



E35

User Experiences with z/VSE Connectors

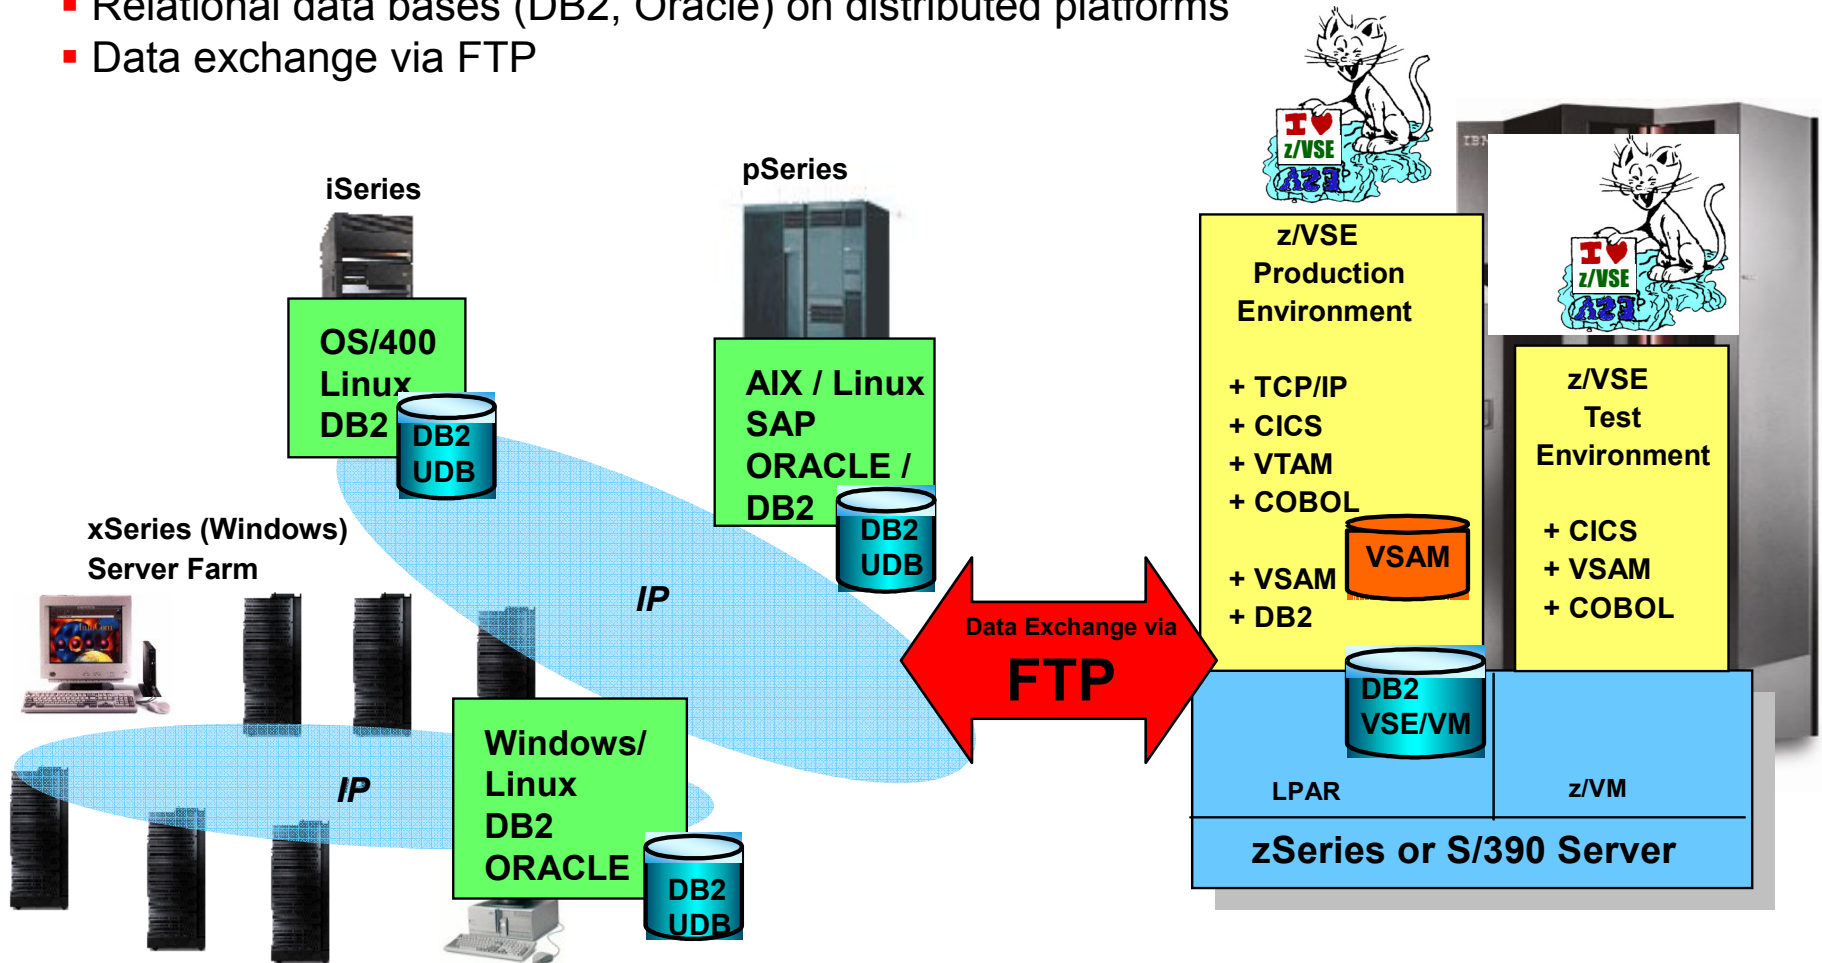
Wilhelm Mild

IBM
SYSTEM z9 AND zSERIES EXPO
October 9 - 13, 2006

Orlando, FL

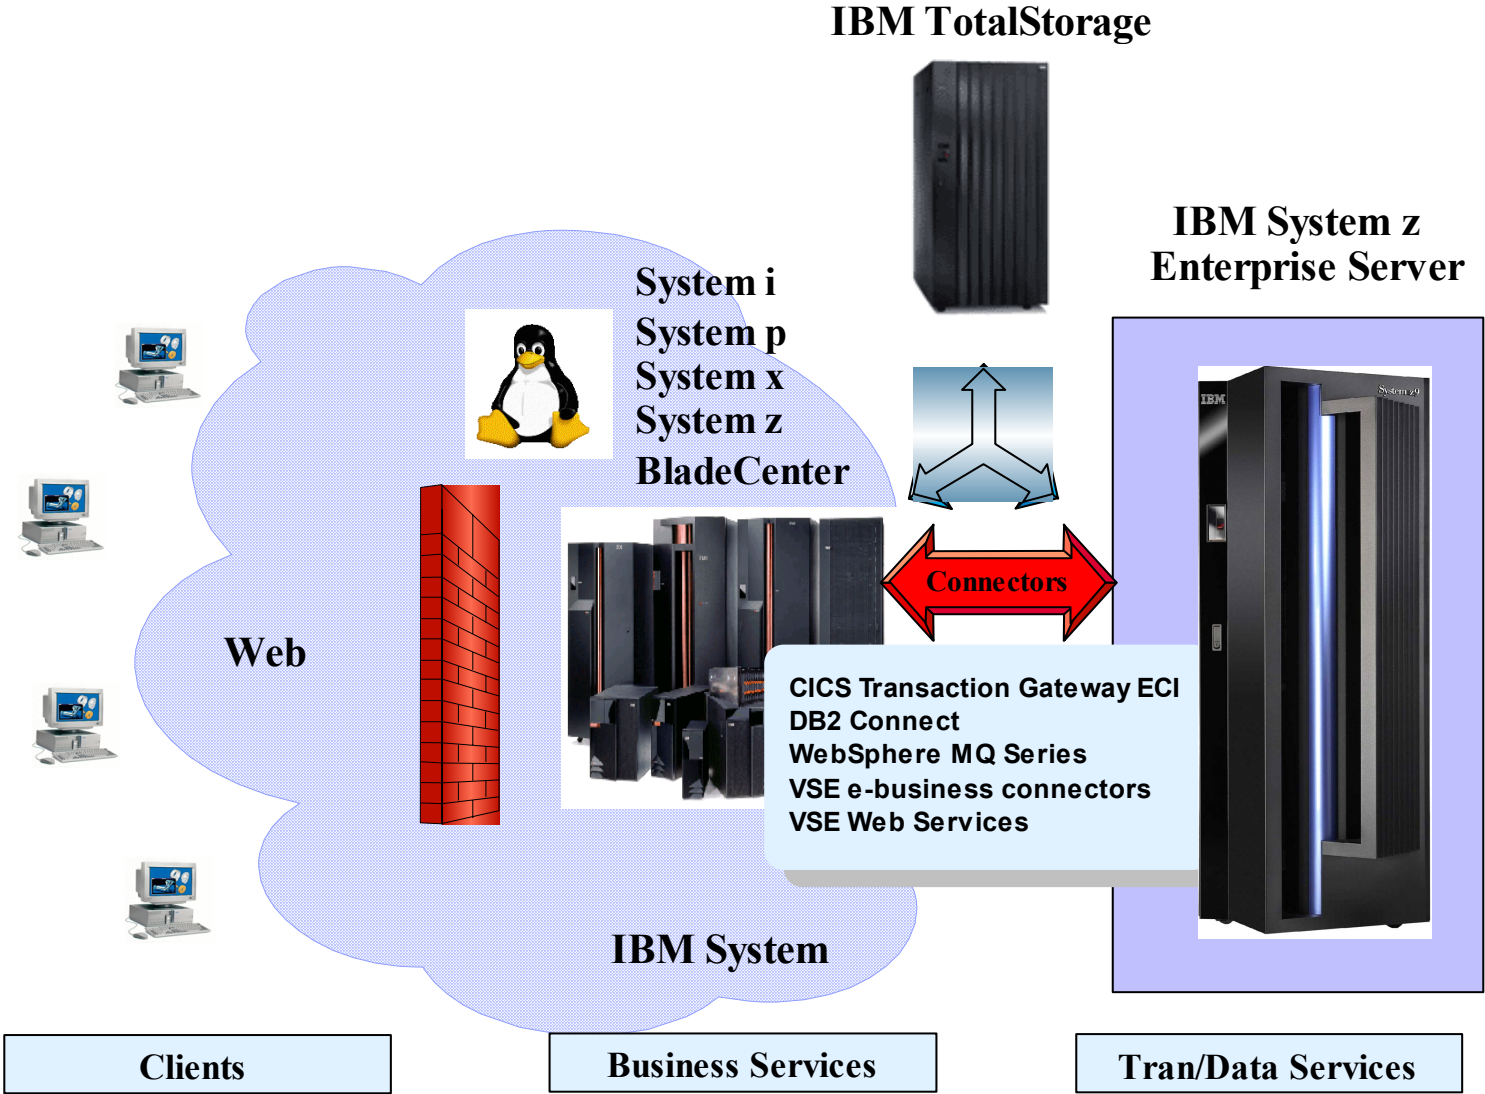
A typical VSE Customer Environment

- Various servers (System z, System p, System i, System x, and competitive)
- CICS and batch programs on VSE
- VSAM data on VSE (plus some DB2 and/or DL/I installations)
