

Linux for zSeries Connectors Architecture Business Intelligence and ERP Solutions

ITSO Workshop - 2003



Viviane Anavi-Chaput viviane@us.ibm.com
Gerard Laumay gerard.laumay@fr.ibm.com

Workshop background



Workshop content based on the following Redbooks:

- Java Connectors for Hybrid Environments on zSeries - Linux to z/OS & VSE
 - SG24-7042 in development
- e-Business Intelligence: Leveraging DB2 for Linux on S/390
 - SG24-5687 <http://www.ibm.com/redbooks/abstracts/sg245687.html>
- e-Business Intelligence: Data mart Solutions with DB2 for Linux on zSeries
 - SG24-6294 <http://www.ibm.com/redbooks/abstracts/sg246294.html>
- SAP on DB2 UDB for OS/390 and z/OS: Application Servers on Linux for zSeries
 - SG24-6847 <http://www.ibm.com/redbooks/abstracts/sg246847.html>
 - <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246847.html?Open>
- SAP on DB2 UDB for OS/390 and z/OS: High Availability Solution Using System Automation
 - SG24-6836 <http://www.ibm.com/redbooks/abstracts/sg246836.html>
- Experiences with Oracle for Linux on zSeries
 - SG24-6552 <http://www.ibm.com/redbooks/abstracts/sg246552.html>

Workshop agenda



- Linux for zSeries update
- WebSphere & e-business connectors architecture
- DB2 UDB for Linux on zSeries
- Linux on zSeries for Business Intelligence solutions
- Linux on zSeries for SAP solutions
- Linux on zSeries for Oracle solutions

Linux for zSeries update





Where are we with Linux?



- Packaged and distributed by distributors
 - Red Hat, SuSE, SCO Group, TurboLinux
 - United Linux Consortium
- UNIX-like operating system developed by Linus Torvalds
 - Basic SMP support - version 2.2 in 1999
 - 8x SMP scalability and initial Enterprise features - v2.4 in 2001
 - 32x SMP scalability, more Enterprise features - v2.6 in Q4/2003
- Open Source is proving reliable and a solid foundation
- Customers are challenged understanding patches & Support
 - Increased commitment to maintenance and support
- Enterprise Versions from Red Hat, UL and SuSE
 - High Availability Features
 - Performance Features
 - 24 X 7 support options
 - Optimized release cycles: 12-18 months, 3-5 year Version support
 - ISV Certification Platform
 - Enterprise rich functionality such as multi-threaded I/O, improved SMP utilization, and more ...

Linux distributors combine their power



- Concept:
 - Binary-compatible Linux distribution, branded "UnitedLinux"
 - Linux distributors maintain brand names
 - "UnitedLinux Inside"
 - SuSE acts as UnitedLinux systems integrator
- Customer Value
 - Drives more rapid adoption and robust versions of Linux
 - Consistency in base kernel regardless of distributor
 - Expanded portfolio of certified hardware and ISV applications
- Implementation:
 - Cross-platform support for:
 - All IBM eServer platforms, key middleware
 - Worldwide language support
 - Standards based (e.g. LSB 1.1, LI18NIX)
 - Offering will be optimized for Enterprise class environments, much like SuSE Linux Enterprise Server 7 is today

**UnitedLinux =
Open Industry
Consortium**



IBM will continue to support Red Hat and United Linux across its key hardware, software and services offerings

Linux market acceptance



Linux will become the dominant server operating system in the United States by 2005.

Stacey Quandt, Giga, Business 2.0, June 17, 2002

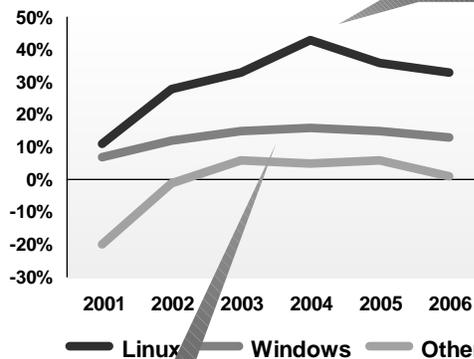
Linux will have a "breakout year" in 2002. Now it seems clear that Linux has become a viable alternative for enterprise use.

IDC, January 2002

By 2006, Linux will be a key foundation for a strategic, cross-development-platform environment, accelerating Unix server consolidation, while creating a powerful alternative to Windows .NET.

Gartner, May 2002

**Server Shipments by OS
% Growth**



**Linux
CAGR
35%**

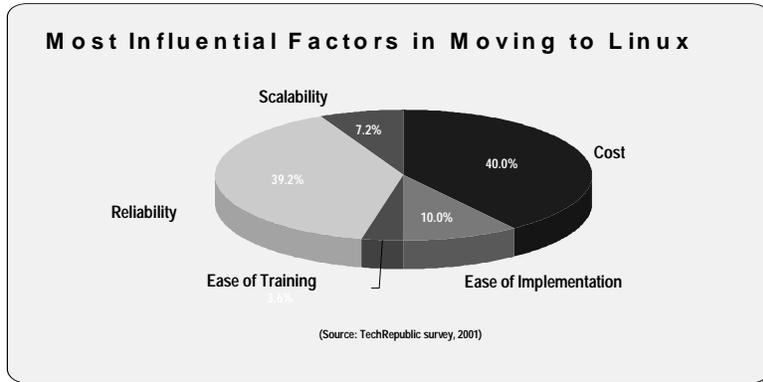
2003 Linux surpasses Unix in server shipments

IDC, Enterprise Server Fundamentals, March 2002

Linux value propositions

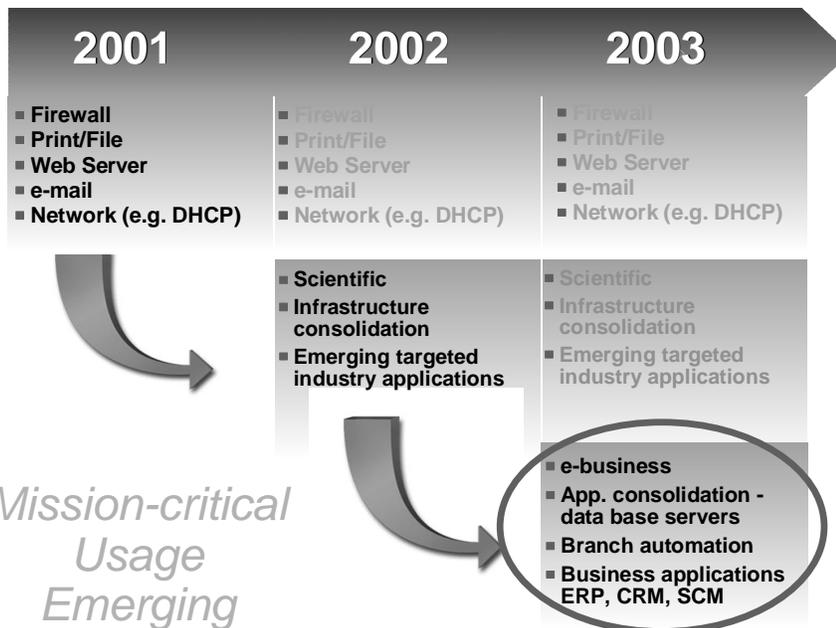


The Linux Marketplace

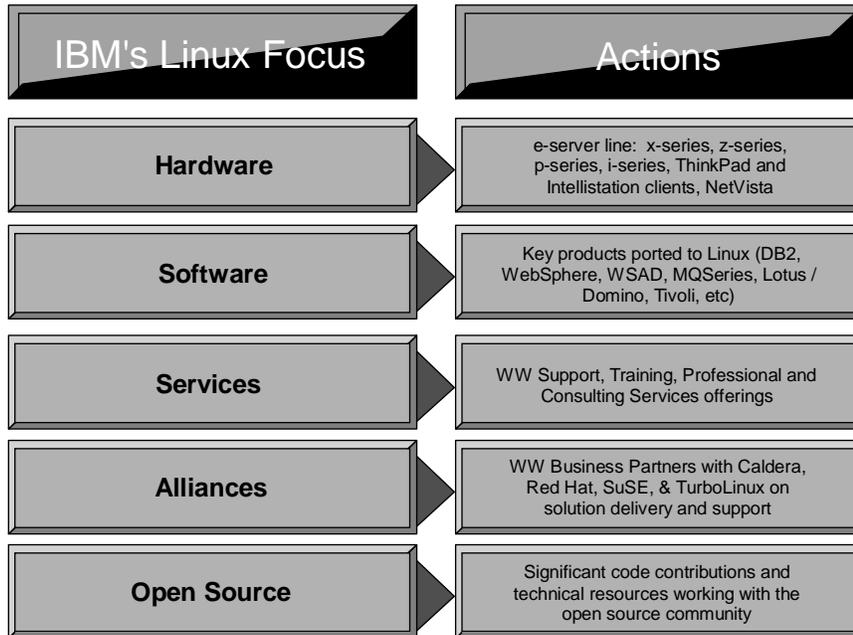


- Flexibility and Choice of Deployment
- Ability to Consolidate on Platform of Choice
- Application Availability & New Business Innovation
- Open Standards
- Price / Performance (Total Cost of Ownership)

Linux adoption



IBM's commitment to Linux



IBM's commitment to Linux ...



A Commitment to Linux across the Entire Business



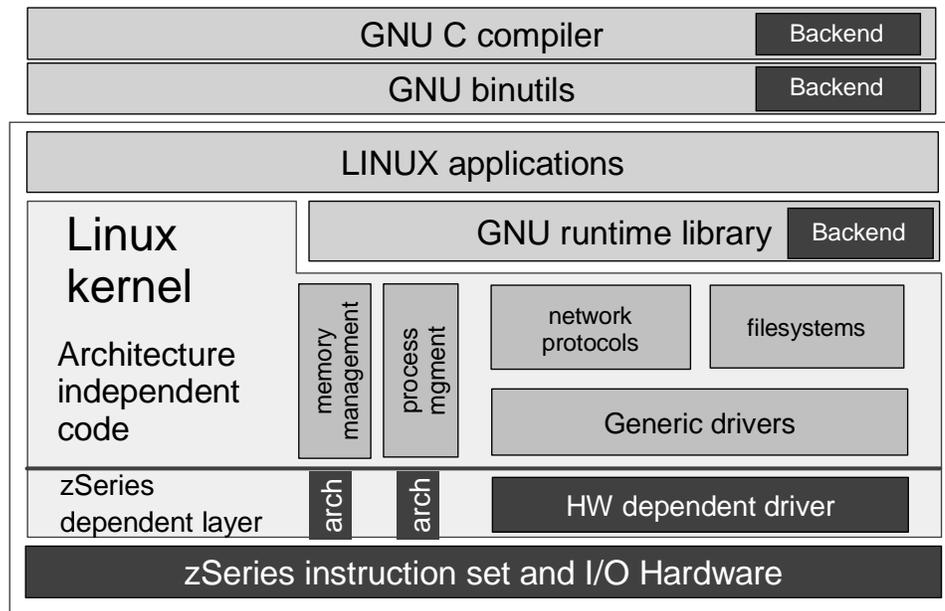
Where are we with Linux for zSeries?



- A native zSeries operating environment
 - Pure Linux, an ASCII environment
 - Exploits IBM zSeries hardware (IEEE FP)
 - Linux for S/390 - 31-bit
 - Linux for zSeries - 64-bit
- Linux run mode:
 - Native, LPAR, z/VM
 - z/VM runs 100's of Virtual Servers
- Processor support
 - Multiprise 3000, 9672 G5, G6, z800, z900 and z990
- Integrated Facility for Linux (IFL)
 - Dedicated to z/VM and Linux workloads
- Not a unique version of Linux or other operating system
- Not a replacement for other IBM zSeries operating systems



Linux for zSeries system structure



IBM contributed

Linux distributions for zSeries and S/390



- RedHat
 - RedHat Linux 7.2 for S/390
 - www.redhat.com/software/linux/s390
 - RedHat Linux 7.1 for zSeries
 - www.redhat.com/software/linux/ibmseries/z.html



- SuSE
 - SuSE Linux Enterprise Server 7 for S/390 and zSeries
 - www.suse.com/us/products/suse_business/sles/sles_s390
 - SuSE Linux Enterprise Server 7 for IBM zSeries
 - www.suse.com/us/products/suse_business/sles/sles_zseries
 - SuSE Linux Enterprise Server 8 for IBM mainframe



- Turbolinux - bought by Software Research Associates - SRA (Aug. 2002)
 - Turbolinux Server 6.5 & Enterprise Server 8 for zSeries and S/390
 - www.turbolinux.com/products/s390/



- Marist College, ThinkBlue Millenium, Debian, Caiman
 - ...
- Build your own

Recent & upcoming distributions



- SuSE Linux Enterprise Server 8
 - Service Pack 2 available in 2Q 2003
 - "Carrier-grade Linux" enhancements
 - Lotus Domino support ("epol")
 - Service Pack 3 scheduled for 4Q2003
 - z990 GA-1 exploitation (06/2003 code drop)

