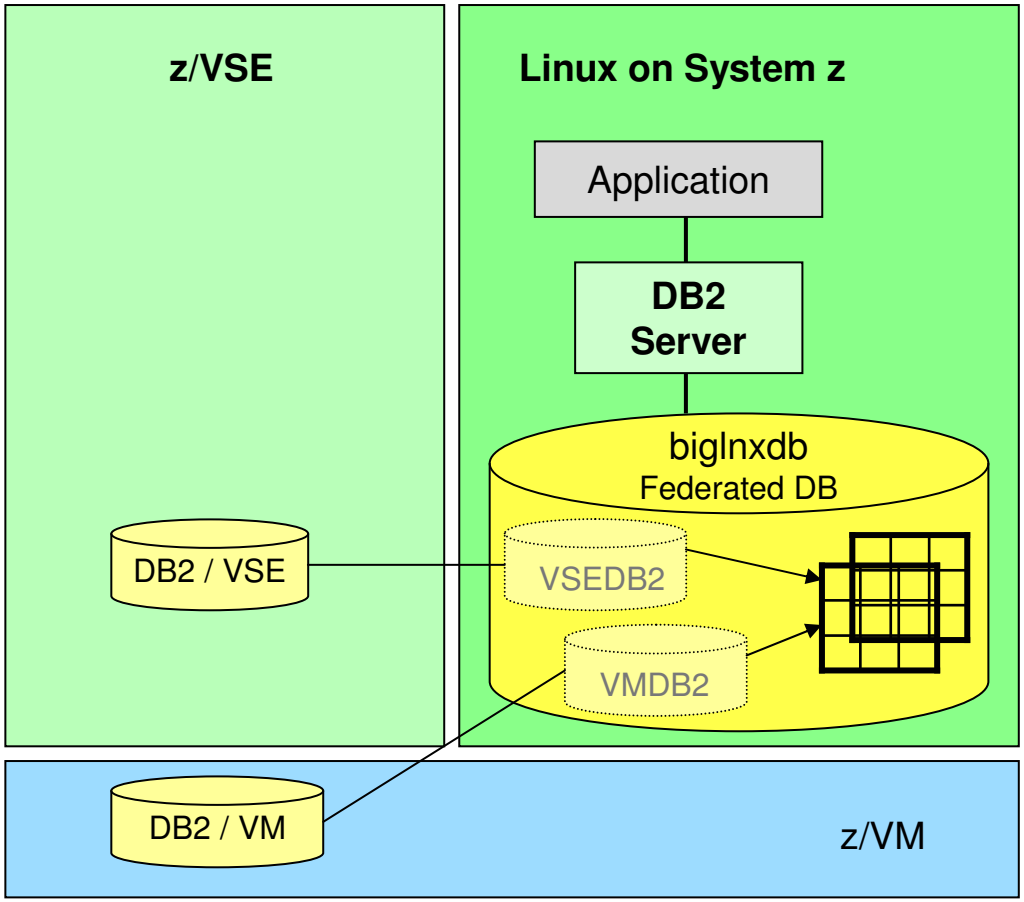

IBM InfoSphere Federation Server

- Integrating at the data layer – Federation of data
  - Read from and write to federated mainframe data sources using SQL
  - Standards-based access via JDBC, ODBC, or Call Level Interface
    - Including for VSAM
  - Multithreaded with native drivers for scalable performance
  - Metadata-driven means...
    - No mainframe programming required
    - Fast installation & configuration
    - Ease of maintenance
  - Works with existing and new...
    - Mainframe infrastructure
    - Application infrastructure
    - Toolsets