- Relational data bases (DB2, Oracle) on distributed platforms
- Data exchange via FTP

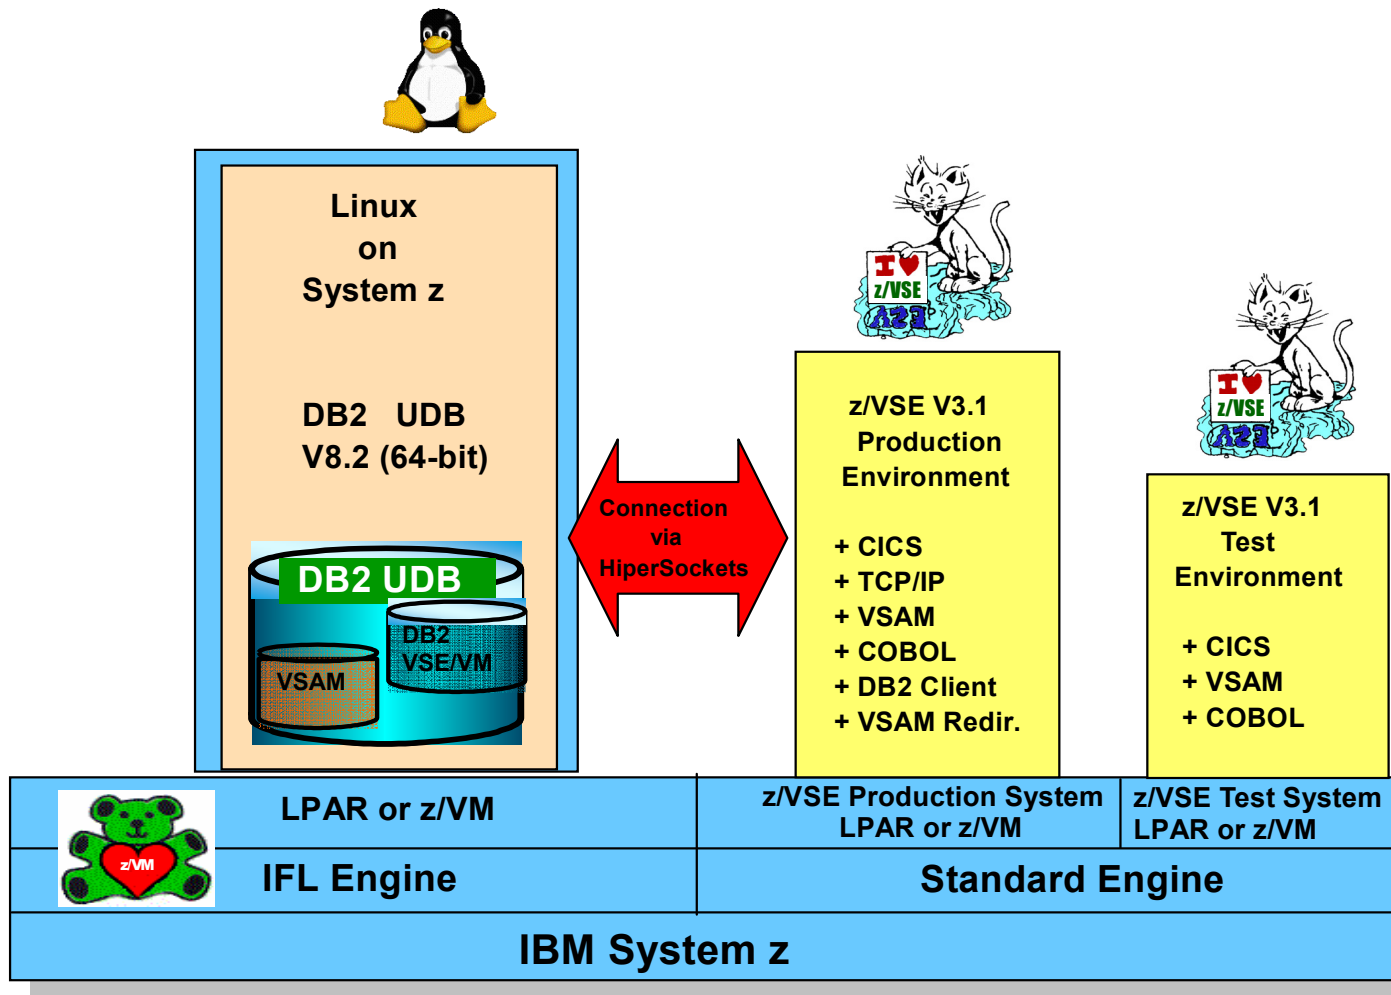


Most important requirement: Common data store!

VSE Strategy



Wanted: Common data store and data warehouse solution with DB2 UDB on System z



How to get started

Information to get started :

z/VSE home page -> solutions

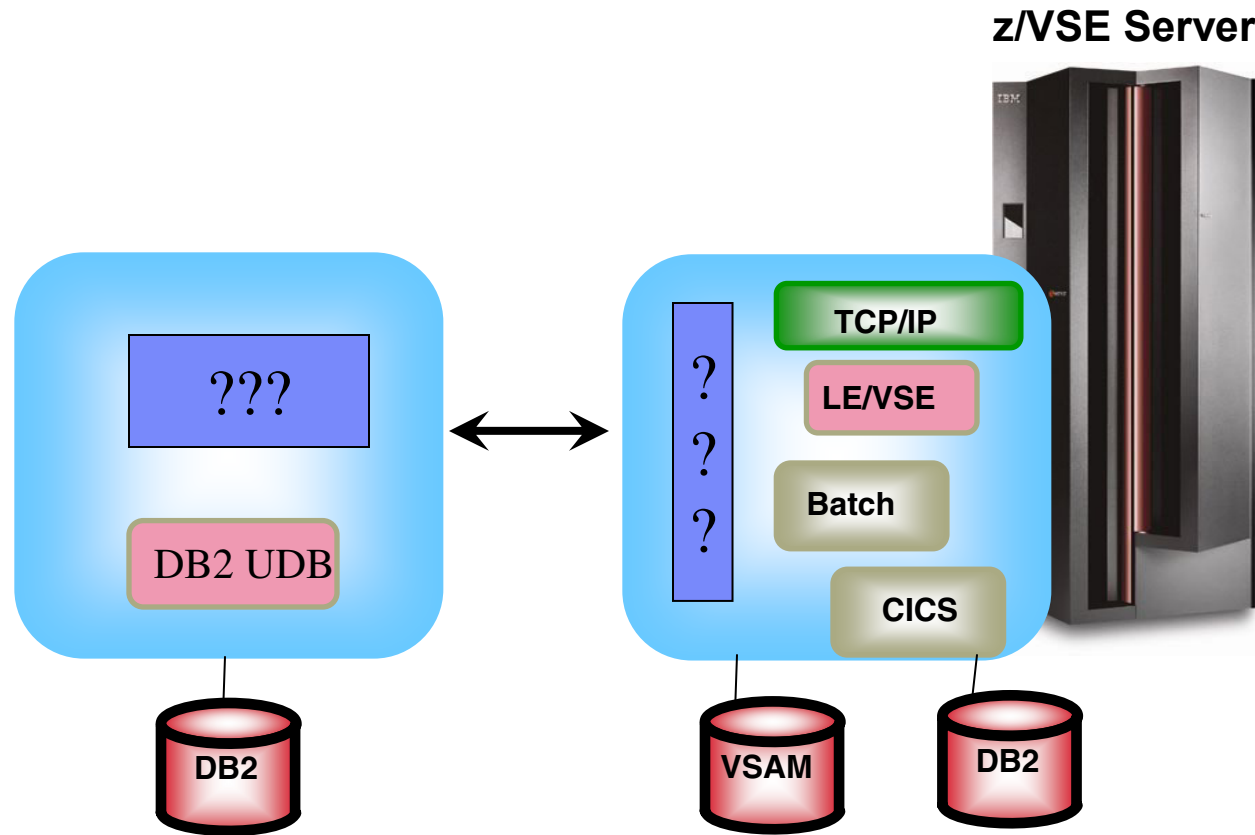
<http://www.ibm.com/servers/eserver/zseries/zvse/solutions>