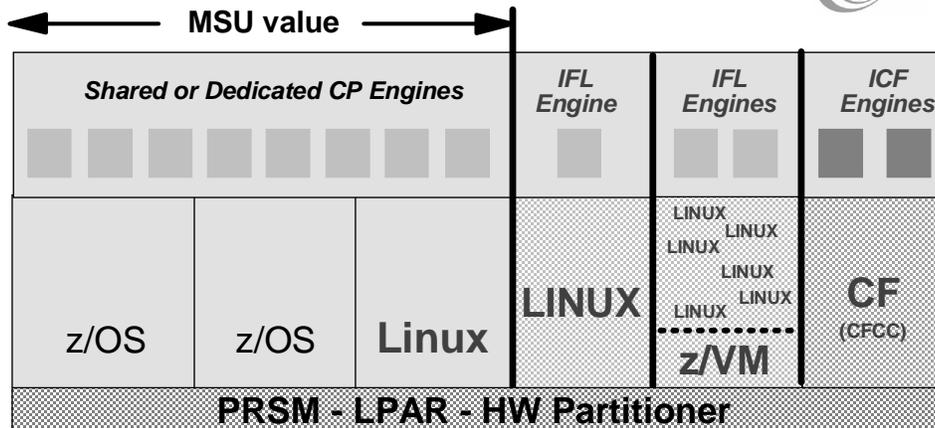


- SuSE Linux Enterprise Server 9
 - "Powered by UnitedLinux 2.0"
 - GA scheduled for 2Q 2004
 - Linux kernel 2.6 (tbc)
 - Full z990 GA1/GA2 exploitation (12/2003 code drop)



- RedHat Enterprise Linux AS 3.0
 - GA scheduled for 4Q2003
 - Linux kernel 2.4.21
 - New POSIX Thread Library (NPTL)

Typical configurations on zSeries



Logical partitions (LPAR)

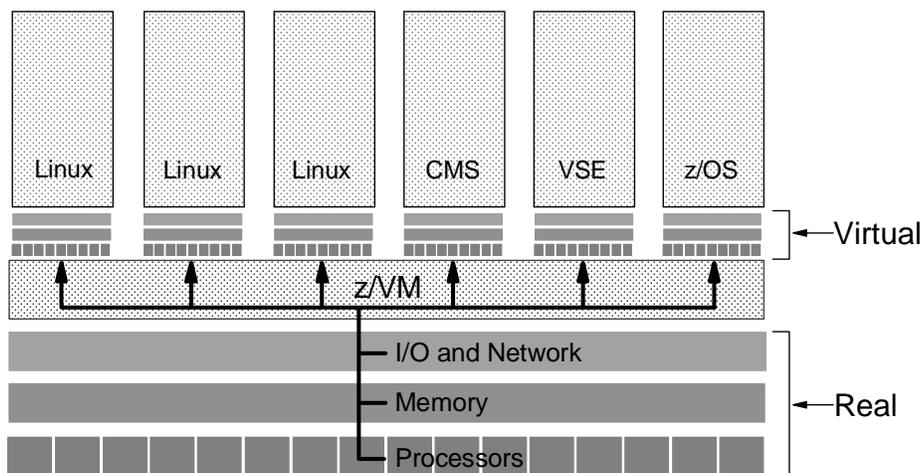
IFL engines

- Engines dedicated to Linux - can **not** run z/OS or any other zSeries operating system
- IFL and CPs cannot be mixed in an LPAR
- IFL capacity does NOT affect the SW cost of IBM & ISV products under z/OS

z/VM virtual machine partitioning

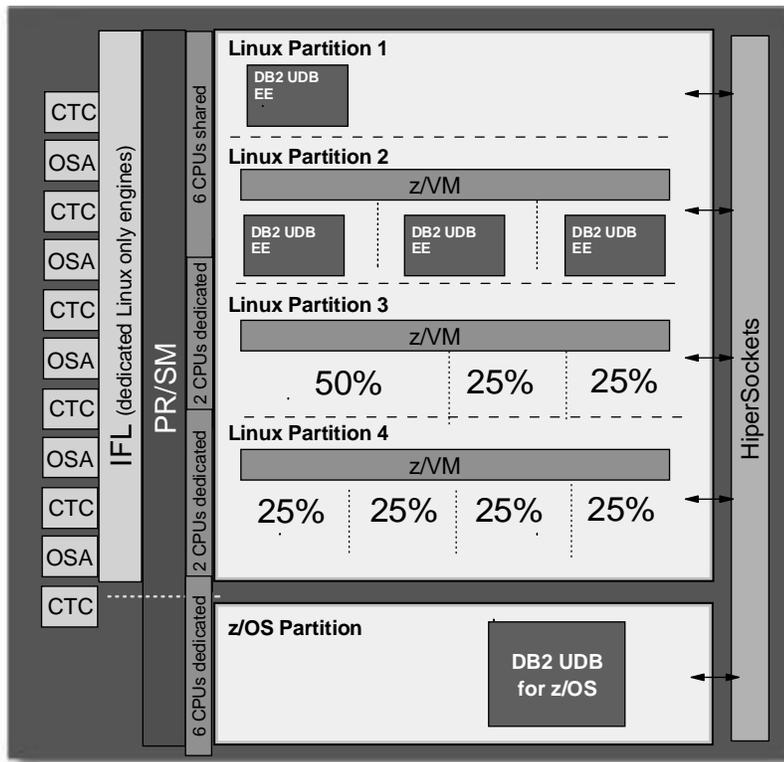


Efficiently Exploiting the Entire Mainframe Complex



A *Virtual Machine* simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.

zSeries partitioning example

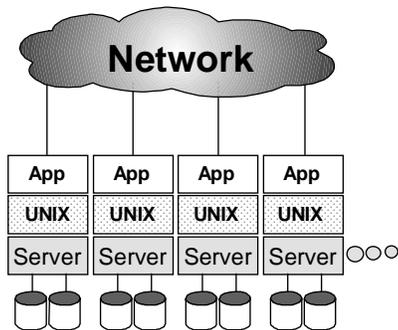


The value of z/VM for Linux



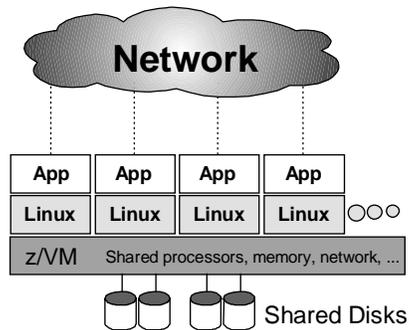
Rapid Horizontal Growth with Linux on z/VM

Typical UNIX environment



- Dedicated processors and disks
- Complex system management
- New servers available in days

Linux on z/VM



- Shared resources
- Simplified system management
- New servers online in minutes

"Just add another server" quickly and easily with z/VM

IBM z/VM statement of direction

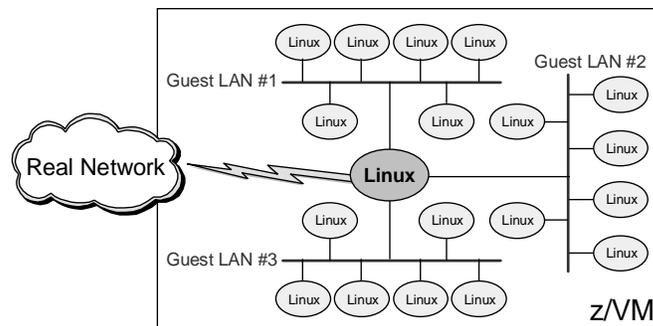


- IBM intends to support greater than 16 processors in a single image of z/VM on future releases. This support is planned to increase the scale of a single z/VM image by taking advantage of the increased number of processors available on a z990 server.
- Future releases of z/VM will impose an Architectural Level Set (ALS) requiring z/Architecture. Consequently, future releases of z/VM are planned to operate only on IBM zSeries 800, 900, and 990, equivalent servers, and future IBM eServer zSeries servers.
- IBM intends for z/VM V4.4 or later to provide support for up to 60 logical partitions (LPARs) with corresponding support on a z990 or future server.
- IBM intends for z/VM V4.4 or later to support four logical channel subsystems with a maximum of 256 channels supported per image with corresponding support on a z990 or future server.
- IBM intends to provide z/VM guest support for the PCIX Cryptographic Coprocessor (PCIXCC).
- It is planned that future performance management enhancements will be made primarily to the Performance Toolkit for VM. z/VM V4.4 is planned to be the last release in which the RTM and PRF features will be available. Both the RTM and PRF features are planned to be withdrawn from marketing in a future z/VM release.

z/VM guest LANs



- A Guest LAN is a "virtual" LAN created by the z/VM Control Program
- OSA-Express (QDIO) and HiperSockets Guest LANs can be created
 - Point-to-point, Multicast, and Broadcast (QDIO) connections are supported
- Linux images can connect to one or more Guest LANs
 - And connect to real network adapters at the same time
 - This enables a Linux image to provide external routing and firewall services for other Linux images



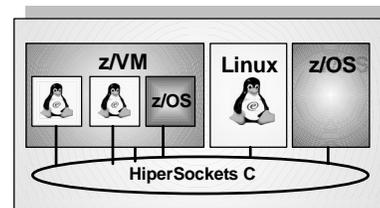
HiperSockets - the network in the box



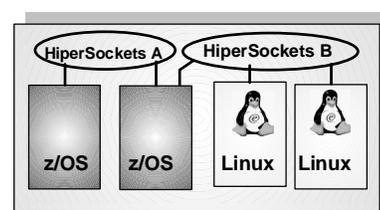
IP networking among virtual servers in a zSeries

- Improved response time due to low latency
 - High Speed connectivity via memory bus
- Highly Secure
 - Data never flows outside the server
- Highly available
 - Integrated zSeries hardware, no external parts
- Cost savings
 - No external network, attachment, or cables
- Flexible
 - Combinations of z/OS, Linux, and z/VM
 - From 4 to 16 HiperSockets
 - From 1024 to 4096 TCP/IP stacks
- Simple to install, operate, maintain
 - Transparent to applications

z990 / z900 / z800



z990 / z900 / z800



Pre-requisites

- z/OS 1.2, Linux kernel 2.4, z/VM 4.2
- z990, z900 or z800 Microcode level

How does Linux benefit from zSeries?



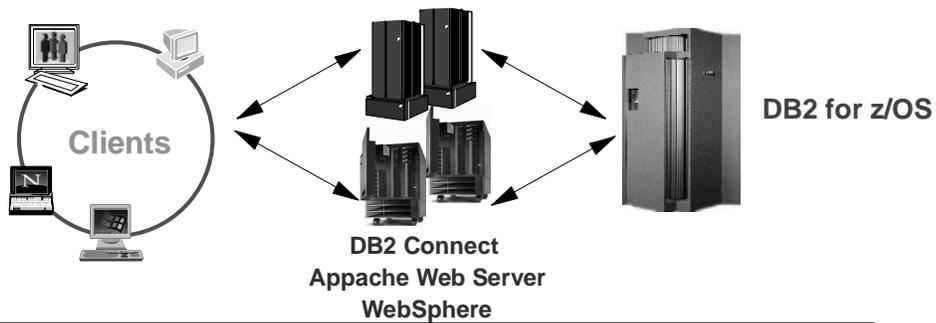
- Reliability
 - The most reliable hardware platform available
 - Mean time between failure over 60 years
- Manageability
 - Centralized Linux systems for zSeries are easier to manage and duplicate
- Scalability
 - Physical - very large processor configurations
 - Logical - hundreds of Linux images on a single footprint
 - Nondisruptive scalability Capacity Upgrade on Demand
- Multiple workloads support
 - Multiple isolated workloads on a single server
 - Efficient use of all server resources
- High speed internal data transfer
 - HiperSockets



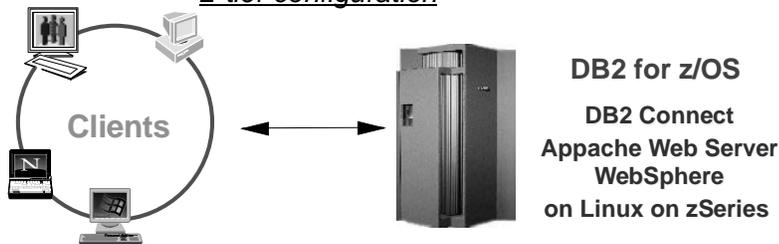
Elimination of middle-tier servers



3-tier configuration



2-tier configuration



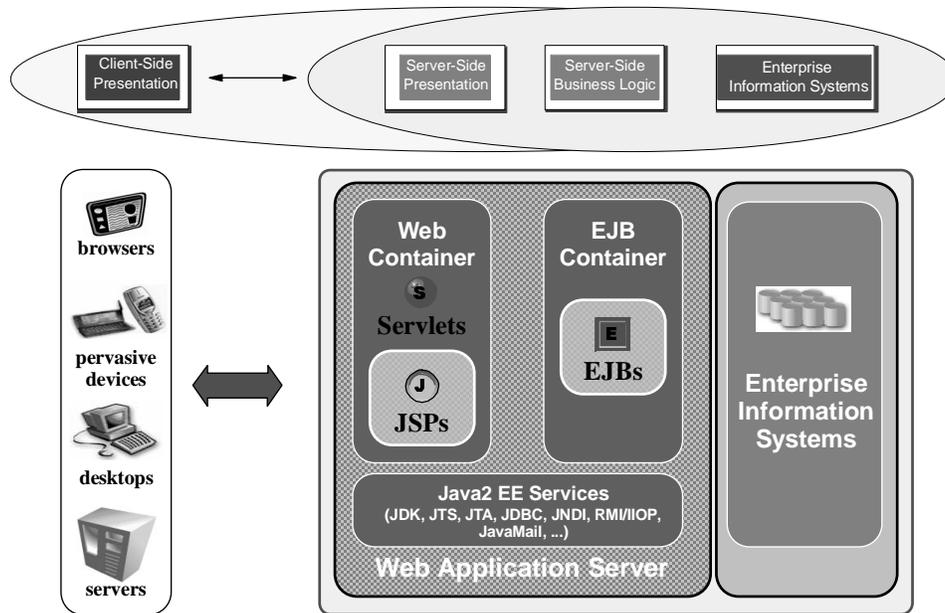
WebSphere and e-business Connectors Architectures