SQL

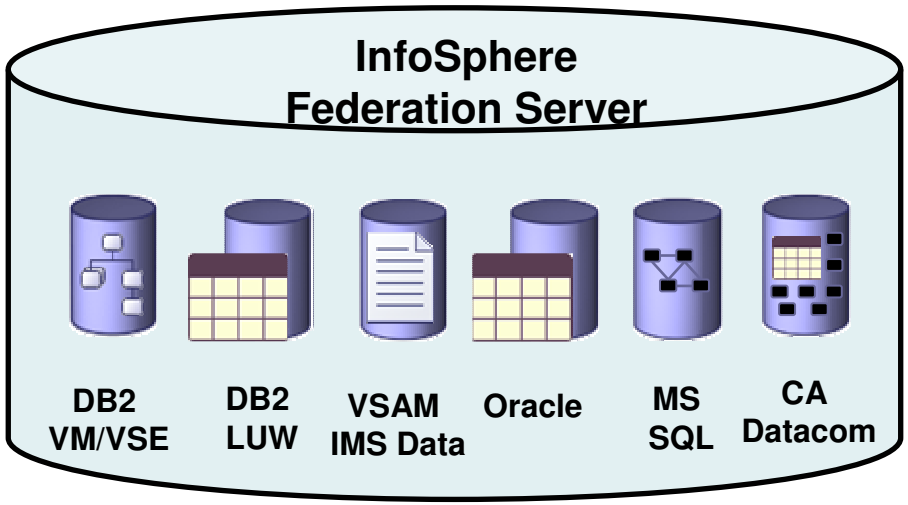

DB2 Coexistence pre-production scenario

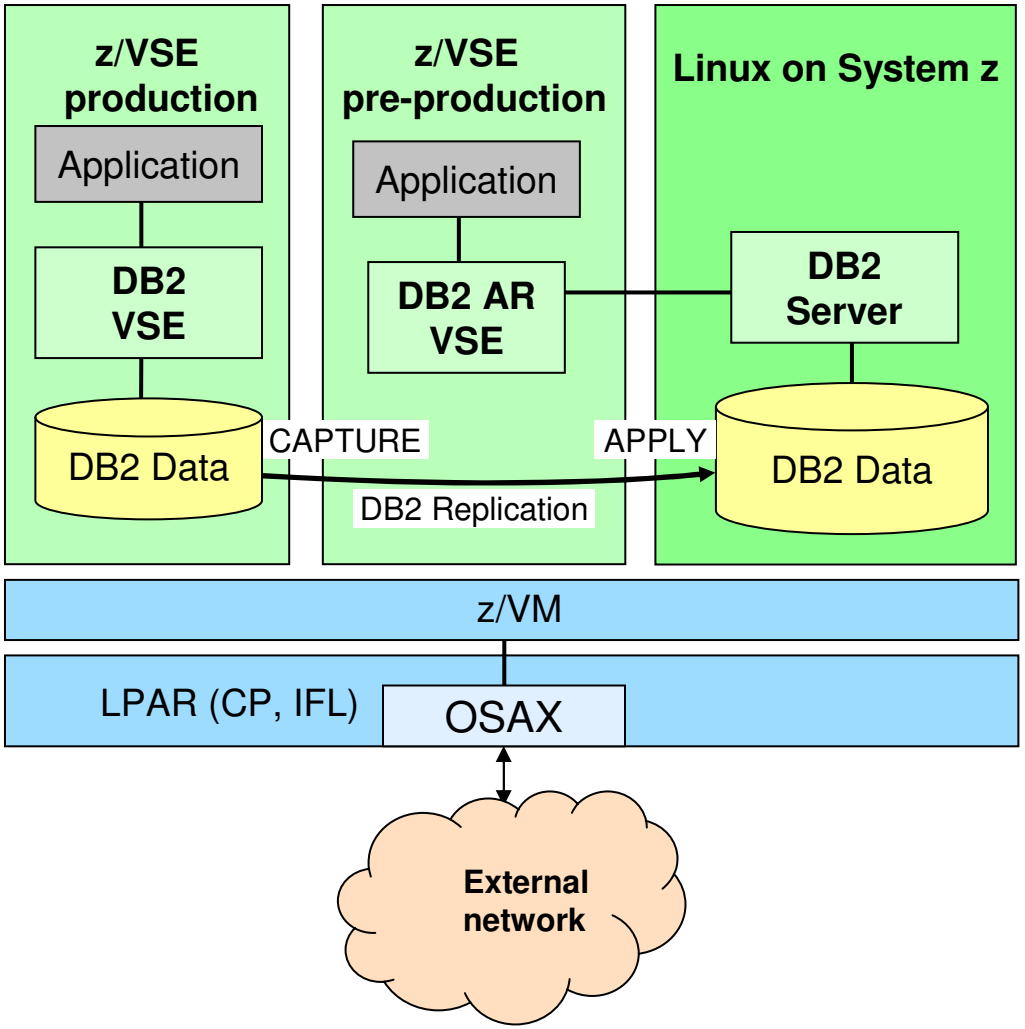

DB2 Redbook

- **Monitoring and Tuning**
  - Monitoring is prerequisite for Tuning
  - DB Monitoring
    - Status-quo of the DB2/VM or DB2/VSE Servers !!!
    - Monitor–Tools necessary
    - DB2/Linux – Snapshots, DB2 Expert, Omegamon XE
  - Application Monitoring (DB)
    - CICS Monitor is recommendable
  - Network Monitoring
    - Network monitors help a lot
    - Troubleshooting – analyze DB2 behavior with Network tools

Customer success samples with DB2 on Linux on System z

- US:
  - Supreme Court of Virginia
- Germany:
  - Wessels & Müller
  - Public Sector
- Slovenia:
  - Impol / Alcad
- Belgium:
  - Securex
- Sweden:
  - Pulsen
- Italy:
  - Olio Carli.