High potential of modernisation exists but is unused because of:

- insufficient communication between departments of different platforms
- management structure inhibits projects in distributed environments
- lack of information about new possibilities
- first Steps are unclear (solutions link above can help)

External consultancy is recommended !!!

VSE applications and DB2 UDB on Linux on zSeries



What is needed to build a such solution

- options
- performance impacts
- risks

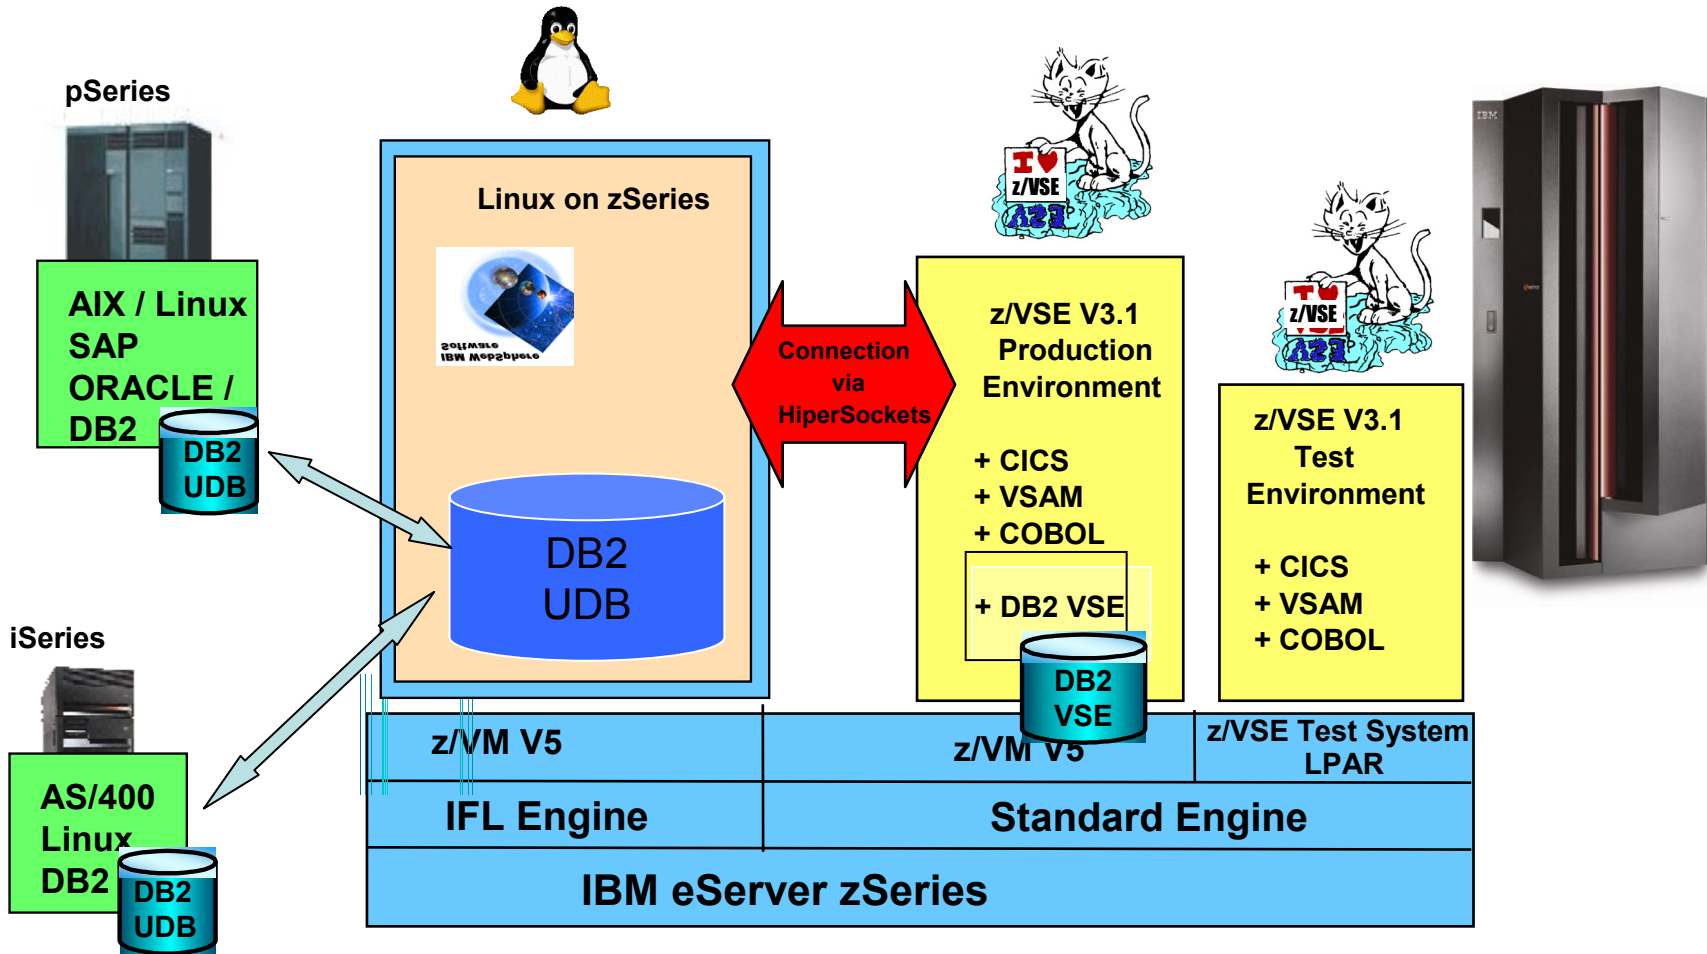
First Steps – first questions

Actual application and data store analysis:

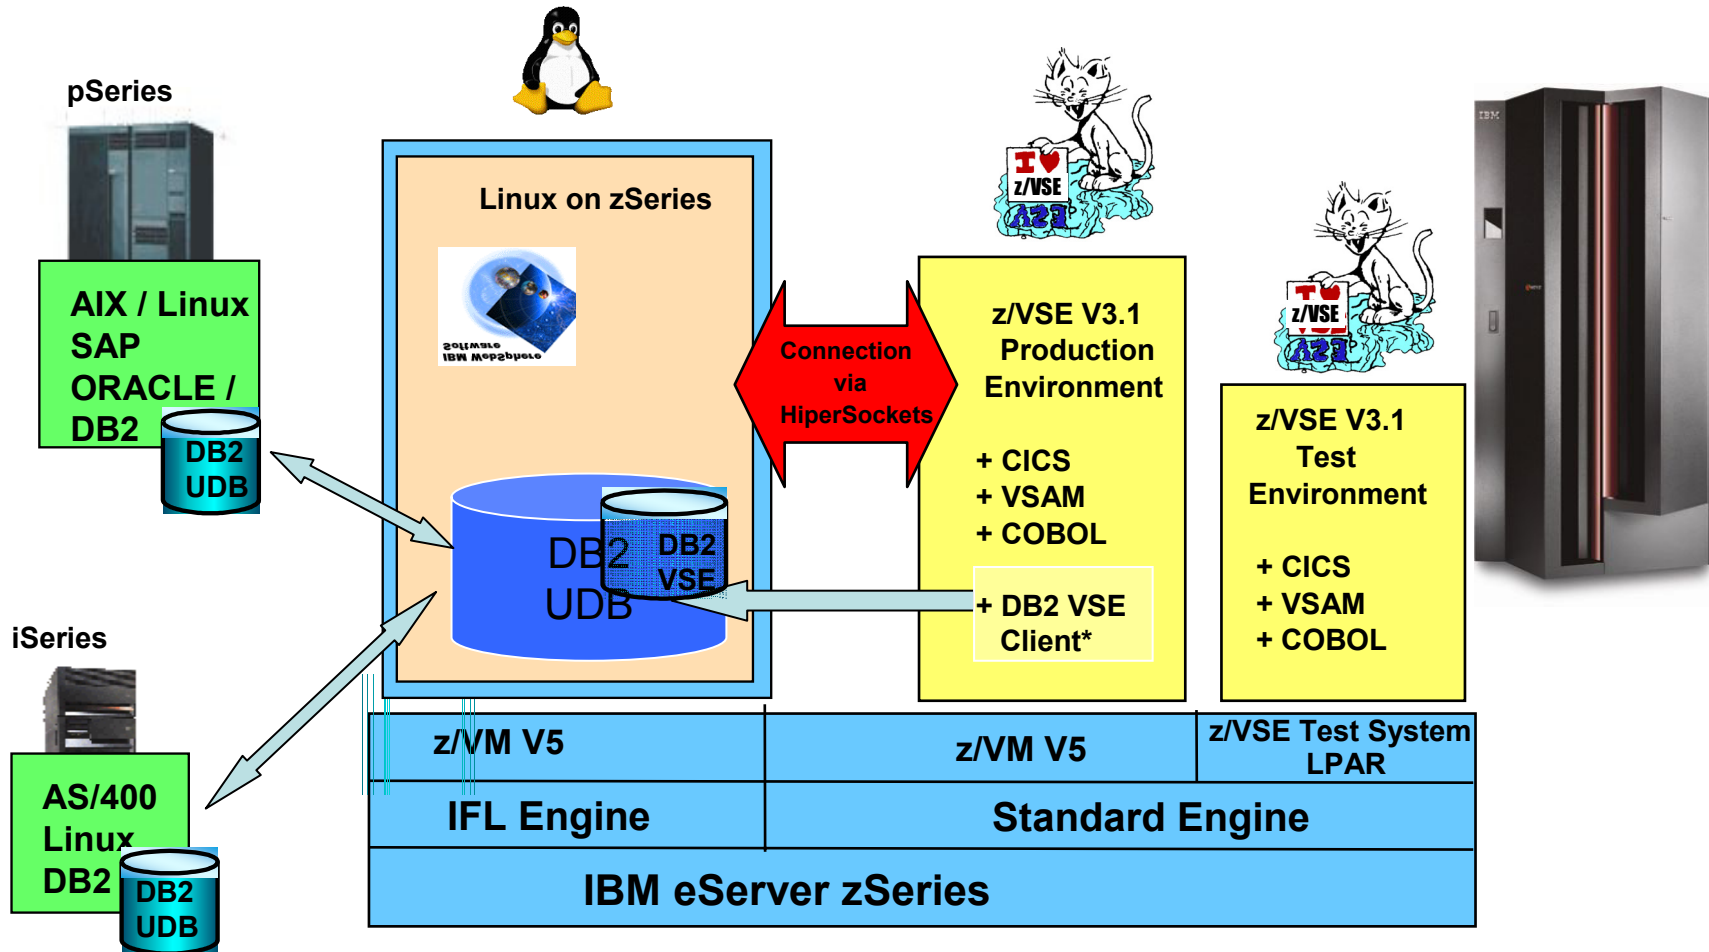
- Where are the data today
 - What is the future long term goal
- What is the amount of applications
 - Important in case of needed changes
 - Types of applications working with the data
 - Online CICS applications
 - Batch application workload
- Performance is always important – but how critical is a longer run of applications

Customer Implementation 1: DB2 VSE

Wanted: DB2 data consolidation – to DB2 UDB on Linux



DB2 VSE data consolidation – to DB2 UDB on Linux



(*) DB2 VSE Client – the client functionality only, can be obtained with [PRPQ P10154](#)

DB2 VSE and DB2 UDB on Linux on System z

Why customer wanted to use DB2 UDB on Linux on System z with VSE Core applications:

- Modern environment in DB2 UDB on Linux on System z
- Existence of lots of tools for:
 - database management
 - Optimization and Tuning
 - Data analysis (Warehouse, Mining, OLAP)
- ASCII environment – easy integration with distributed DB2 UDBs
- Consolidation of DB2 UDB databases from distributed platforms
- **Savings: DB2 CONNECT is not needed on Linux on zSeries**

DB2 VSE and DB2 UDB on Linux on System z

Arguments for DB2 UDB on Linux on zSeries with VSE Core applications:

- VSE applications access to DB2 UDB on Linux via HiperSockets
 - reliable network – no wires
 - fast network (memory copy speed)
 - transparent

- Core applications on VSE (CICS and batch):
 - can be used unchanged with considerations of EBCDIC – ASCII code pages (i.e. sorts with low values)
 - can show performance degradations if mass single row processing is done – these applications might need adoptions

Environment and Database design

Configuration steps for CICS applications and the remote DB2 UDB environment

■ VSE environment

- configure DB2 VSE database directory
 - configure ARISDIRD (IP, port, DBname of remote database)
- enable DRDA code (batch and online)
 - configure ARIS74LD (batch), ARIS745D (AR)
 - new transaction in CICS to bind packages (CBND) to remote AS (done during program preparation)

■ zLinux environment

- configure database manager on DB2 UDB zLinux
 - change some DBM parameters to allow implicit connect from within CICS
- configure VSE batch and ISQL options (create remote packages)
 - ARIISQL for ISQL and ARIDSQL for Batch

- **Note: DB2 CONNECT is not needed on Linux on System z**

Environment and Database design

Configuration for CICS applications and remote DB2 UDB database


Application considerations:

- migrate tables from DB2 VSE to DB2 UDB zLinux
 - UDB export/import options
 - use of federated DB2 UDB options and a cursor application

- existing CICS/DB2 VSE applications
 - no changes to the source code required (except Code page issues)
 - the SQL precompile creates new packages on the remote DB2 UDB)

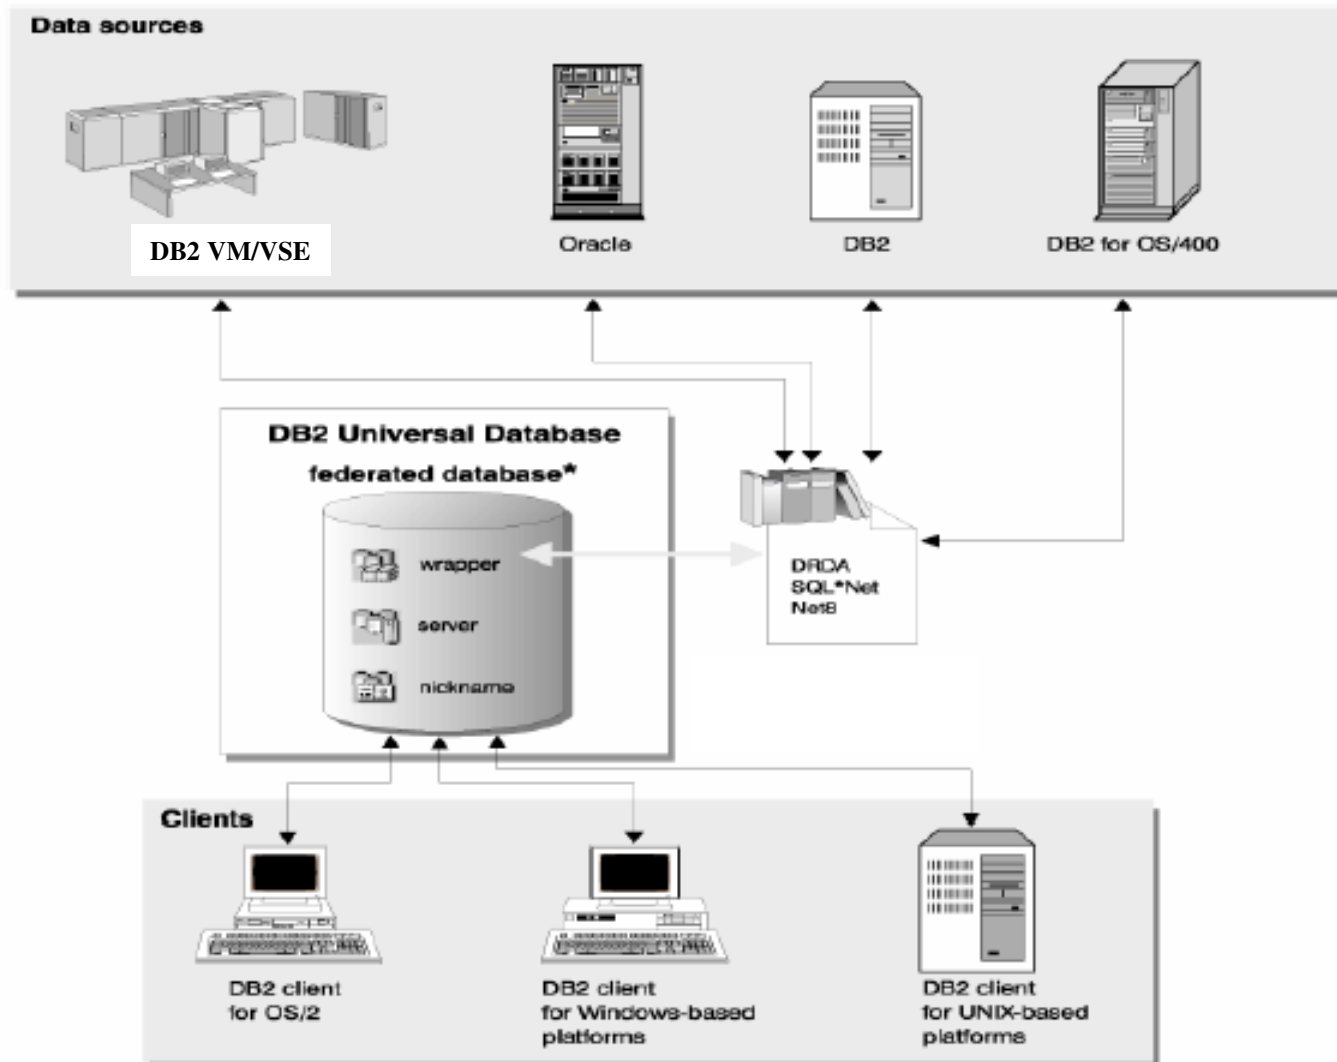
- existing VSE batch DB2 VSE applications
 - no changes to source code required
 - adapt CONNECT statements to access remote DB2 UDB