WebSphere architecture



Systems are best designed using a multi-tier application architecture
Separate the logic for Presentation, Business logic and Data access



WebSphere Application Server themes

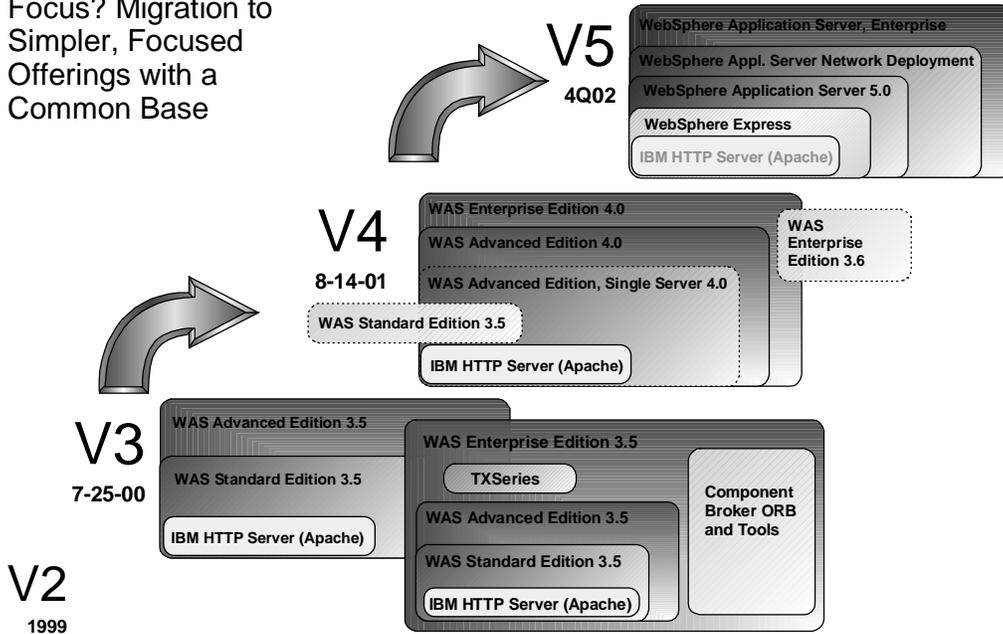


- J2EE 1.3 Compliance
 - EJB 2.0, Servlet 2.3, JSP 1.2
 - Integrated, built-in JMS provider
 - J2EE 1.3 Security: Java 2 Security, JAAS, Enhanced Pluggable Authentication
- Extended WebServices Support
 - Enhanced SOAP support, Private UDDI, and WebServices Gateway
- New, more open, flexible administration model
 - Based on Java Management Extensions (JMX)
 - Provides improved failover capability and high availability
 - User interface enhancements and application management
- Edge Components Integration
- Autonomic Computing
- Extensions: beyond the standard programming model

WebSphere Application Server evolution



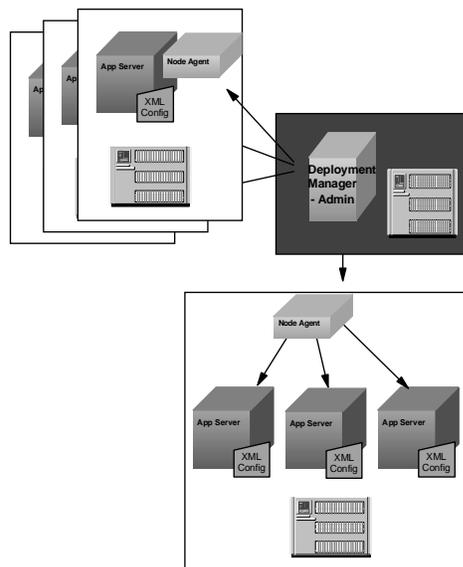
Focus? Migration to
Simpler, Focused
Offerings with a
Common Base



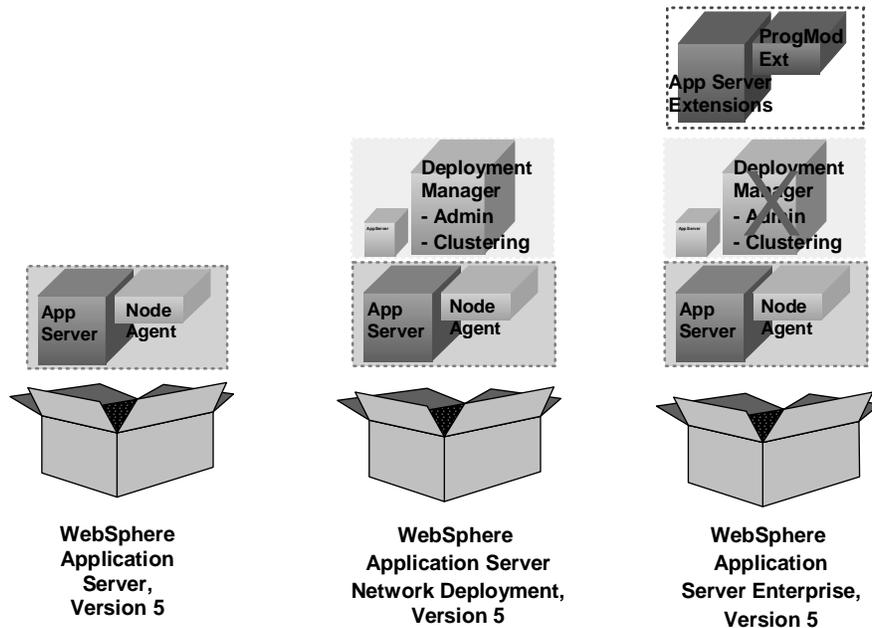
Product packaging: new terminology



- *Managed Process or Server*
 - Each server running in its own JVM
 - Application Servers
 - JMS Server
- *Node Agent*
 - Resides on a single node
 - Manages the servers running on the node
- *Deployment Manager*
 - Manages the multiple nodes in a distributed topology
- *Cell*
 - Network of multiple nodes in a single logical administration domain



Product packaging: 5.0 release



WebSphere Application Server 5.0

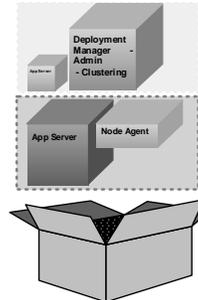


- Security
 - Java 2, JAAS, CSiv2
 - Support for crypto card Eracom CSA 8000
- Problem determination and tracing
 - FFDC First failure data capture
 - End-to-End Transaction Trace
 - Messaging improvements
- Application Assembly Tool
 - Supports J2EE 1.3
- Enhanced migration tooling
 - Also, 3.5/4.x interoperability/coexistence
- IBM HTTP Server
 - IHS 1.3.26
 - IHS 2.0
- JDK 1.3.1
- Single server environment
- Full J2EE 1.3 support
- Web Services Update
 - SOAP 2.3, WSIF
 - UDDI clients
- Extended relational DB support
 - Add Oracle 9i
 - Ships Cloudscape for samples
- New Administrative Model
 - Based on Java Management Extensions (JMX) framework
 - Browser-based Administration Console
 - XML-based configuration repository
 - No RDB is required
 - Role based administration security

WAS Network Deployment 5.0



- Targets multiple servers, multiple nodes environment
 - Focus is on scalability, availability, and performance
 - Clustering and failover support
 - Web server plugin supports static weighted load distribution
- Distributed Administration
 - Browser based administration console for the entire topology
 - XML-based repository
- Ships with some Edge Components on "DMZ CD"
 - Network Dispatcher
 - Switch Consultants for CISCO and Nortel switches
 - Site Selector for scaling Network Dispatcher
 - Caching Proxy
 - Reverse Proxy, Content Based Routing,...
 - Dynamic content caching: integration with WAS servlet caching
- Enhanced Web Services functions
 - UDDI Registry - for private UDDI registry
 - Web Services Gateway

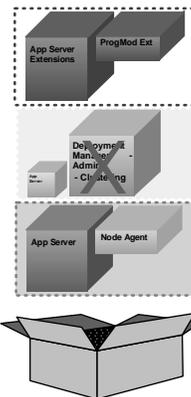


**WebSphere
Application
Server Network
Deployment
(WAS-ND)**

WAS Enterprise 5.0



- Focus is on extending the J2EE Programming Model
 - To enable implementation of sophisticated application functions
 - To provide performance advantages
- Functions provided in WAS Enterprise 5.0
 - Dynamic Access Intent and Application Profiling
 - As opposed to static Access Intent
 - Extended Unit of Work
 - Activity Sessions, Last Participant Support
 - Dynamic Query Capabilities
 - Extends static and standard EJB QL
 - Container Managed Messaging
 - WebSphere Workflow
 - Synchronous / Asynchronous runtime support
 - Asynchronous Beans
 - Scheduler
 - Startup Beans



**WebSphere
Application
Server
Enterprise
(WAS Enterprise)**

WAS standard checklist



| Standards | Level | WebSphere 5.0 |
|-----------------|-------|--|
| J2EE | 1.3 | ✓ Fully certified and part of Sun's JCEE list |
| EJB | 2.0 | ✓ EJB 2.0 and EJB 1.1 support |
| JDK | 1.3 | ✓ JDK 1.3 |
| Servlet | 2.3 | ✓ Servlet 2.3 |
| JSP | 1.2 | ✓ JSP 1.2 |
| JTS/JTA | 1.0.1 | ✓ w/distributed transactions |
| JMS | 1.0.2 | ✓ With Native Provider, and MQ plug-in |
| JDBC | 2.0 | ✓ 2PC across heterogeneous databases |
| JNDI | 1.2 | ✓ JNDI 1.2 for EJB lookup and CosNaming |
| RMI/IIOP | 1.3.1 | ✓ Fully supported |
| JavaMail/JAF | 1.2 | ✓ Plus Domino support |
| SSL Security | 2.0 | ✓ JSSE 1.0.2 and JCE 1.2.1 |
| XML JAXP | 1.1 | ✓ XML in EJBs |
| J-IDL/CORBA | | ✓ IIOP 1.2 |
| J2C | 1.0 | ✓ Bean and container managed |
| LDAP | | ✓ Directory Server, SecureWay, iPlanet, ActiveDirectory |
| HTTP | 1.1 | ✓ Yes, plus across multiple Web servers |
| SOAP | 2.3 | ✓ Soap support for WebServices. |
| SOAP-SEC | 1.0 | |
| COM/ASP Support | 1.0 | ✓ w/Java wrapping & proxy |
| JMX | 1.0 | ✓ |
| XML4J | 4.2 | ✓ XML support: includes JAXP1.1, XMLSchema1.0, Xerces1.4.1 |
| XSL | 2.3.6 | ✓ XSL parser |

<http://www-4.ibm.com/software/webservers/appserv/doc/latest/prereq.html>

WebSphere family for Linux on zSeries



| | | |
|--|---|-------------------|
| WebSphere AS V5 (Express) | ✗ | Not in plan |
| WebSphere AS V5 (Base) | ✓ | |
| WebSphere AS V5 (Network Deployment) | ✓ | |
| WebSphere AS V5 (Enterprise) | ✓ | |
| WebSphere Studio | ✗ | Not in plan |
| WebSphere Portal for Multiplatform Enable V5.0 | ✓ | |
| WebSphere Portal for Multiplatform Extend V5.0 | ✓ | |
| WebSphere Commerce Business Edition V5.5 | ✓ | |
| WebSphere Commerce Professional Edition V5.5 | ✓ | |
| WebSphere MQ V5.3 | ✓ | Client and Server |
| WebSphere Host on Demand 8.0 | ✓ | |

IBM JDK support matrix



| Distribution | Kernel | Mode | SDK & JRE 1.3.1 31-bit | SDK & JRE 1.3.1 64-bit | SDK & JRE 1.4.1 31-bit | SDK & JRE 1.4.1 64-bit |
|-----------------------|--------|--------|------------------------|------------------------|------------------------|------------------------|
| SuSE 7.0 | 2.2.16 | 31-bit | ok | - | - | - |
| SuSE SLES7 (1) | 2.4.7 | 31-bit | ok | - | ok | - |
| SuSE SLES7 (1) | 2.4.17 | 64-bit | ok | - | - | - |
| SuSE SLES8 (2) | 2.4.19 | 31-bit | ok | - | ok | - |
| SuSE SLES8 (2) | 2.4.19 | 64-bit | ok | ok | - | ok |
| United Linux 1.0 | 2.4.19 | 31-bit | ok | - | ok | - |
| United Linux 1.0 | 2.4.19 | 64-bit | ok | ok | - | ok |
| Turbolinux Server 6.5 | 2.2.16 | 31-bit | ok | - | - | - |
| RedHat 7.1 | 2.4.9 | 64-bit | - | - | - | - |
| RedHat 7.2 | 2.4.9 | 31-bit | ok | - | - | - |