Customer implemetation(1): Public sector customer,Germany

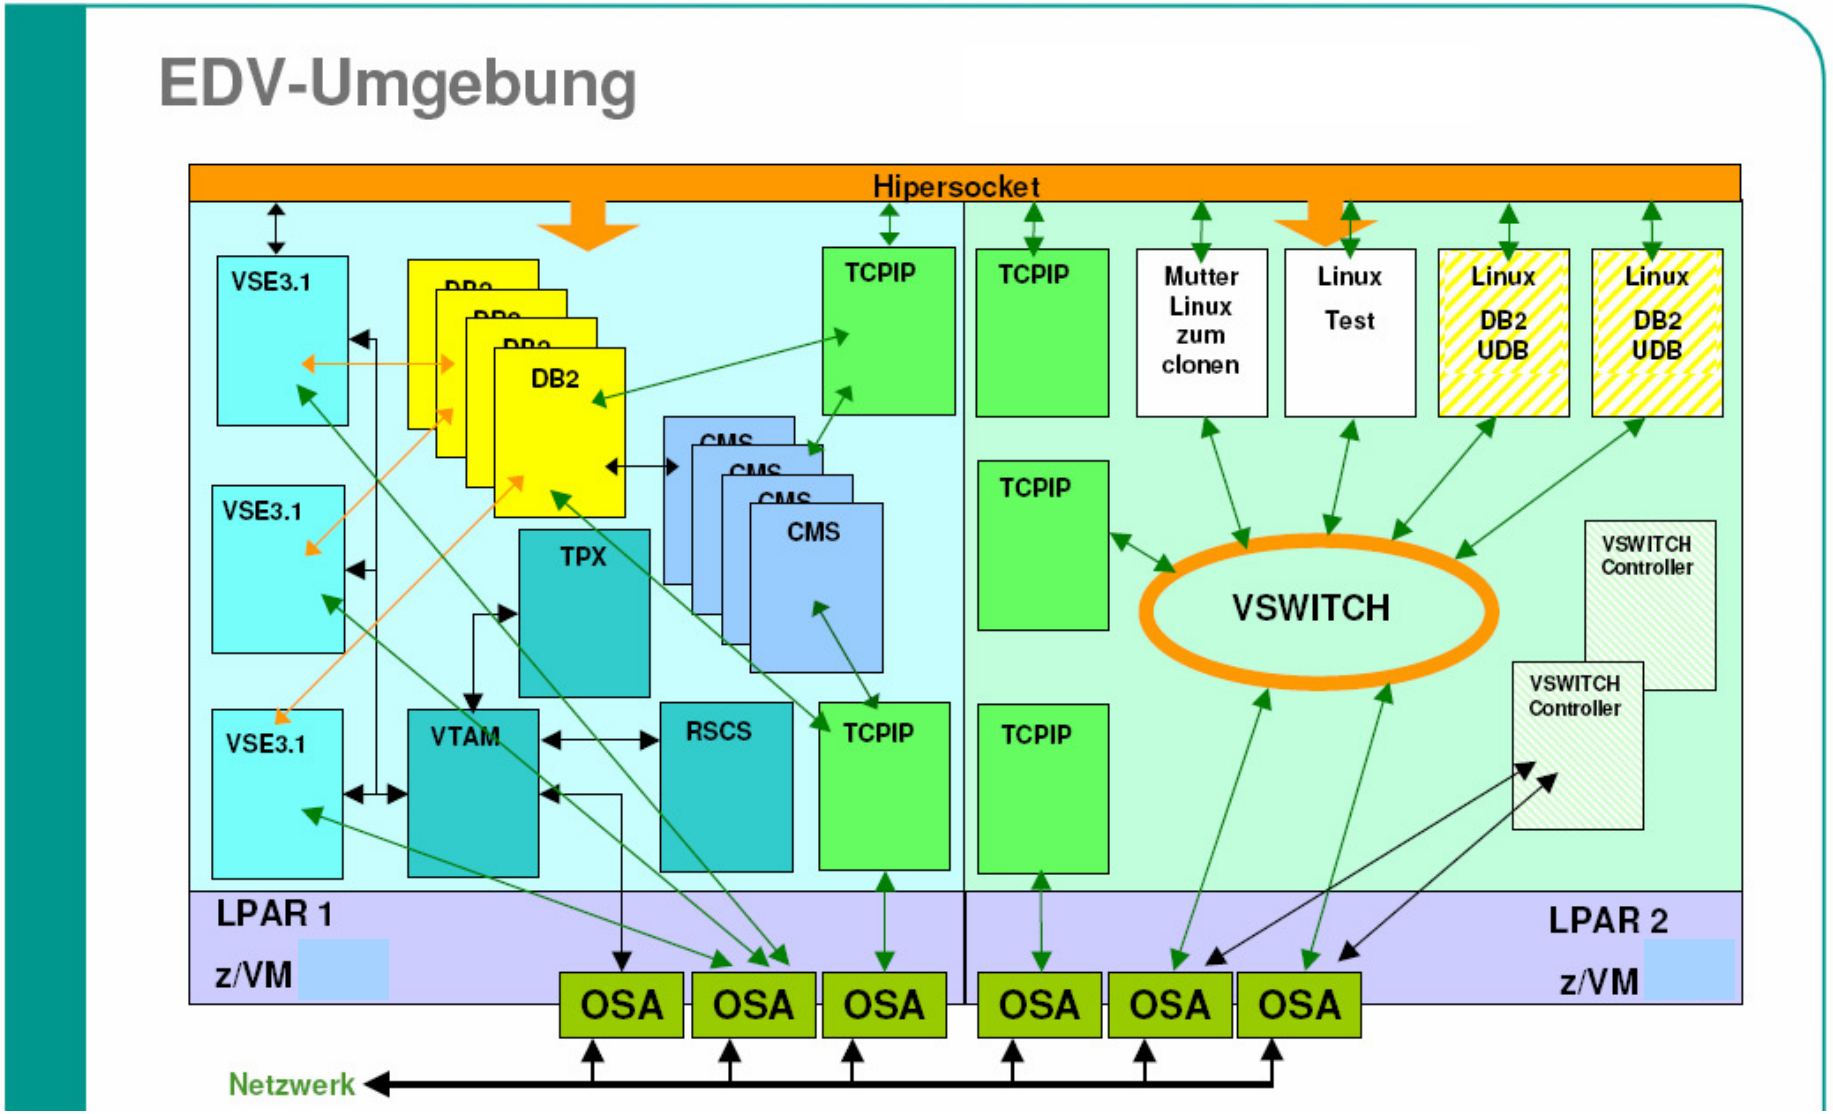

VSE Customer References(1) Impol /ALCAD Slovenia

Design, Applications and Solutions

Alcad