Move DB2 VSE data to DB2 UDB on Linux on System z via Federated option in DB2 UDB



- ▶ Minimum changes – maximum combination
 - ▶ DB2 UDB for Linux on System z with Federated Option – includes DB2 VSE logically
 - ▶ DB2 UDB Applications have transparent access to DB2/VSE

Federated Database design



Federated Database design

- Define federated database nicknames on the DB2 UDB LUW server for real tables that exist on a DB2/VSE&VM server
- Then new tables are created on the DB2 UDB LUW server, one for each nickname
- The table definitions are identical to the table definitions on the DB2/VSE&VM server.
- A small program with cursor is defined with a select statement from the nickname.
- The DB2 UDB federated server will get the table data from the remote DB2/VSE&VM server using a DRDA protocol query over a TCP/IP connection.
- The LOAD command specifies the cursor name, which results in obtaining the DB2/VSE&VM data, and the target of the load is the new table that was created in the DB2 UDB LUW server.
- The net result is that the table data is copied from the DB2/VSE&VM source table into the DB2 UDB LUW target table.
- The copy is performed in one step.
- The cursor select is fairly efficient since the DRDA blocking protocols can potentially return many table rows in each network send.
- The LOAD is efficient because it uses DB2 UDB fast load feature.

Federated Database design

Other methods moving the data using:

DBSU DATAUNLOAD and DBSU DATALOAD on DB2 VM/VSE

DB2 UDB EXPORT and DB2 UDB IMPORT commands.

These have the disadvantage of requiring two distinct steps.

The DBSU DATALOAD option also requires a DRDA connection to load the data which will be transmitted one row at a time and inserted into the DB2 UDB LUW database one row at a time. This is a huge network overhead and will also incur all of the logging overhead of each insert.

The performance test done of transferring data from a DB2/VSE database to a DB2 UDB for Linux on zSeries server is shown in the table below.

The test consisted of copying the data from 4 different tables of various sizes.

The table below shows the result of running each of the data movement utilities. A is a 510,000-row table; B is a 1,530,000-row table; C is a 10,000-row table; and D is a 1000-row table. Your actual results will vary based on the size and speed of CPUs on your hardware, load on your systems and network and the amount of data that you are moving. The table just shows the relation of the 3 methods of moving the data.

	VSE	VSE	Linux	Linux	Linux
	DATAUNLOAD	DATALOAD	EXPORT	LOAD	CROSS LOADER (federated)
	(SQLDS)	(DRDA)	(DRDA)		
A	1383s	1718s	46.276s	6.1s	44.642s
B	4238s	5084s	235s	15.3s	148.128s
C	29s	57s	1.2s	1.7s	2.3s
D	4s	17s	0.17s	1.5s	1.5s

Experience

Applications can have big variations in execution time

Application considerations:

- a high percentage of applications worked with the remote DB2 UDB on Linux at least as fast as with the local DB2 on VSE
 - lots of sorts slow down the performance
 - ineffective SQL statements show long running times

- DB2 utilities should not be considered for test
 - the productional processes showed an overall increase of 20% in response time

- VSE resources showed higher CPU utilization
 - higher CPU utilization were seen in TCP/IP partition
 - the DRDA Code is running with the application – therefore in high used CICS times – CPU increased in CICS partition

Test Results in a customer proof of concept
VSE application running against
DB2 UDB on Linux on System z

Testcase	Time	Factor	Hardware
SQLDS	5:40	1	z890 - 160
UDB LE	8:40	1,53	z890 - 160 + IFL
UDB (DRDA Asm IF)	6:55	1,22*	z890 - 160 + IFL
SQLDS	3:36	0,64	z890 - 170
UDB (DRDA Asm IF)	4:23	0,77*	z890 - 170 + IFL

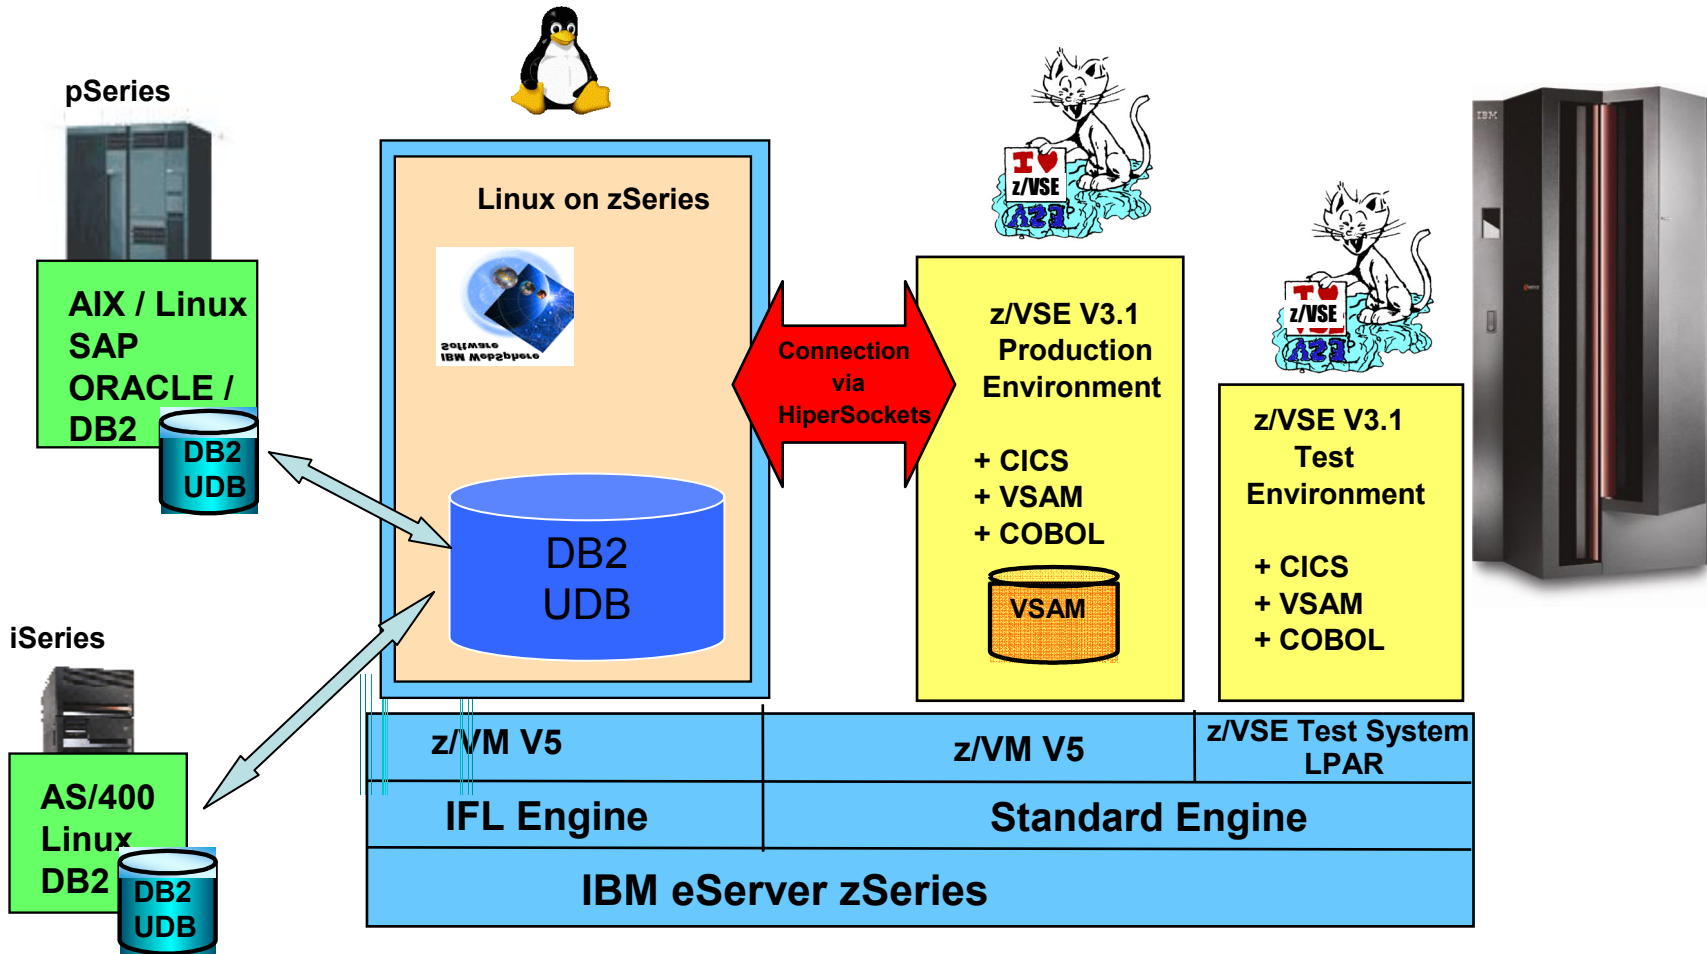
DRDA Asm IF – is the Assembler interface for DRDA in VSE (available via DB2 PTF)