- (1) SuSE SLES7 Patch CD1 needed
- (2) SuSE SLES8 Service Pack 2 needed

IBM WAS V4 and V5 support matrix



No WAS 64-bit version available

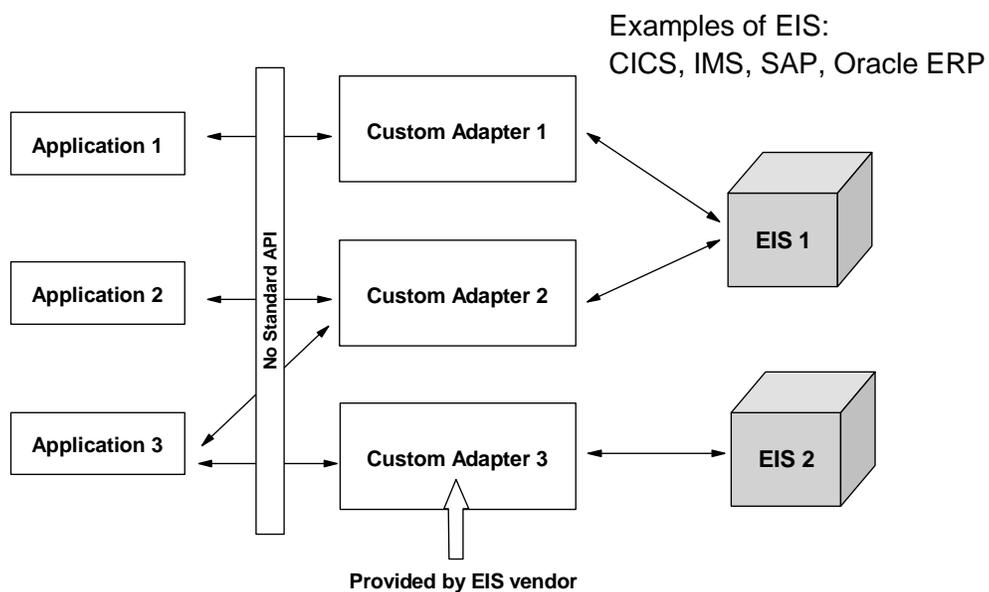
| Distribution | Kernel | Mode | WAS AE 4.0.6 31-bit | WAS 5.0.2 31-bit Enterprise |
|-----------------------|--------|--------|---------------------|-----------------------------|
| SuSE 7.0 | 2.2.16 | 31-bit | ok | - |
| SuSE SLES7 (1) | 2.4.7 | 31-bit | ok | ok |
| SuSE SLES7 (1) | 2.4.17 | 64-bit | - | - |
| SuSE SLES8 (2) | 2.4.19 | 31-bit | - | ok |
| SuSE SLES8 (2) | 2.4.19 | 64-bit | - | - |
| United Linux 1.0 | 2.4.19 | 31-bit | - | ok |
| United Linux 1.0 | 2.4.19 | 64-bit | - | - |
| Turbolinux Server 6.5 | 2.2.16 | 31-bit | - | - |
| RedHat 7.1 | 2.4.9 | 64-bit | - | - |
| RedHat 7.2 | 2.4.9 | 31-bit | - | ok |

- (1) SuSE SLES7 Patch CD1 needed
- (2) SuSE SLES8 Service Pack 2 needed

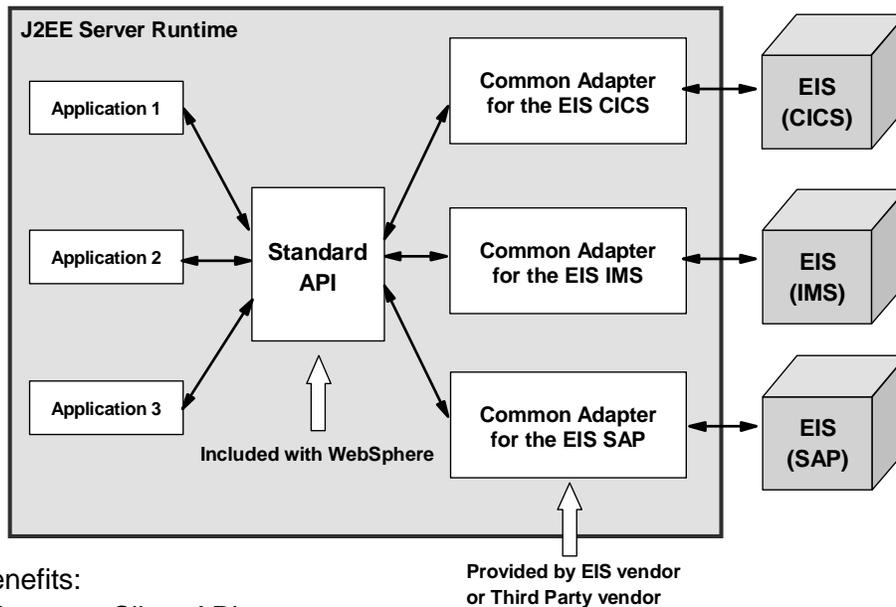


J2C Connectors Framework

Connecting to EIS - past



Framework architecture



Benefits:

- Common Client API
- Common EIS Adapters

Provided by EIS vendor
or Third Party vendor

Java2 connectors

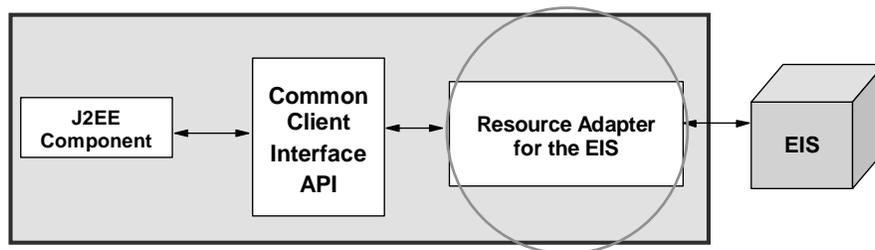


- J2C 1.0 is part of J2EE 1.3 spec
 - Implemented in WebSphere as technology preview
- A standard architecture for connecting J2EE components to heterogeneous EIS
- Benefits:
 - Simplifies the integration of diverse EIS
 - Each EIS requires just one implementation of the J2EE Connector Technology.
- J2C Components:
 - Resource Adapter (RA)
 - Supplied by EIS or third party vendor
 - Common Client Interface (CCI)
 - Packaged with WebSphere

J2C: Resource Adapter (RA)



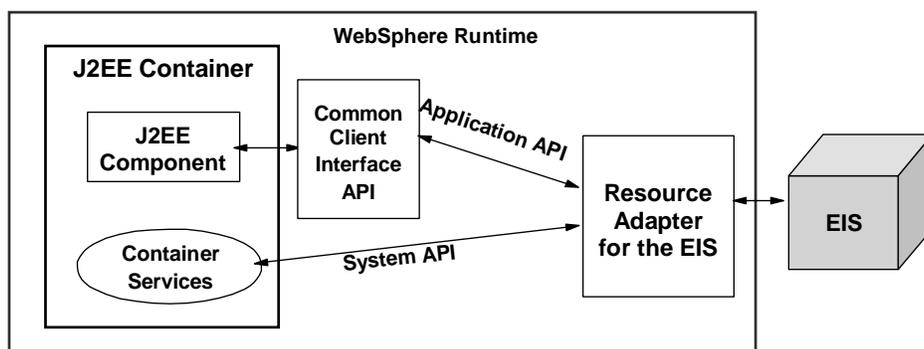
- Provides connectivity between J2EE Component and EIS
- Plugs into an application server
- Collaborates with the Application server to provide important services
 - Connection, transactions, security, etc.
- Stored in a .rar (Resource Adapter archive) file



J2C: Resource Adapter Details



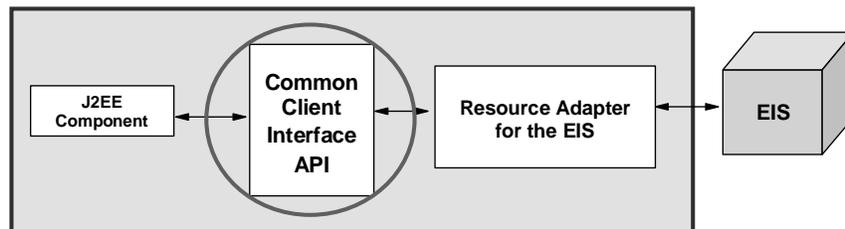
- Two types of API implemented by a Resource Adapter
 - Application API
 - Defines the API through which a J2EE component accesses the EIS.
 - System API links the Resource Adapter to J2EE server services
 - Connection, transaction, security, etc.



J2C: Common Client Interface (CCI)



- CCI API provides access from J2EE components to an EIS through a Resource Adapter
 - Simplifies writing code to connect to an EIS's data store
 - J2EE components interact with CCI only - not directly with Resource Adapter



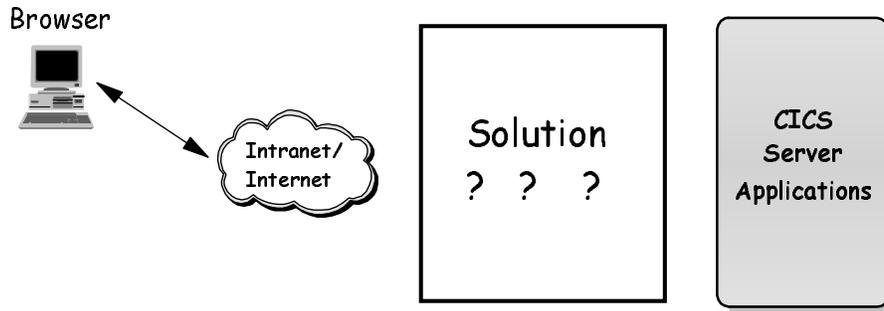
CICS Transaction Gateway for Linux on zSeries - a J2C Perspective



Business needs



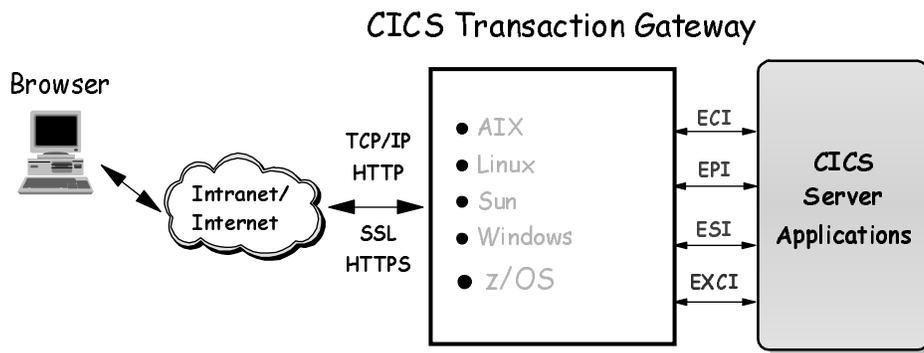
- End users need secure and easy access from a browser to business-critical applications running on a CICS Server



Solution - CICS Transaction Gateway



- The business-critical applications running on a CICS Server can be accessed securely and easily from a browser using Java via commarea



CTG Overview



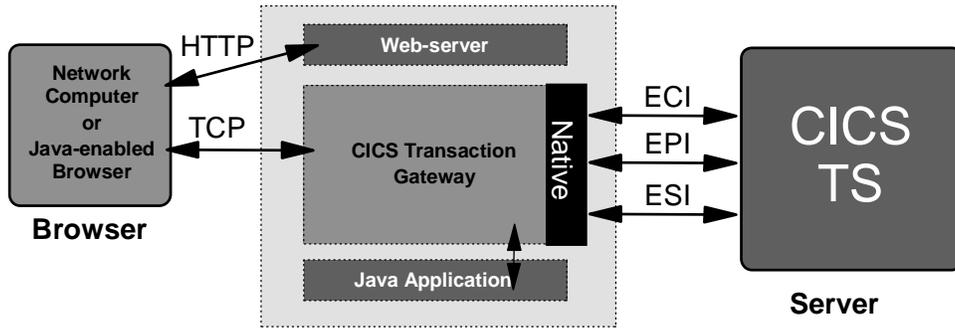
- The CICS Transaction Gateway ...
 - Provides an interface between the world of Java & HTML and the proven capabilities of CICS, so..
 - ... allows State-of-the-art access to CICS applications from a machine which simply has a Java Virtual Machine or Web Browser installed
- Can be used from ...
 - Java Applets
 - Java Applications
 - Java Servlets

CICS Transaction Gateway



- "The purpose of CICS Transaction Gateway is to provide efficient integration of middle tier application servers with CICS ... by providing a multi-user gateway which supports different programming interfaces... for use by Web applications in Java and other languages."
- "CICS Transaction Gateway enables customers to implement robust and efficient end-to-end configurations and... provides considerable flexibility in the configurations which may be adopted."
- "The option to deploy CICS Transaction Gateway on OS/390 or z/OS, supporting Web applications locally on OS/390 or z/OS, or remotely on a physical middle tier server... enables high scalability to be achieved and has been tested at transaction rates in excess of 1,000 transactions per second."