Database design

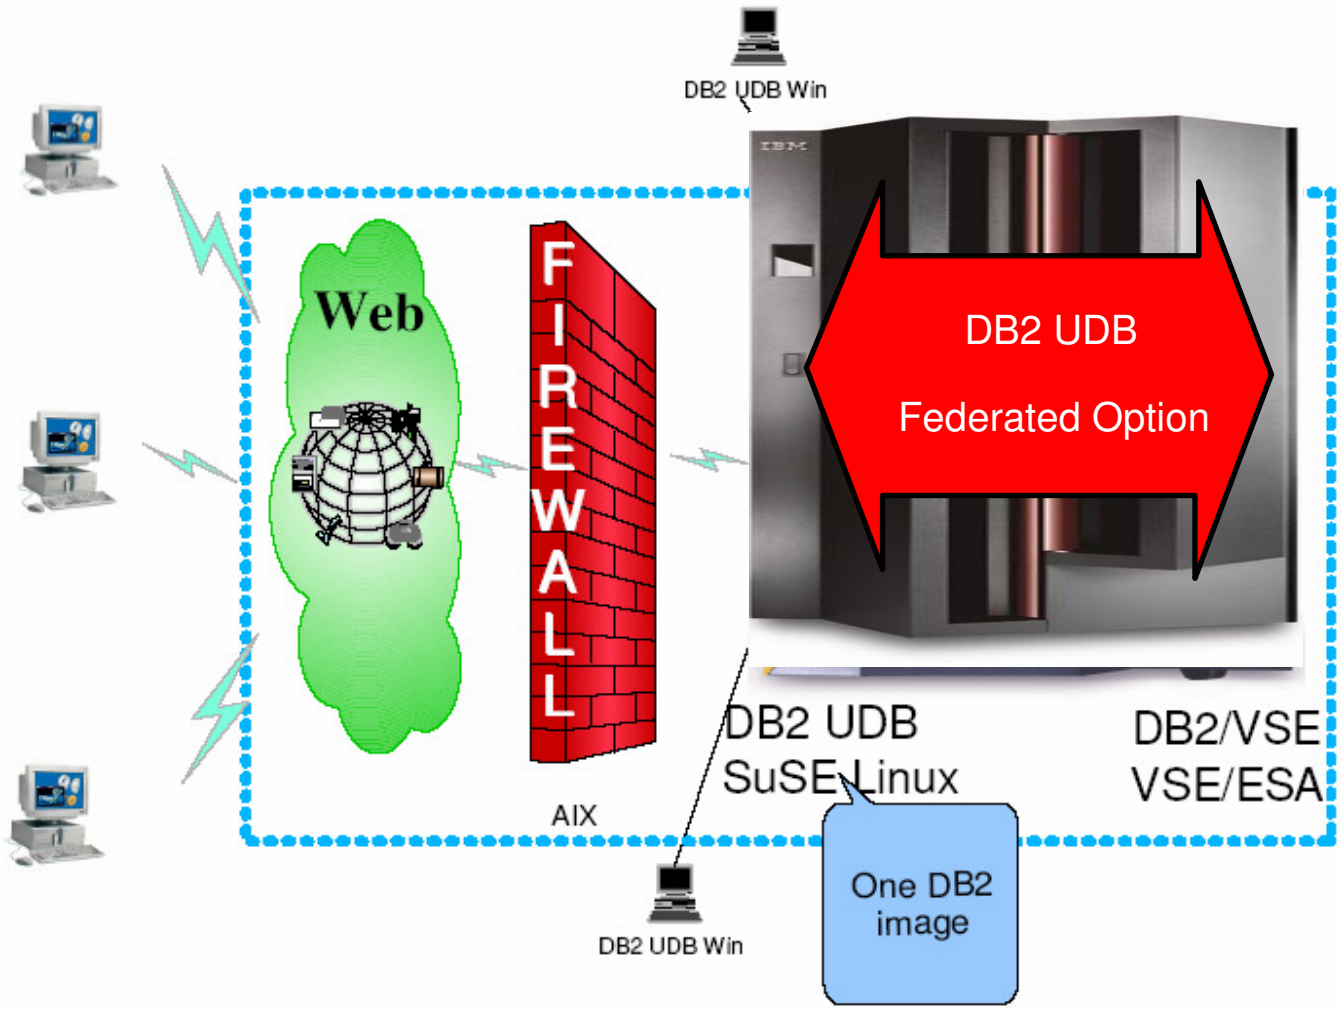

Customer Refrence (2): Pulsen, Sweden

PULSEN

Technical Platform

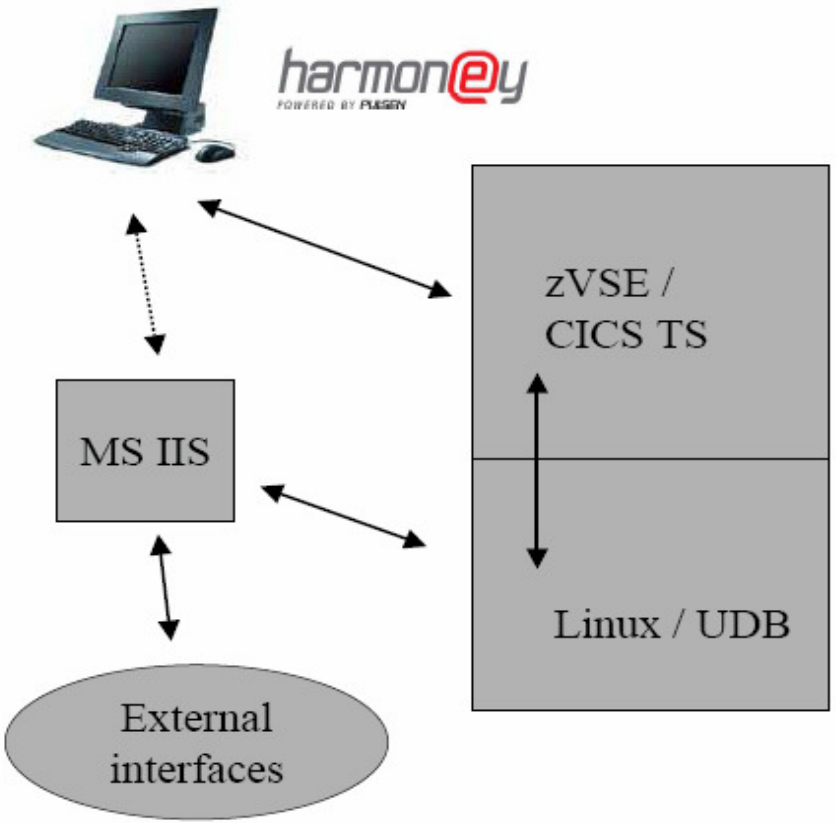

User interface – Windows/.Net

Data transfers between client and host in XML

CICS Web Services

DBM - DB2 UDB under Linux

Business logic in z9BC, partly "traditional" PL/1 programs, partly Stored Procedures / UDFs in UDB