*an overall increase of 20% response time was measured

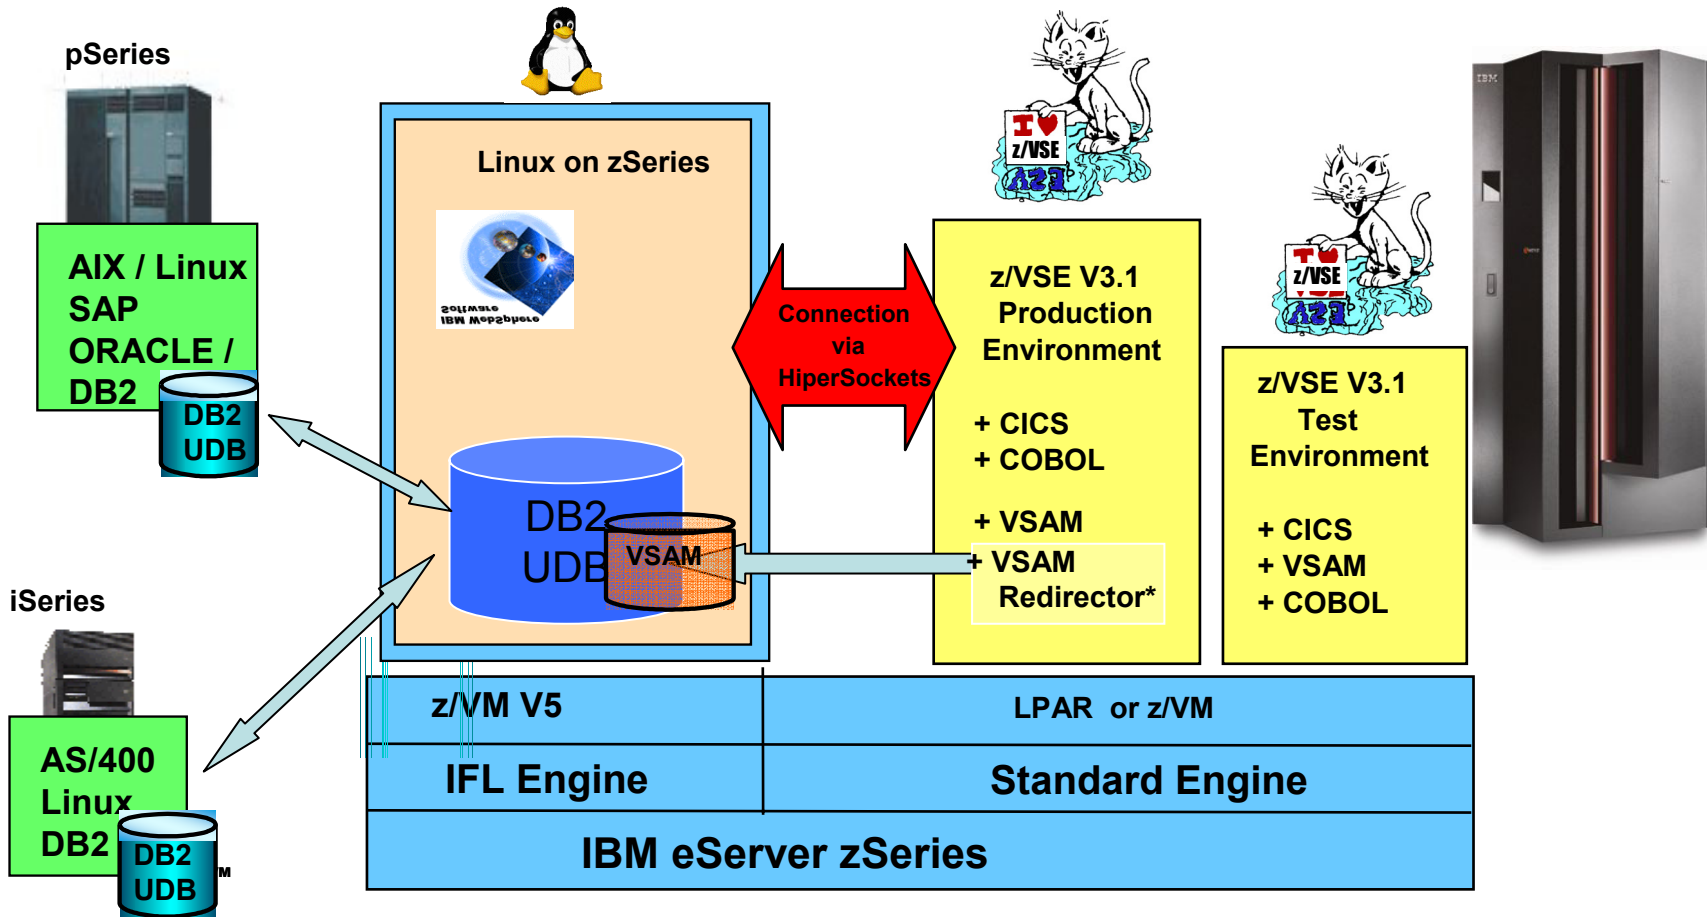
Conclusion: Every application base is different and no general response time statement can be made

Customer Implementation 2: VSAM data

Wanted: DB2 data consolidation – to DB2 UDB on Linux



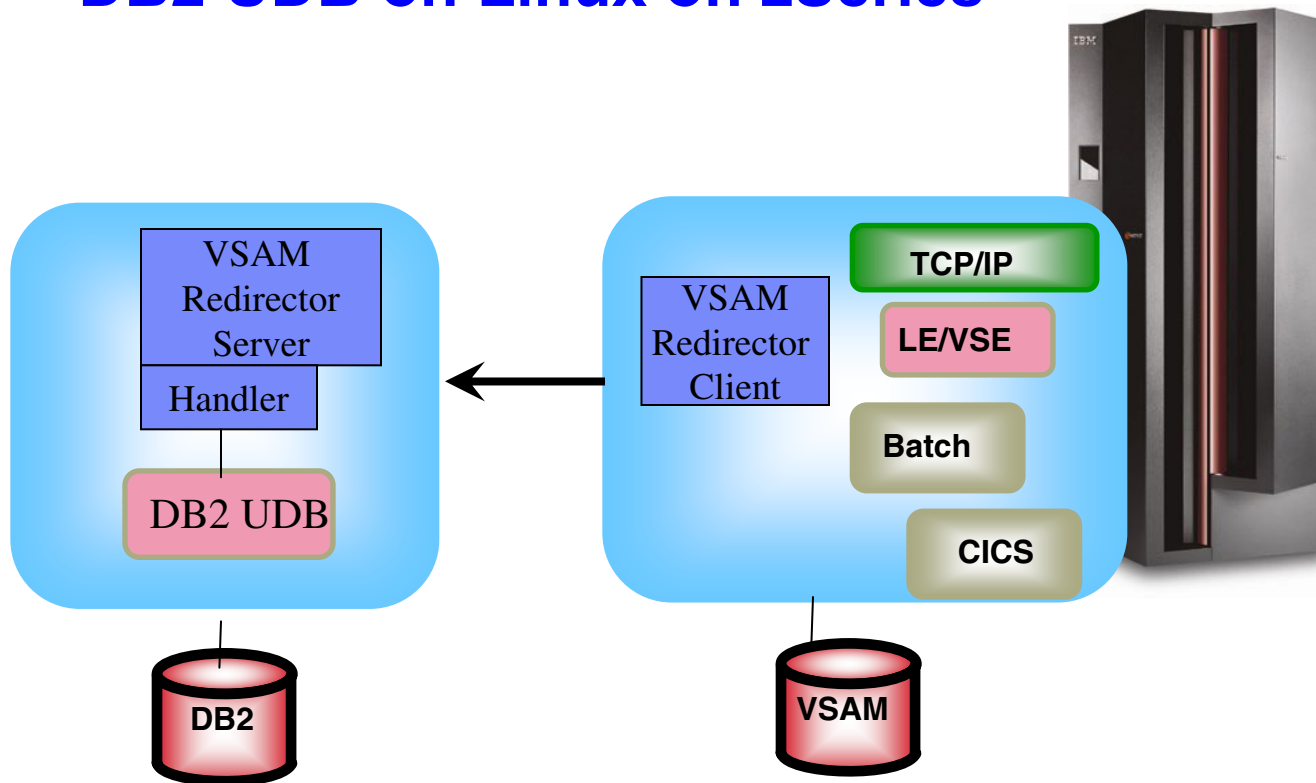
Transparent Work of VSAM Programs with DB2 UDB on Linux on System z