CICS Transaction Gateway

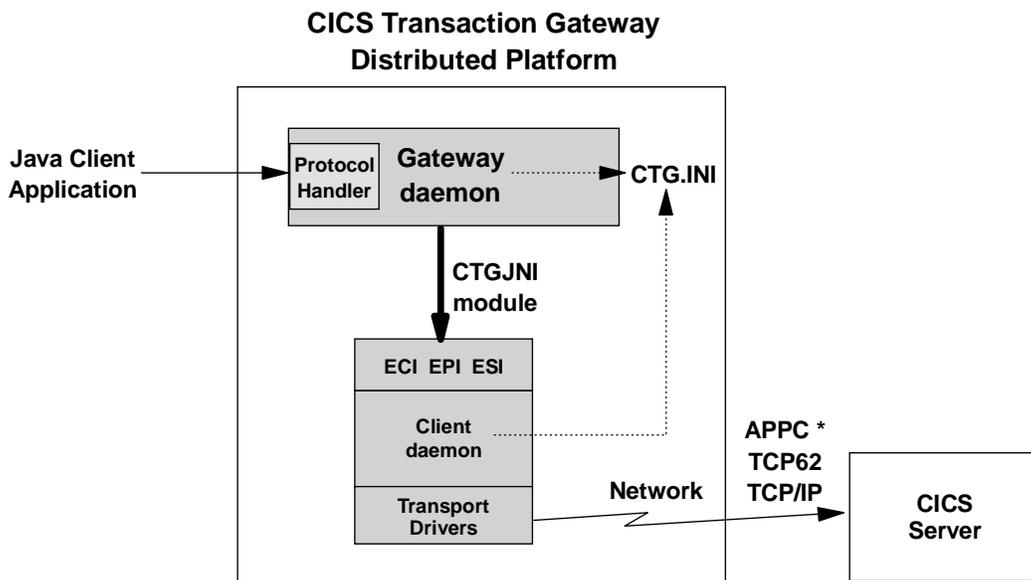


Current Version
for all platforms
is **CTG V5**

CICS Transaction Gateway is:

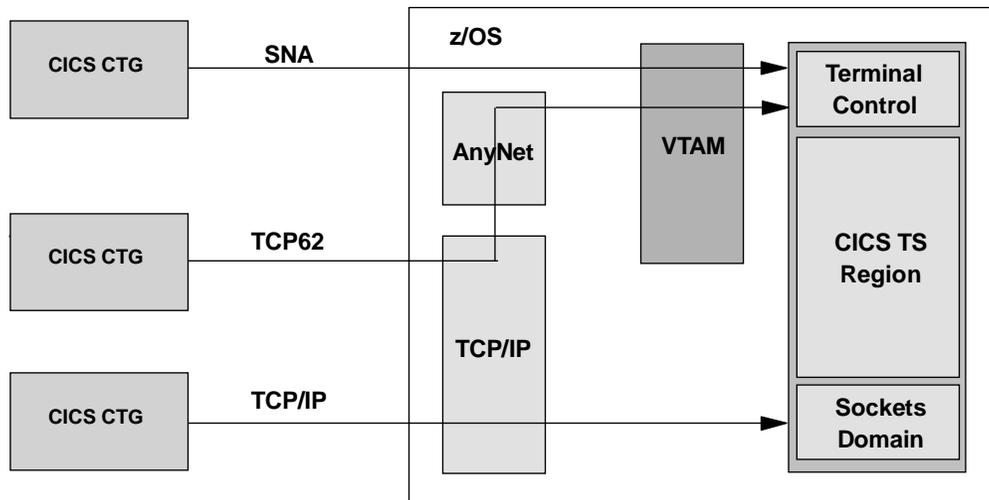
- ▶ Gateway daemon
- ▶ A set of Java classes which can be used to connect to CICS (from applet, application or servlet)
- ▶ Native code which handles communication from Gateway to CICS server

CICS Transaction Gateway

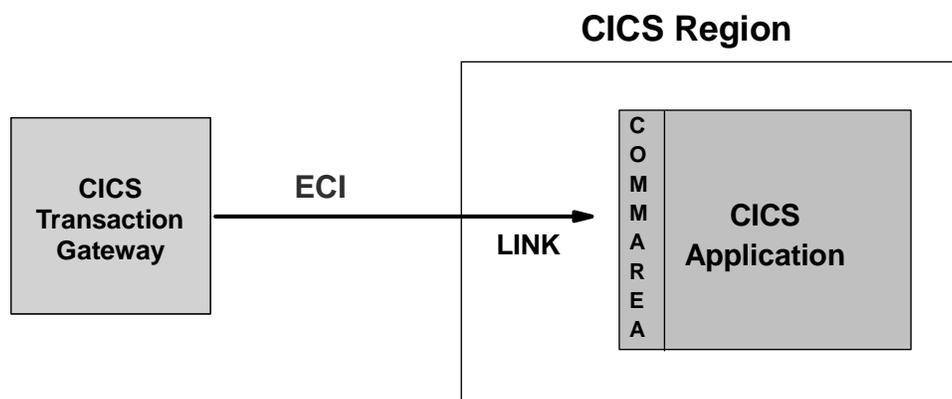


* APPC not supported on Linux for zSeries environment

Network communications with CICS

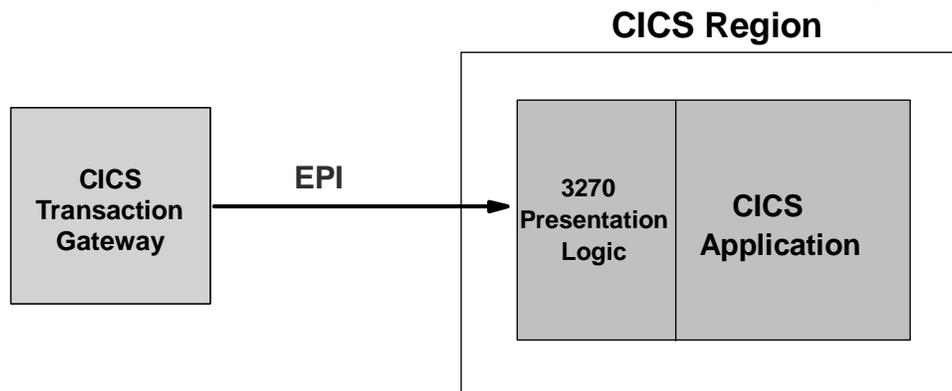


CTG interfaces - ECI



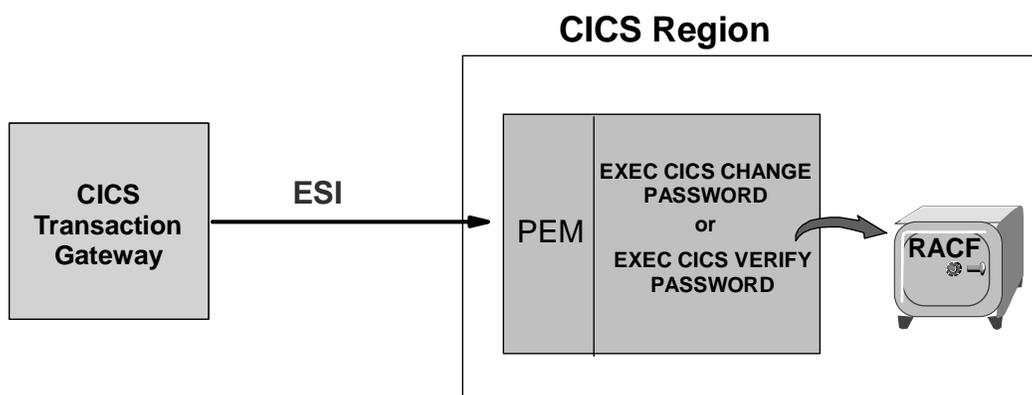
- The CTG has 3 external interfaces
- The External Call Interface (ECI)
 - ▶ Link to CICS programs
 - ▶ Available on all platforms

CTG interfaces - EPI



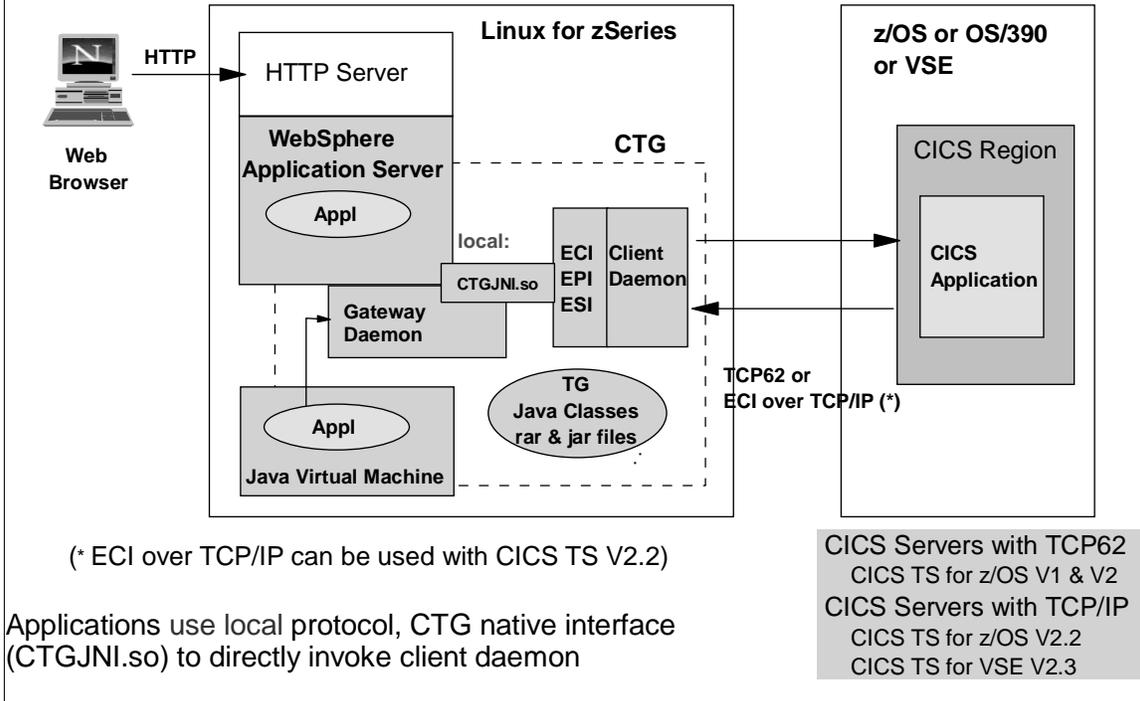
- The External Presentation Interface (EPI)
 - ▶ Invoke 3270 based transactions
 - ▶ Available on distributed platforms only
 - ▶ EPI support classes and EPI beans provided
 - ▶ Do not need detailed knowledge of 3270 data streams to use

CTG interfaces - ESI

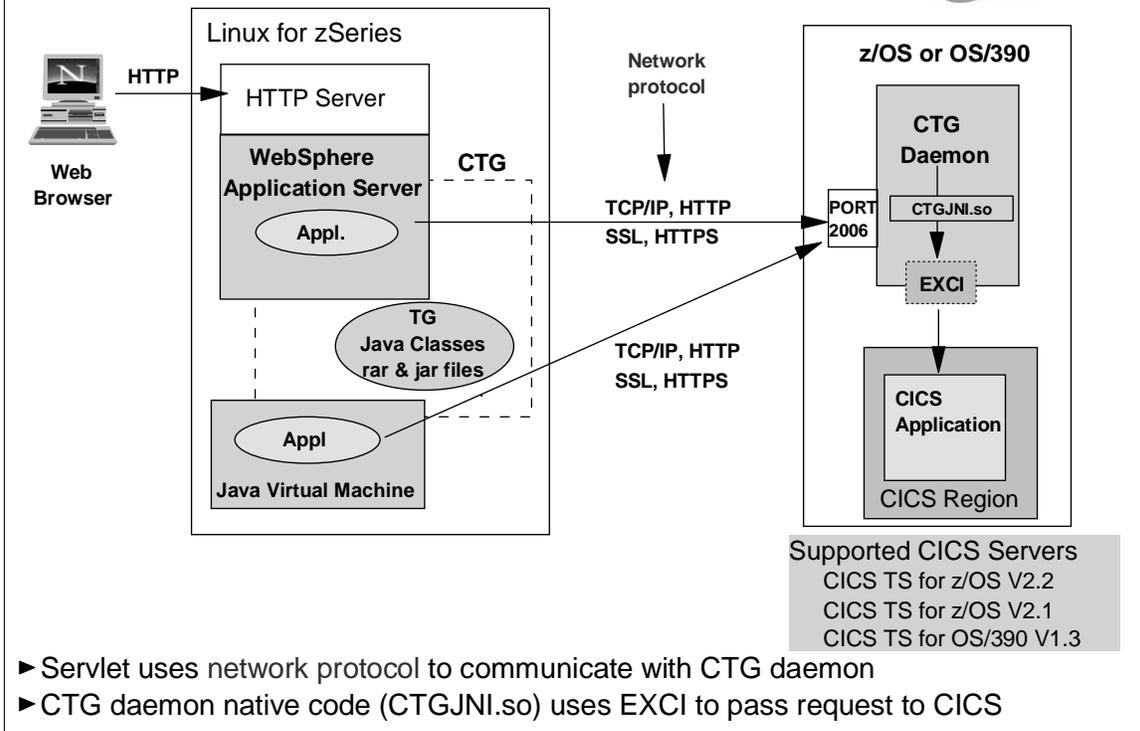


- The External Security Interface (ESI)
 - ▶ Verify and change password information
 - ▶ Available on distributed platforms only

CTG with WAS on Linux on zSeries



Using the CTG daemon on z/OS



CICS Transaction Gateway V5



- Generally available July 2002
- Enhancements include:
 - Support for the J2EE Connector Architecture
 - First introduced in CTG V4.01 (dist. platforms) and CTG V4.02 (for use with WebSphere for z/OS)
 - Support for J2EE Asynchronous ECI calls on z/OS
 - Operational improvements e.g.
 - Dynamic control of tracing
 - ARM restart for CTG/390
 - JSSE support (Java Secure Sockets Extension) for 128-bit encryption
 - Enhanced logging
 - No default logging of client connections
 - Logging of EXCI return codes on z/OS
 - Enhanced TCP62 support
 - Including support for TCP62 on the Linux for S/390 platform

IBM CTG V5.0.1 support matrix



| Distribution | Kernel | Mode | IBM CTG 5.0.1 31-bit |
|-----------------------|--------|--------|-------------------------|
| SuSE 7.0 | 2.2.16 | 31-bit | ok |
| SuSE SLES7 (1) | 2.4.7 | 31-bit | ok |
| SuSE SLES7 (1) | 2.4.17 | 64-bit | - |
| SuSE SLES8 (2) | 2.4.19 | 31-bit | ok |
| SuSE SLES8 (2) | 2.4.19 | 64-bit | - |
| United Linux 1.0 | 2.4.19 | 31-bit | ok |
| United Linux 1.0 | 2.4.19 | 64-bit | - |
| Turbolinux Server 6.5 | 2.2.16 | 31-bit | ok |
| RedHat 7.1 | 2.4.9 | 64-bit | - |
| RedHat 7.2 | 2.4.9 | 31-bit | - |