Customer Refrence (3): Securex, Belgium

DB2 linux

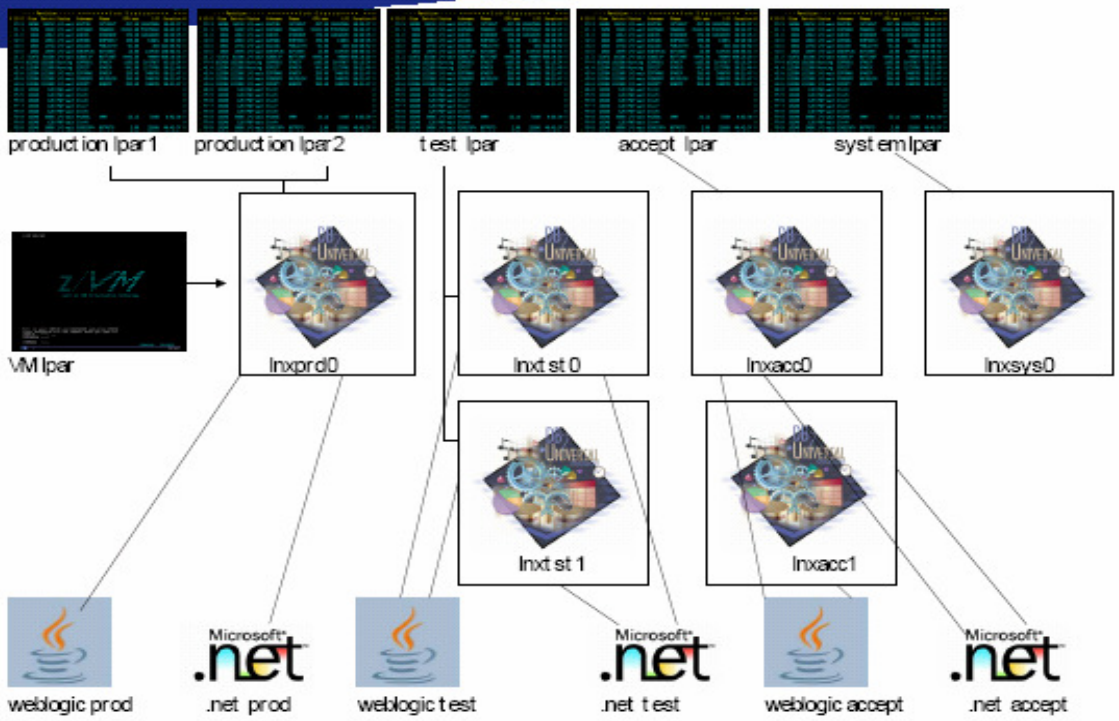

Other

QUEST SOFTWARE · sas · Health Center · webMethods · DbVisualizer

securex

Customer Refrence (5): Supreme Court, USA

The Magistrate Environment Today

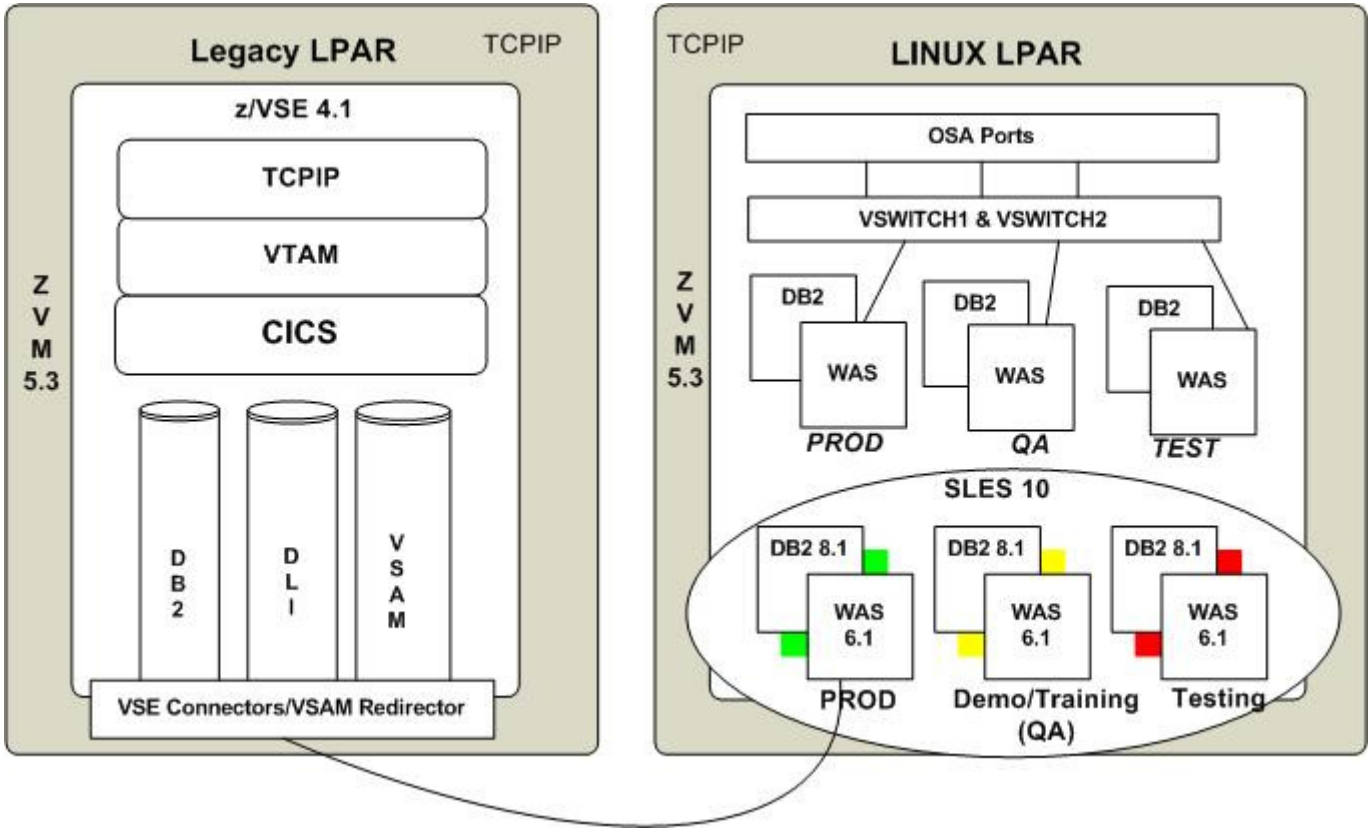

125 locations
2,800 processes per day
Direct interface with CMS application systems

The Lessons Learned (a work in progress)

- Have a plan! Linux on System z gets along well with everyone so long as you involve them.... Network, remote apps.....
- Document and then document some more