(*) VSAM Redirector – Common data store solution – with DB2 on Linux on zSeries Solutions without changes to VSAM programs

VSE/VSAM applications access to DB2 UDB on Linux on zSeries

z/VSE Server



Redirector considerations – data access

Redirection decision points based on the different flavors and needs

- 1) Redirection with DB2 access only in real time
 - ✓ For Online CICS applications recommended
 - ✓ Might have an impact on performance for mass data process in batch

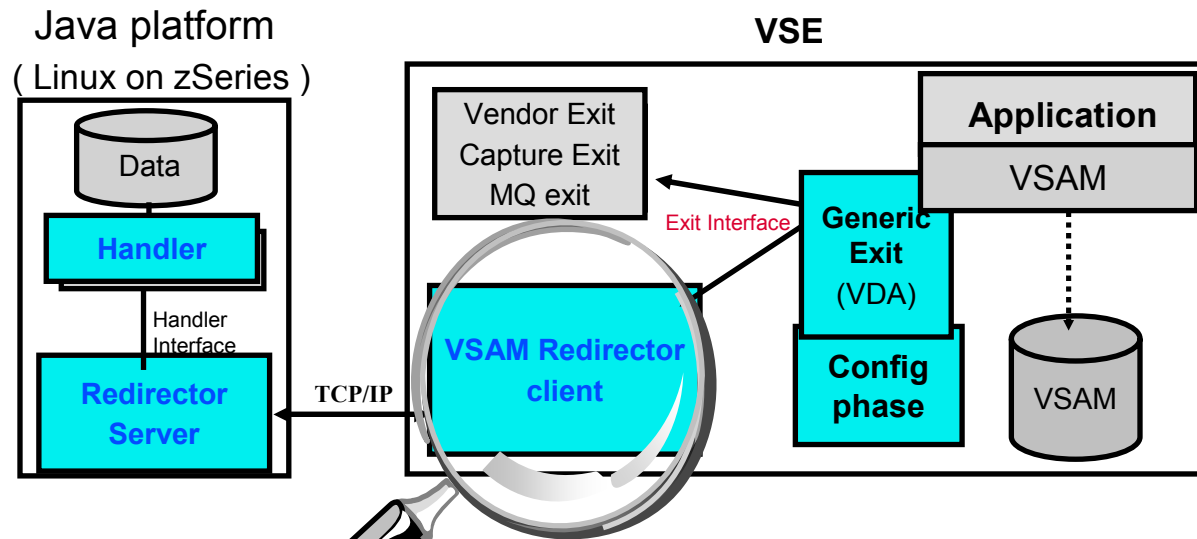
- 2) Redirection with synchronization of VSAM and DB2
 - ✓ For propagation of changes into DB2 on Linux
 - ✓ Less Impact on VSE performance

- ✓ Special Considerations for program failures (crashes)
- ✓ Network reliability (except Hipersockets)
- ✓ Dependent on database availability and response time

A proof of concept or pilot implementation is always recommended !

VSAM application working with DB2 UDB

Remote Access / synchronization



- ▶ **VSAM Redirector Client / Redirector server / Handler**
 - ▶ Redirector client - running in the context of calling program
 - ▶ Redirector server / Handler – are Java components
- ▶ Network dependent approach (fastest with System z Hipersockets)

Experiences for real time VSAM redirection

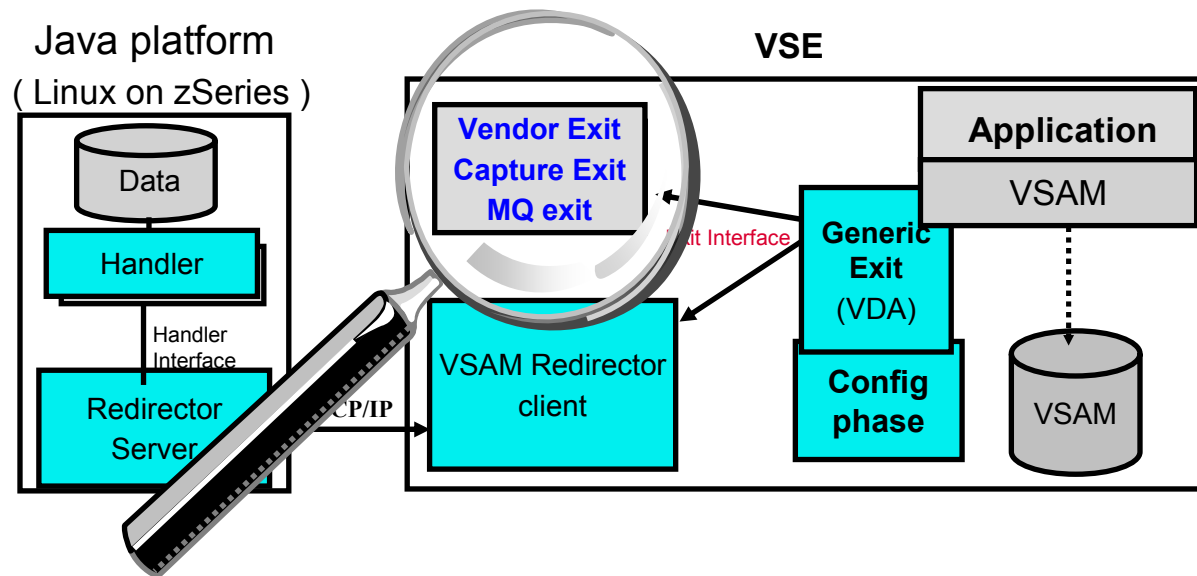
Real time Redirection enables real time processes across systems.

- 1) Data distribution to data warehouse in real time
 - ✓ Online CICS applications Update DB2 real time
 - ✓ No VSAM data anymore
 - ✓ VSAM recovery processes must be replaced with DB2 recovery combination

- 2) Redirection with synchronization of VSAM and DB2
 - ✓ For propagation of CHANGES only into DB2 on Linux
 - ✓ Data held in VSAM and UPDATES are propagated to DB2 .
 - ✓ Two phase Commit mechanism to keep VSAM and DB2 in sync