(1) SuSE SLES7 Patch CD1 needed

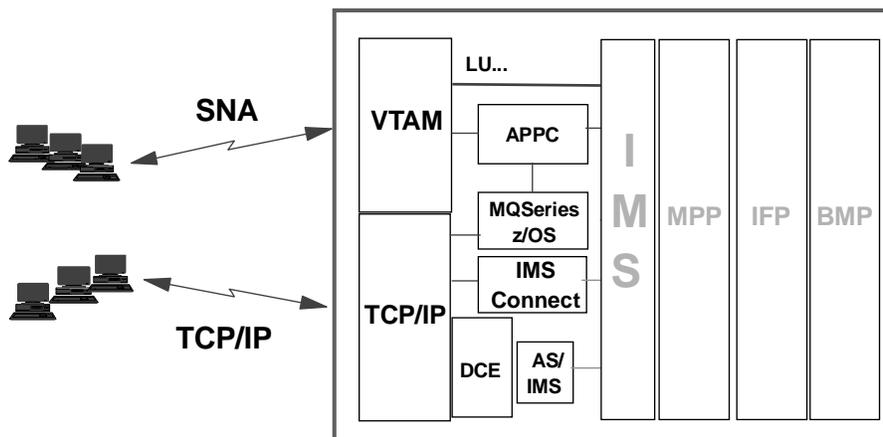
(2) SuSE SLES8 Service Pack 2 needed



IMS Connect



Access to IMS Java application programs

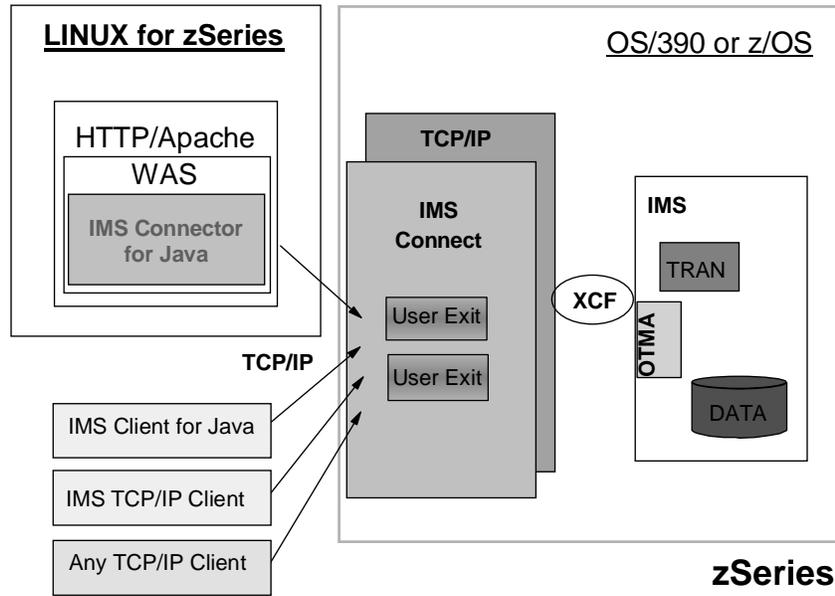


MPP : Message Processing Program
BMP : Batch Processing Program
IFP : IMS Fastpath Program

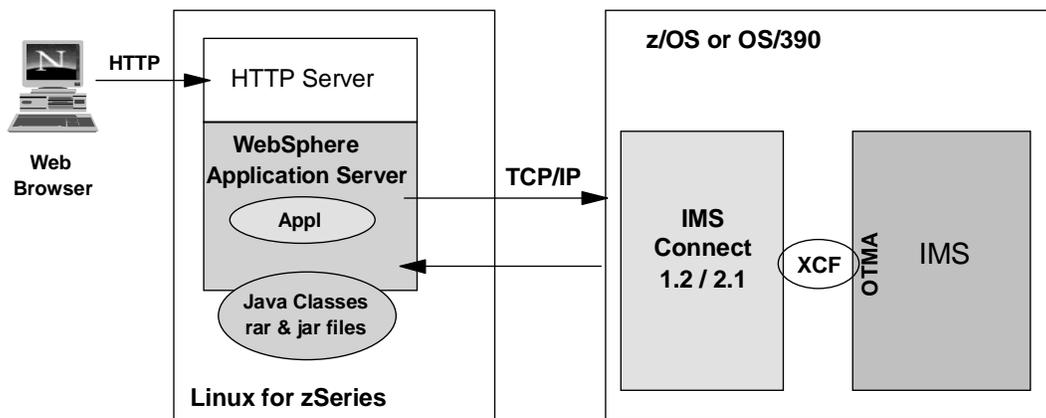
IMS TCP/IP connection possibilities



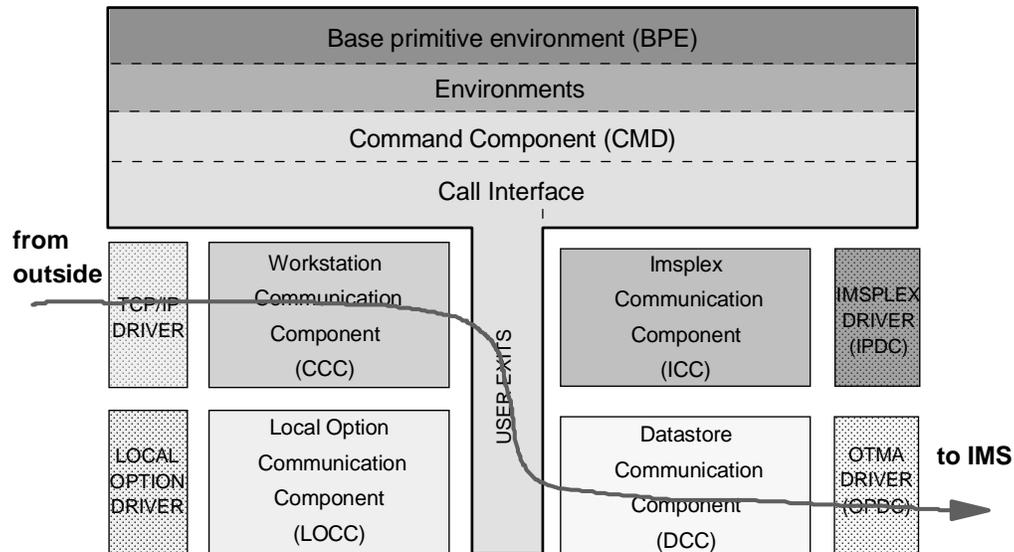
OTMA : Open Transaction Manager Access



IMS connector on Linux for zSeries



IMS Connect 2.1 gateway



IMS Connect for z/OS V2.1 enhancements



- **IMS Connect for z/OS V2.1 major enhancements:**
 - Two-phase commit support across TCP/IP for distributed environments to expand transactional integrity (Linux zSeries)
 - Security enhancements
 - Auto reconnect to IMS
 - Internet Protocol Version 6 (IP V6) network protocol support for scalability
 - WAS for z/OS V5 support

IMS Connect for z/OS requirements



- IMS Connect for z/OS V1.2 requires:
 - IMS V6 or V7 Transaction Manager (TM)
 - Connector for Java J2EE runtime support requires IMS V7 TM
 - IMS Connector for Java WebSphere V5 requires WAS V5
- IMS Connect for z/OS V2.1 requires:
 - IMS V7 or IMS V8

IMS documentations



- SG24-6514-00 IMS e-business Connectors
- SC27-1559-00 IMS Connector for Java User's Guide and Reference
- SC27-0946-01 IMS Connect Guide and Reference
- SG24-6123-00 IMS Version 7 and Java Application Programming
- SG24-6536-00 IMS Version 7 Java Update
- J2EE Connector Architecture and Enterprise Application Integration



Using JDBC on Linux for zSeries

Introduction

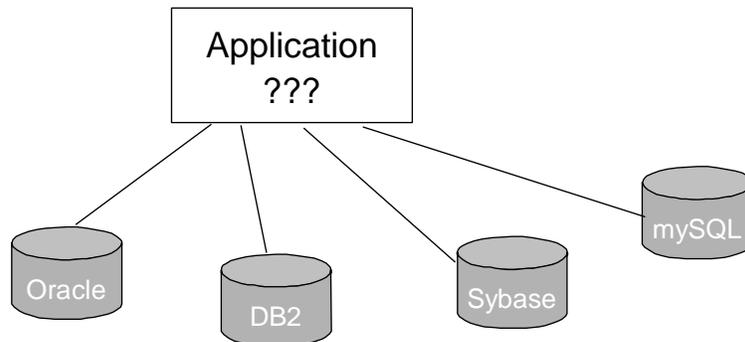


- JDBC and SQLJ are standard Application Programming Interfaces for creating platform independent Java applications that can access any relational database
- JDBC and SQLJ are Open standards
- Sun Microsystems developed the specifications
- Major DBMS vendors are now supporting
 - IBM, Oracle, Sybase, Informix ...
- JDBC provides support for dynamic SQL, it uses a Java call interface to pass SQL statements to DB2
- SQLJ provides support for static SQL and use embedded SQL statements in Java applications and servlets.
- In JDBC Applications SQL are built and prepared at runtime prior to execution .
- SQLJ applications require precompiles and binds at build time.

Compatibility of database applications



Back to old days of applications development

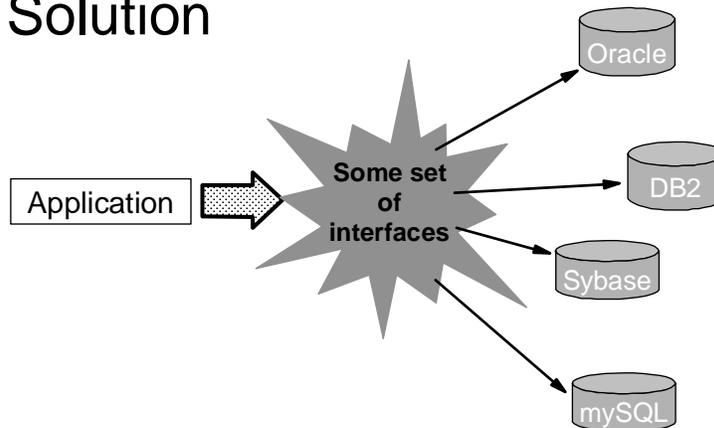


Migration to other database required huge amount of change in the application code

Standardized set of interfaces required...



Solution



First came ODBC...



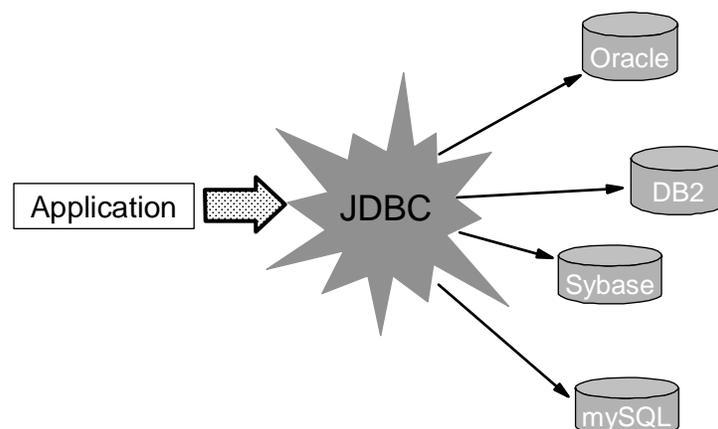
- ODBC – Open Database Connectivity, specifies a standard set of interfaces for database access (developed by Microsoft™)
- Limitations
 - Only for windows platform
 - C Language
 - Complex design
- Something better still needed...

Then came JDBC for Java



A little thought on Java (write once, compile once, run everywhere)

And so enters JDBC...

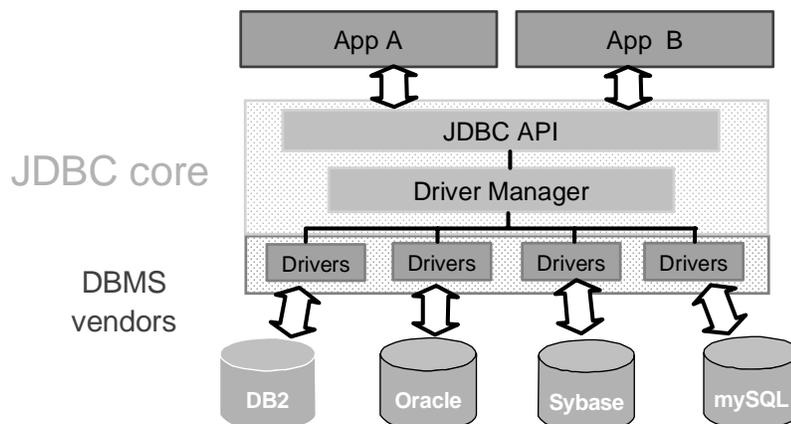


What is JDBC?



- JDBC is a standardized interface for accessing databases from the the Java programming Langage.
- JDBC™ is short for Java Database connectivity
- Sun Microsystems developed the specifications : JavaSoft JDBC Specification - 1.2
- Many compagnies work together to define new and updated versions of the specifications: Three versions have been defined this far called : JDBC1, JDBC2 and JDBC3.
- IBM DB2 Universal Database V8.1 supports the JDBC v2.1 specification and portions of JDBC 3.0.