  - WAS settings
  - Passwords (root, wasadmin, wasmon, db2inst1 etc etc)
  - FAQs – build and maintain to help the next in line
- Managing and controlling changes for application deployments and system fix packs?
  - Test / QA / Production – keeping things in sync

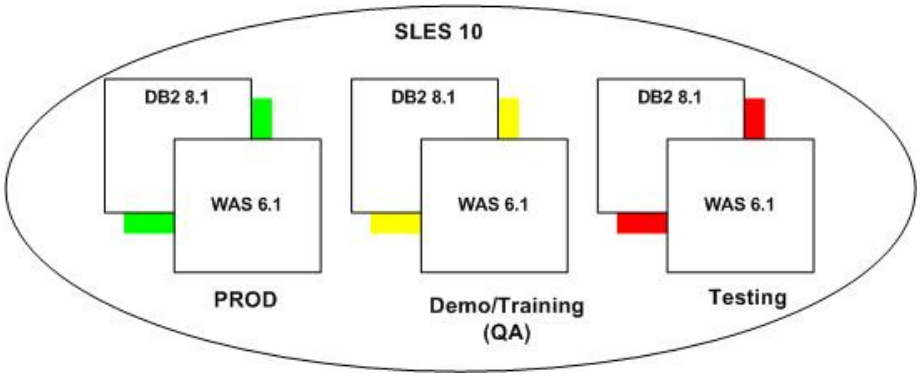

The Lessons Learned (con't)

- Have a good monitor and know what it's telling you
  - Helps with sizing and tuning
  - Quickly pinpoints out potential or growing problems areas
  - Virtual Disk works great for swap volumes
  - Shows management they are getting their money's worth