VSAM Data collection / transformation / journaling on VSE

Vendor Exit



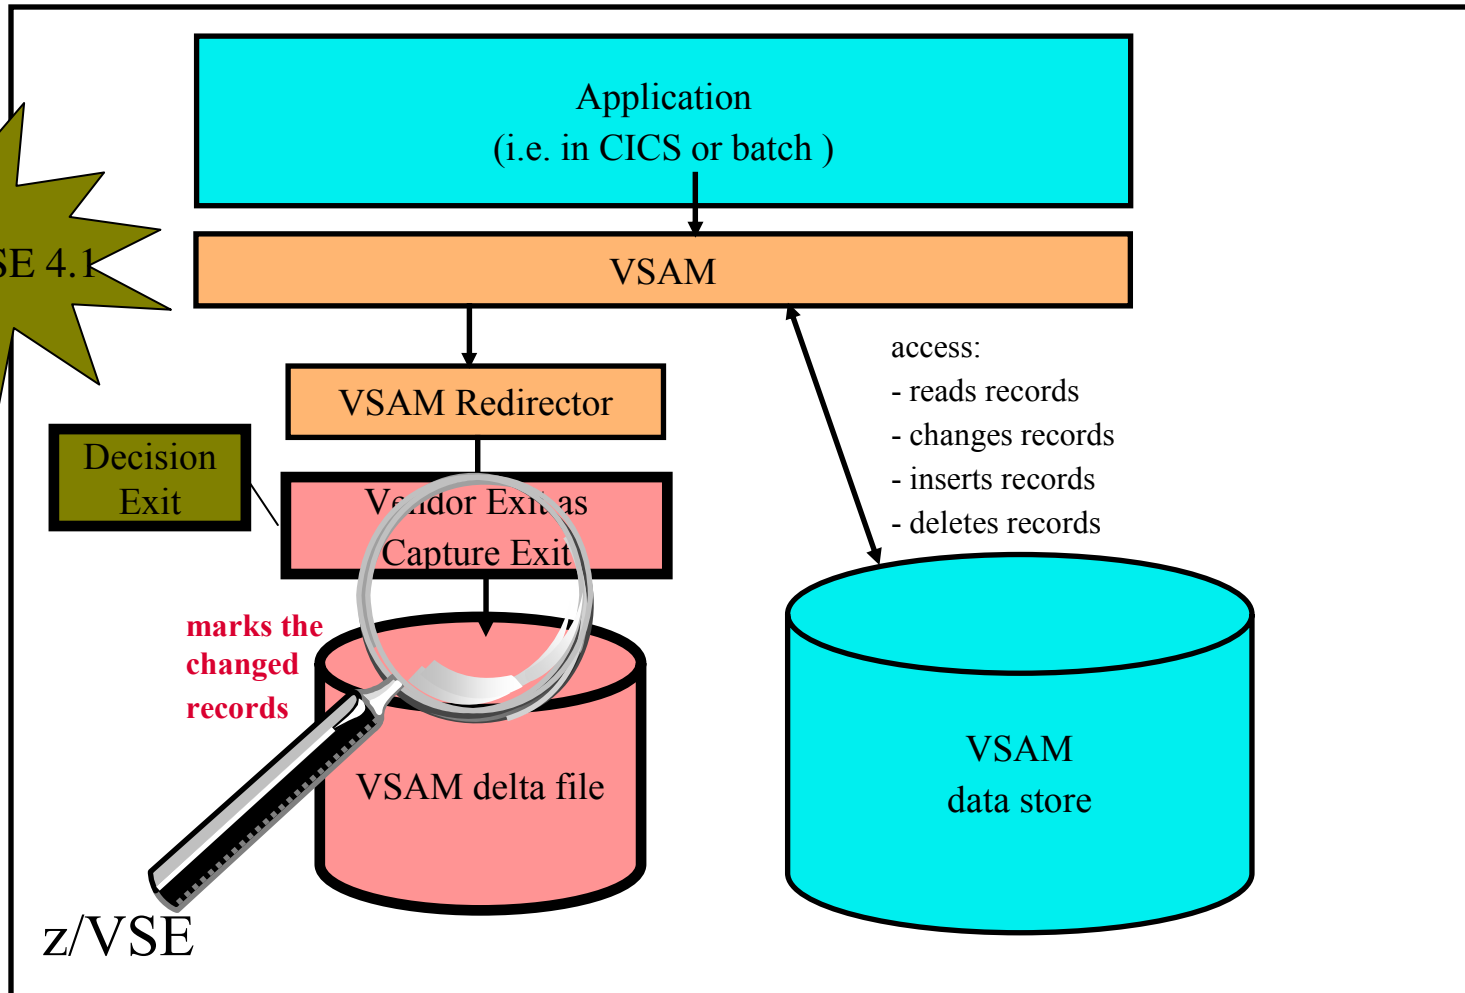
► Vendor Exit

- user (vendor) written phase for data collection/transformation
- has to comply with the documented **Exit Interface**

Note: No chaining of Vendor Exit with VSAM Redirector client supported

Redirection with Capturing

New in z/VSE 4.1



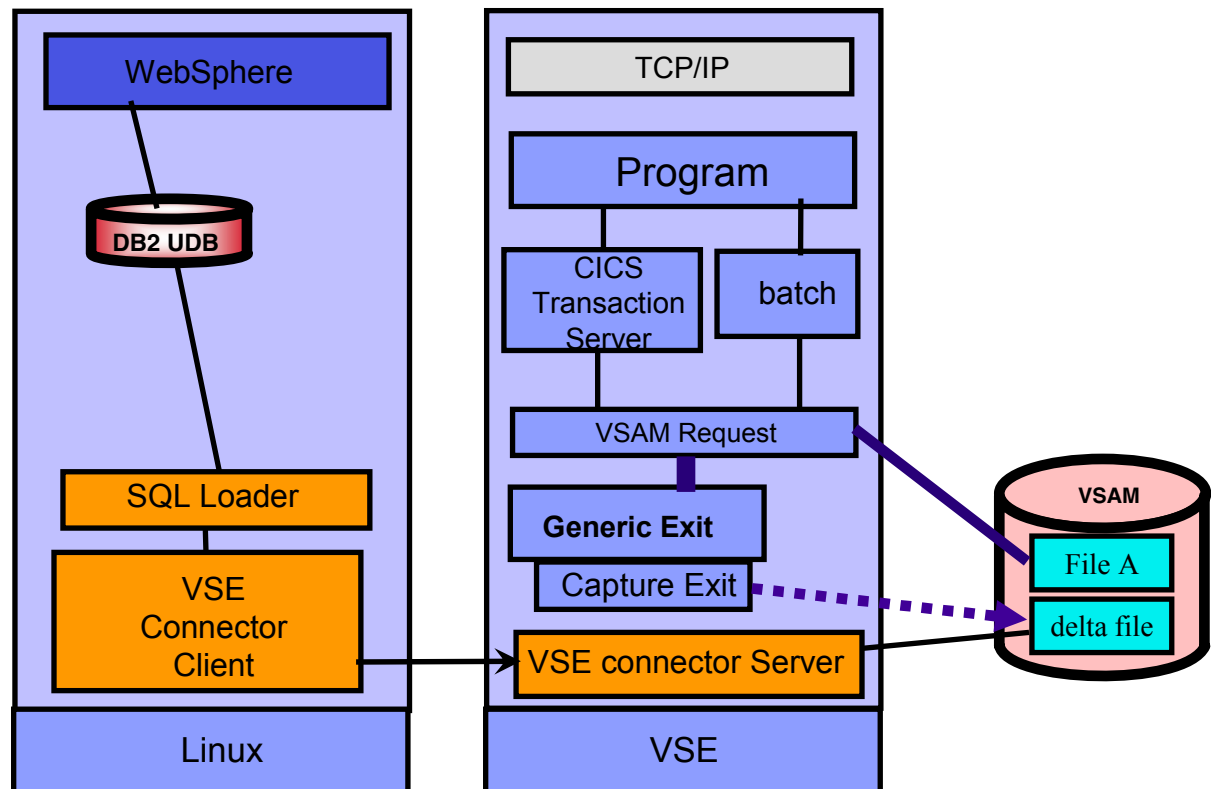
z/VSE

Capture Exit and Java-Based Connector

Asynchronous data pull from remote

Reduce network traffic, save time

- ▶ Collect the changed records in a separate VSAM file
- ▶ Possibility of cleansing
- ▶ Connector Client reads the delta file and inserts them into a database
- ▶ Transparent Journaling of data changes if wanted



Experiences for Redirection and Capture Exit

When VSE is the main production environment – performance is highest priority

- ✓ With Redirection and Capture – the VSE environment is decoupled from the database response time and network.
- ✓ Very low performance impact of the overall VSE processing
- ✓ The Capture approach, combined with the ‘Pull’ scenario establishes a complete solution for VSAM data distribution
- ✓ The customer could automate data distribution together with the realization of a common data store in a relational database.
- ✓ With the pre-announcement, we specified that in z/VSE 4.1 the Capture functionality will be enhanced with additional decision exits and database handling methods

Address <http://www.ibm.com/servers/eserver/zseries/zvse/>

Links IBM Business Transformation Homepage

Country/region [select] | Terms of use

IBM

Home | Products | Services & solutions | Support & downloads | My account

Servers > Mainframe servers > Operating systems >

z/VSE

About VSE

How to buy

News

Events

Solutions

Products & components

Documentation

Service & support

Downloads

Education

Partners

FAQ

Contact VSE

Related links

- Linux on zSeries
- z/OS
- z/VM
- IBM Storage
- IBM Printing Systems

z/VSE

z/VSE is designed to help provide robust, cost-effective solutions for customers with a wide range of capacity needs, in most industries, worldwide. z/VSE is built on a heritage of ongoing refinement and innovation that spans four decades. It brings the value of innovative IBM eServer zSeries and IBM TotalStorage technology to VSE clients.

Learn more

- [About VSE](#)
- [News](#)
- [History of VSE](#)

We're here to help

Easy ways to get the answers you need.

[E-mail us](#)

Mark your calendar

Guide Share Europe
April 18-20, 2005
Berlin, Germany

[Register](#)

WAVV conference
May 20-24, 2005
Colorado Springs, Colorado, USA

[Catch the WAVV](#)

Spotlights

- [IBM eServer zSeries](#)
- [Infrastructure simplification](#)
- [VSE Recommended Service Level](#)

Middleware

- [WebSphere software](#)
- [Information management software](#)

Announcing z/VSE V3.1

Built on a heritage of ongoing refinement and innovation that spans four decades.

Redesigned z/VSE homepage

You may have already noticed that the z/VSE home page has changed. We've redesigned this entire web site and included additional information. The objective is to provide you with a more useful business tool, as well as to offer you a more enjoyable experience. We encourage you to use, or to simply explore, the enhanced z/VSE web site. If you have questions, suggestions, or comments, please contact the [VSE team](#).

z/VSE Version3 Release 1

[z/VSE Version 3 Release 1 \(z/VSE V3.1\)](#) is designed to support:

- [IBM @server zSeries 890 and 990](#) (31-bit mode only)
- SCSI disks attached to zSeries FCP channels
- [OSA-Express2](#) and [FICON Express2](#) adapters
- [Crypto Express2](#) and CP Assist for Cryptographic Function (CPACF)
- IBM TotalStorage [3494 Virtual Tape Server](#)
- improved support for [IBM 3494 Tape Library](#)
- IBM TotalStorage [DS8000](#) and [DS6000](#) series Storage Servers
- enhanced Advanced Copy support

z/VSE is designed to enable network integration and infrastructure

New Web presence: ibm.com/servers/eserver/zseries/zvse

Additional Information

- z/VSE Home Page

<http://www.ibm.com/servers/eserver/zseries/zvse/>

- z/VSE Solutions

<http://www-1.ibm.com/servers/eserver/zseries/zvse/solutions/>



- e-business Solutions for VSE/ESA

SG24-5662

- e-business Connectivity for VSE/ESA

SG24-5950

- CICS Transaction Server for VSE/ESA

CICS Web Support

SG24-5997-00

- WebSphere V5 for Linux on zSeries Connectivity Handbook

SG24-7042

We appreciate your comments at : zvse@de.ibm.com