JDBC architecture



JDBC drivers

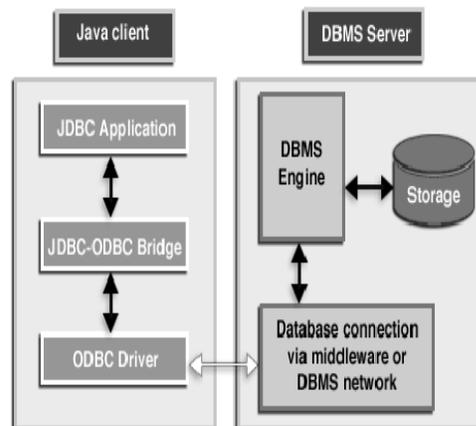


- The JDBC interface offers the possibility to access a variety of database without recoding the application by means of drivers
- Drivers are provided by the DBMS vendors
- Java defines 4 types of JDBC drivers
 - **Type 1** -- Implements JDBC as a layer on top of ODBC.
 - **Type 2** -- Uses native method calls (JNI) to the database engine native API's (DLLs/SO's, specially written for JDBC).
 - **Type 3** -- Communicates with a middleware server using a database independent protocol.
 - **Type 4** -- Is a network driver that communicates with the database using the native network protocol built into the database engine.
- Type 1 drivers generally have more CPU overhead than type 2
- Type 3 and type 4 drivers must route through a network layer, not generally efficient for local JDBC connectivity.

JDBC driver Type 1



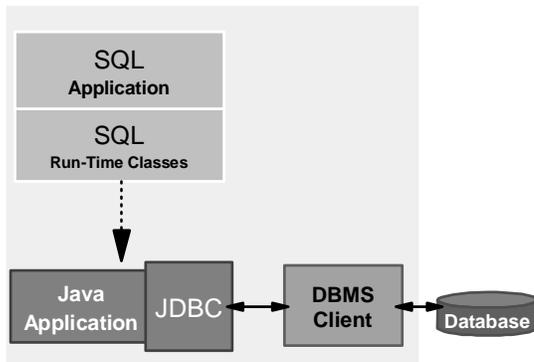
- A JDBC Type 1 driver is based on a JDBC-ODBC bridge.
- An ODBC driver is used in combination with this JDBC driver (provided by Sun - supplied as part of standard java packages).
- IBM does not provide a Type 1 driver.
- It is a "deprecated" and not recommended environment.



JDBC driver Type 2



- A JDBC Type 2 driver relies on a specific shared library to communicate with the DBMS.
- The application loads the JDBC driver and the driver uses the shared library to communicate with the server.



- DB2 UDB for Linux, UNIX®, and Windows® V8.1 provides two different type 2 drivers:
 - The legacy/CLI JDBC driver in the file db2java.zip.
 - Used for the current J2EE certifications.
 - The Universal JDBC driver in the file db2jcc.jar.
 - New to DB2 UDB V8.1 for Linux, UNIX, and Windows
 - Initial implementation was for direct Java connectivity to servers, using a type 4 driver architecture.
 - Can use this driver with DB2 V8.1.2 in a type 2 architecture for local application performance and for distributed transaction support.

JDBC driver Type 3

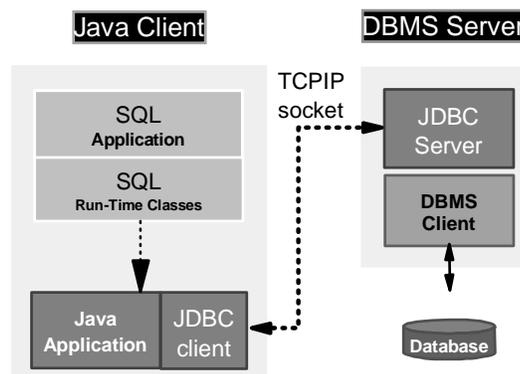


- JDBC Type 3 driver is a pure Java implementation that must communicate with a JDBC Applet Server to access database data.
- Often referred to as the "net driver", since the application talks to another machine where a DBMS client has been installed.

IBM DB2 UDB Version 8.1 supports the net driver. It can be used within JDBC applications.

The type 3 driver (db2java.zip) must be at the same maintenance level as the Applet Server. Except when the driver is used within applets, the browser will download the appropriate file during application execution.

Currently, the WebSphere Application Server and other J2EE servers do not support the IBM type 3 driver, because that driver does not support distributed transactions (JTA) (2 phase commit)



JDBC driver Type 4

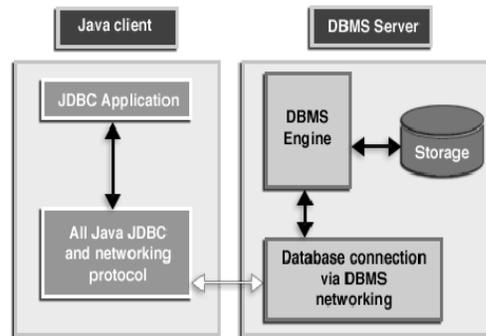


- A type 4 driver is a Java-only JDBC driver that connects directly to the database server.
- DB2 UDB for Linux, UNIX, and Windows v8.1 introduces a type 4 driver known as the "Universal JDBC driver". The Universal JDBC driver is provided in the file db2jcc.jar.

The DB2 UDB Universal JDBC driver is an architecture-neutral JDBC driver for distributed and local DB2 UDB access. It is independent of any particular JDBC driver-type connectivity or target platform. It supports both all-Java connectivity (type 4 driver) or JNI-based connectivity (type 2 driver) in a single driver instance to DB2 UDB.

This driver can be used for standalone Java applications or multi-tier applications.

Requires a license JAR file to be in the CLASSPATH along with the db2jcc.jar

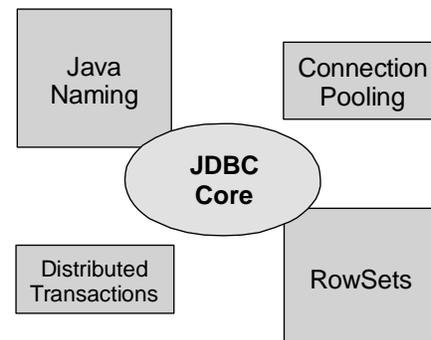


J2EE - JDBC Standard extensions



With J2EE java environment JDBC standard extensions add more fonctionnality to the JDBC core.

Enterprise applications exploite the JDBC standards extensions to access database, inititate distributed transacions..., by means of the Application servers (ie IBM WebSphere WAS 5)



Data Sources



- Is the alternative to Driver and DriverManager class
 - Problem was: the connection URL and the driver class need to be known at the application code level.
 - Based on Naming and Directory service (JNDI).
- Properties defined for DataSource
 - *dataSourceName*
 - *serverName, portNumber*
 - *databaseName, user, password*

Data Sources (continued)



- "*DataSource*" interface in JDBC represents a data source (database location, connection pool..)
- Its attributes are assigned when the DataSource instance is bound to its JNDI directory
- APIs
 - *get___()* and *set___()* methods for properties
 - *getConnection()*

Transaction



- ...One atomic unit of work...
- A Transaction is a series of one or more statements followed by either a COMMIT or a ROLLBACK, on the Connection
- When commit() or rollback() is invoked, the current transaction ends and a new one begins on that Connection
- Default mode is auto-commit:
 - commit() automatically called after each statement
- Disabled auto-commit
 - requires an explicit call to either commit() or rollback() to complete a transaction
- All statements of a transaction is committed or rolled back as a group

Connection Pooling



- Up until now
 - Open connection >> talk to database >> close connection
 - Problem : creating a connection is an expensive operation
- Solution – Connection Pooling
 - Open database connections are held in cache: pool of connections
 - get one from pool >> use >> release
- DataSource used
 - *ConnectionPoolDataSource*
- Configurable properties
 - *initialPoolSize, minPoolSize, maxPoolSize, maxIdleTime*

Programming with Java to access to DB2

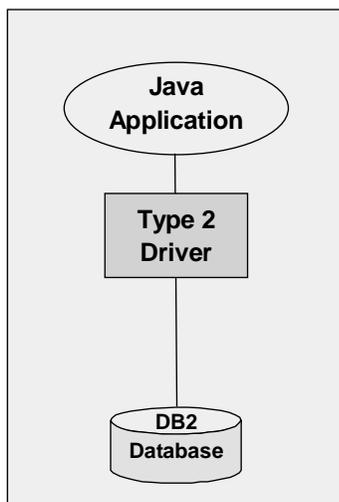


- DB2 has rich support for the Java programming environment. You can access DB2 data by putting the Java class into a module in one of the following ways:
 - DB2 Server
 - Stored procedures (JDBC or SQLJ)
 - SQL functions or User-Defined Functions (JDBC or SQLJ)
 - Browser
 - Applets based on JDBC (JDBC)
 - J2EE Application Servers (such as WebSphere® Application Server)
 - Java ServerPages (JSPs) : JDBC
 - Servlets : SQLJ or JDBC
 - Enterprise JavaBeans (EJBs) : SQLJ or JDBC

DB2 UDB Java connectivity



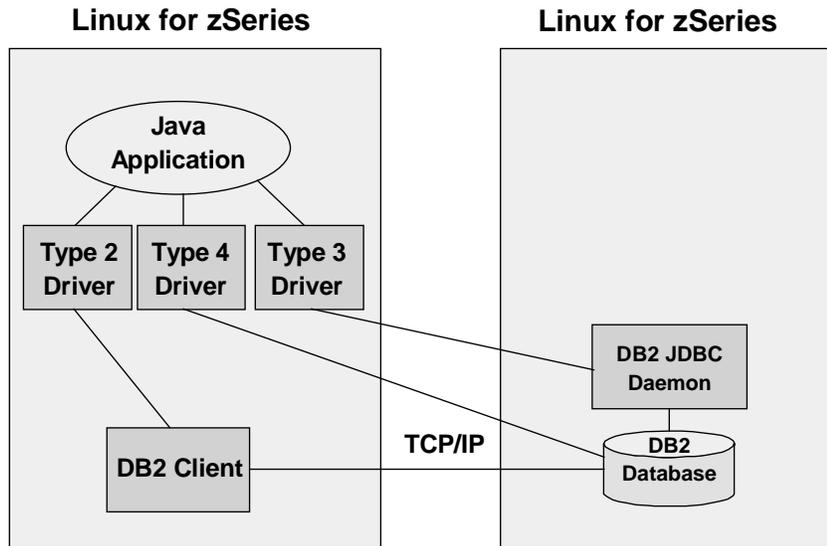
Linux for zSeries



DB2 driver provided as part of DB2 UDB product

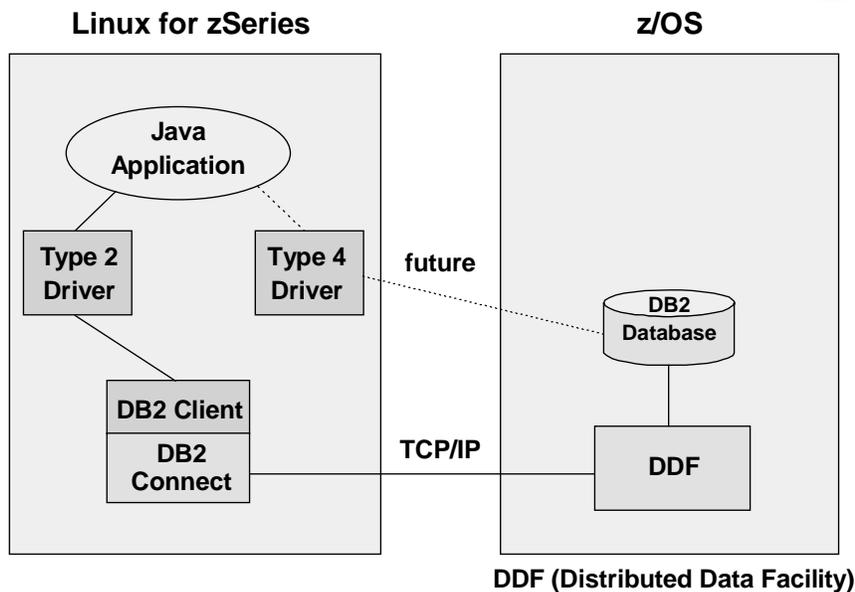
Local connection to DB2 UDB on Linux for zSeries

DB2 UDB Java connectivity



Remote connection to DB2 UDB on Linux for zSeries

DB2 UDB Java connectivity



DDF (Distributed Data Facility)