Customer Refrence (6): Olio Carli, Italy

Internal Connections

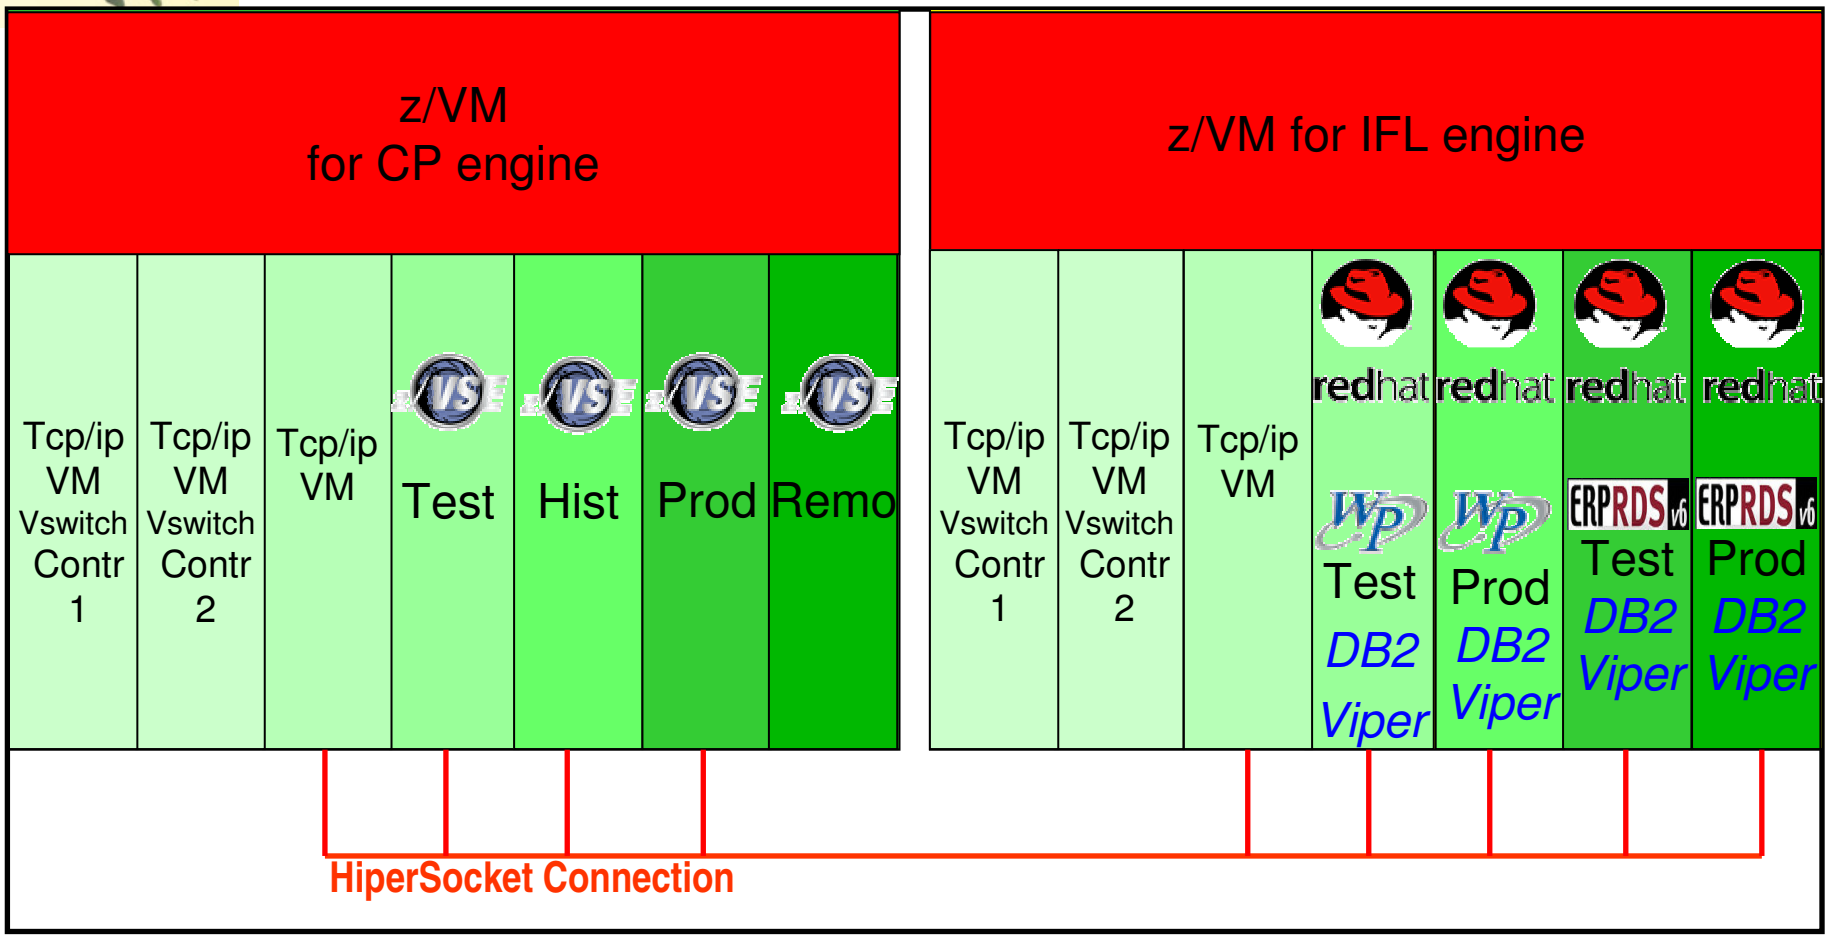

Olio Carli

More information

- DB2/Linux on System z

  http://www.ibm.com/developerworks/linux/linux390/perf/tuning_rec_database.html

  http://www.ibm.com/developerworks/data/library/techarticle/dm-0509wright/

- DB2 Server for VM and VSE

  http://www-01.ibm.com/software/data/db2/vse-vm/

- Documentation

  http://www-01.ibm.com/software/data/db2/vse-vm/directory.html#VSE7.5

  http://www-01.ibm.com/support/docview.wss?rs=71&uid=swg27009727

- Redbooks:

  http://www.redbooks.ibm.com/