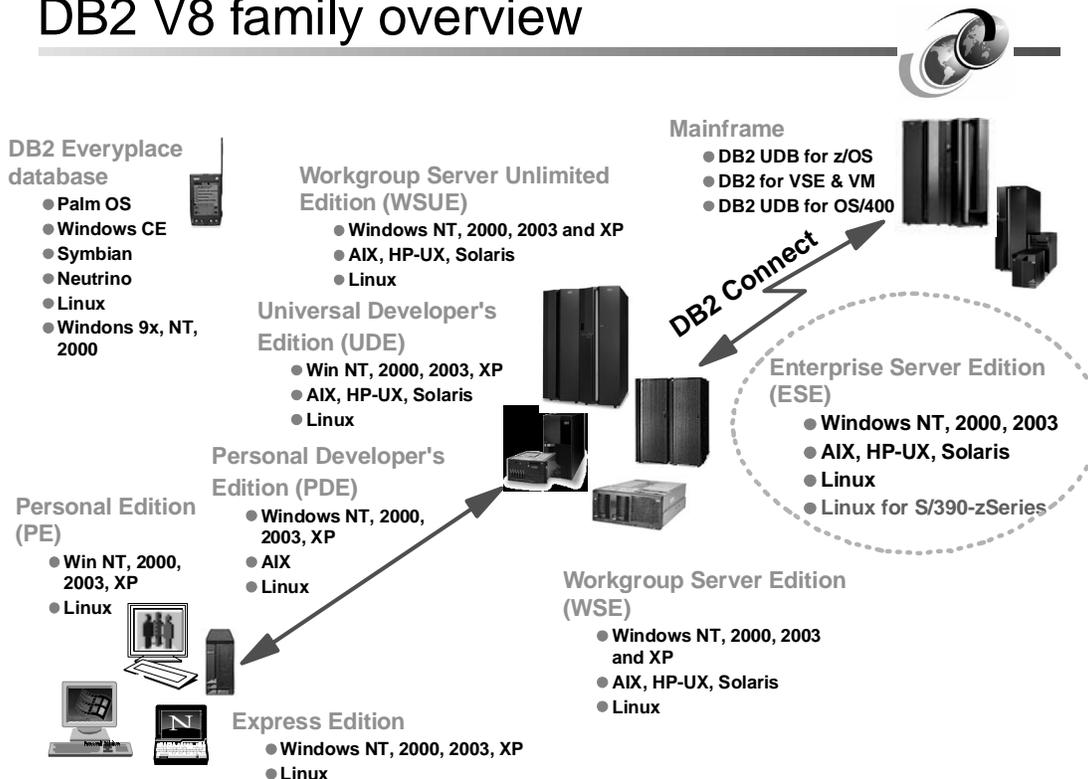
DB2 UDB V8 Type 4 driver supports DRDA protocol

Remote connection to DB2 UDB on z/OS

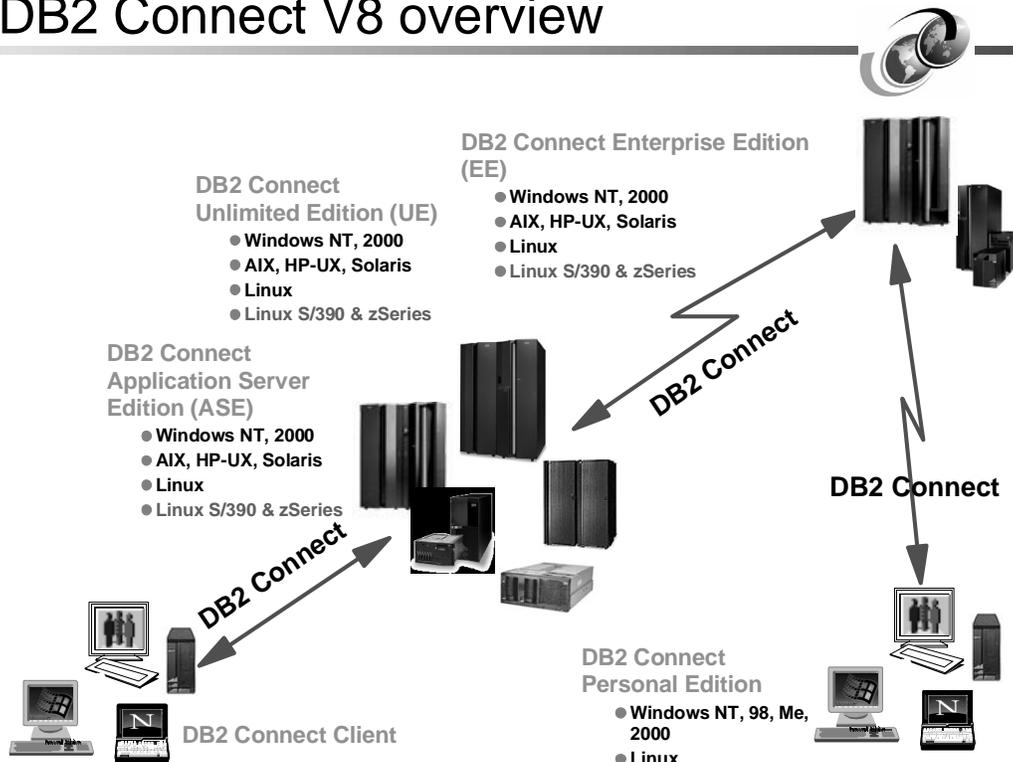
DB2 UDB for Linux on zSeries



DB2 V8 family overview



DB2 Connect V8 overview



DB2 UDB ESE V8 components

- DB2 UDB Enterprise Server Edition for Linux for S/390 and zSeries includes :
 - DB2 Administration Client
 - DB2 Connect Application Server Edition
 - DB2 Connect Enterprise Edition
 - DB2 Connect Unlimited Edition
 - DB2 HTML Documentation
 - DB2 PDF Documentation
 - DB2 Run-Time Client
 - DB2 Spatial Extender
 - DB2 UDB Enterprise Server Edition

DB2 UDB V8 ESE packages



Basic features:

- DB2 Enterprise Server Edition *
- Administration Client *
- Run Time Client *
- IBM Developer kit Java
- Distributed Debugger for Java
Stored Procedures
- Audio, Image, and Video Extenders
- WebSphere Studio Site Developer
Advanced
- WebSphere MQ
- QMF for Windows
- Data Management Tools
- DB2 Connect ASE, EE, UE *

Options Available:

- Database Partitioning Feature *
- DB2 Net Search Extender *
- DB2 Spatial Extender *
- DB2 Data Links Manager
- DB2 Warehouse Manager
- DB2 Intelligent Miner Scoring *
- DB2 Intelligent Miner Modeling *
- DB2 Intelligent Miner
Visualization *

* Available for Linux S/390 & zSeries

DB2 for Linux - Why and what?



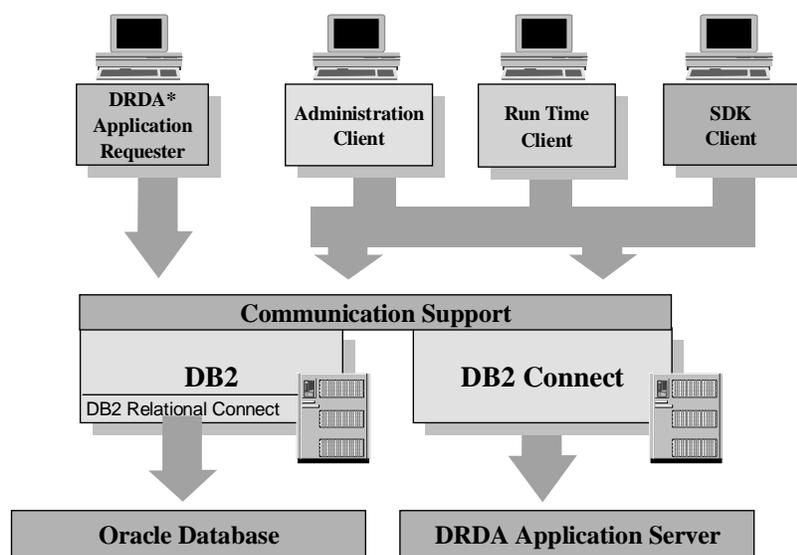
- DB2 UDB for Linux is a relational database management system ready for any e-business and anywhere.
- Is fully Web-enabled, scalable from single processors to symmetric multiprocessors.
- Is one of the databases that can be used by WebSphere, and the first commercial DB software available for Linux on the mainframe.
- You can access, manage and analyze all forms of information across the enterprise and on all major platforms..
- DBM provides Java support, as well as SQLJ, JDBC, Perl, Net.Data to access data from DB2 databases.
- It supports the Open Standards i.e. Java, XML, JDBC2.0 and Web Services.

J2EE coverage



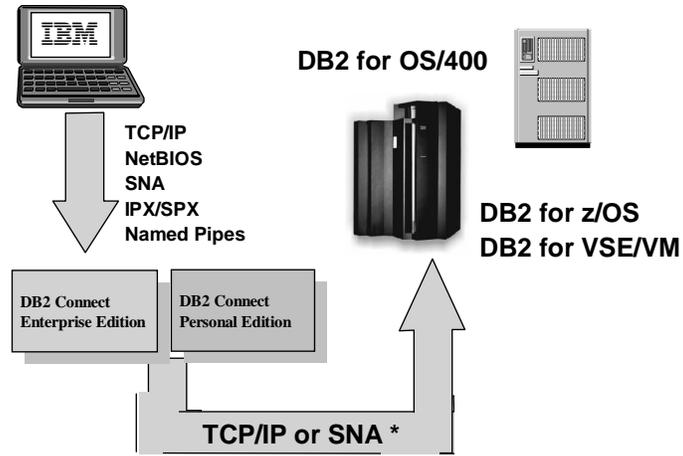
- J2EE 1.3 (spec was finalized in 09/2001) is the de facto industry standard for building java application includes new technology for deploying Web services, which allow application work together better over the Internet.
- IBM WebSphere V5 and DB2 UDB V8.1 have been certified compatible with the latest version of J2EE (Java 2 Enterprise Edition 1.3).
- It means these products support Web Services development, Java Messaging Services, increased integration with support for the CORBA (Common Object Request Broker Architecture) and Java applications.

DB2 UDB components



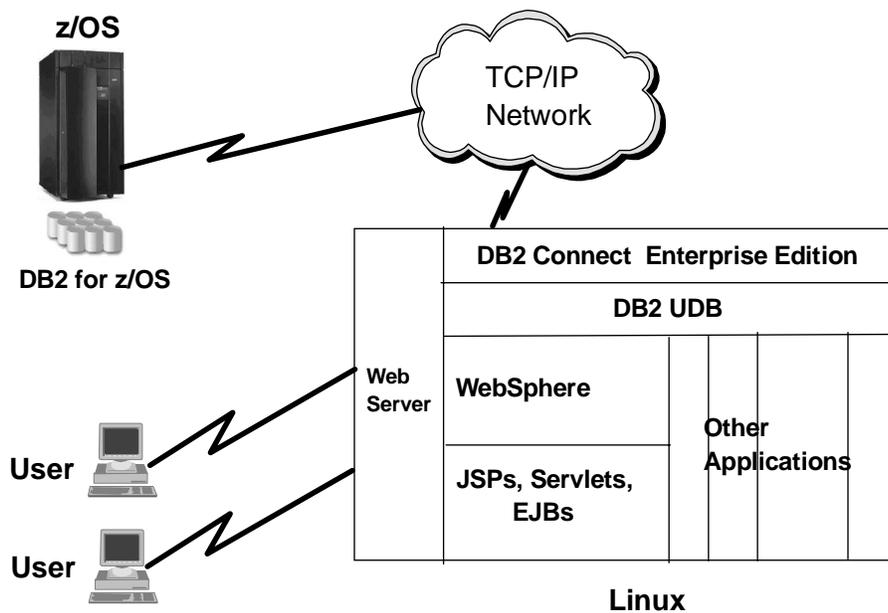
DRDA = Distributed Relational DataBase Architecture

DB2 UDB Application Requester support



* SNA not supported on Linux for S/390 and zSeries

DB2 Connect for Linux



DB2 UDB V8.1 & DB2 Connect matrix



No 64-bit DB2 version available

| Distribution | Kernel | Mode | DB2 UDB & DB2 Connect 7.2 31-bit | DB2 UDB & DB2 Connect 8.1 31-bit |
|-----------------------|--------|--------|----------------------------------|----------------------------------|
| SuSE 7.0 | 2.2.16 | 31-bit | ok | - |
| SuSE SLES7 (1) | 2.4.7 | 31-bit | ok | ok |
| SuSE SLES7 (1) | 2.4.17 | 64-bit | ok | - |
| SuSE SLES8 (2) | 2.4.19 | 31-bit | - | ok |
| SuSE SLES8 (2) | 2.4.19 | 64-bit | - | - |
| United Linux 1.0 | 2.4.19 | 31-bit | - | ok |
| United Linux 1.0 | 2.4.19 | 64-bit | - | - |
| Turbolinux Server 6.5 | 2.2.16 | 31-bit | ok | - |
| RedHat 7.1 | 2.4.9 | 64-bit | - | - |
| RedHat 7.2 | 2.4.9 | 31-bit | - | ok |

- (1) SuSE SLES7 Patch CD1 needed
(2) SuSE SLES8 Service Pack 2 needed

DB2 UDB V8.1 technical pre-requisites



■ DB2 UDB V8.1 + Fixpak 3 for Linux for z Series requirements are :

- IBM S/390 9672 G5 or higher or IBM zSeries or IBM Multiprise 3000 hardware
- Disk space for full features installation: ~ 1.235 GB (without JDK)
- JDK 1.3.1 service release 1 for Java support with 2.4 kernel
- RedHat 7.2 (31-bit)
 - kernel 2.4.9-38 downloaded from RedHat website and related IBM OCO modules for RedHat Linux 7.2 S/390 from IBM website
 - glibc 2.2.4
 - and latest recommended patches from IBM for gcc, binutils and glibc
- SuSE Linux Enterprise Server (SLES) 7 (31-bit)
 - kernel 2.4.7-58
 - glibc 2.2.4
- SuSE Linux Enterprise Server (SLES) 8 * (31-bit only)
 - kernel 2.4.19
 - glibc 2.2.5
- United Linux 1.0
 - kernel 2.4.19
 - glibc 2.2.5

* **United Linux 1.0**

DB2 on Linux on zSeries experience



- DB2 V7.2 recommendations:
 - Amount of contiguous memory for buffers pools is limited
 - For small to medium sized datamarts (up to 50GB)
 - Small OLTP workloads
 - DB2 Connect Gateways consolidation
- DB2 V8.1 recommendations:
 - Large query workloads (up to 100-150GB with 2 IFLs and 2 GB memory)
 - Moderate OLTP workloads
 - DB2 Connect Gateways consolidation
- Optimize DB2 Configuration
- Spread DB2 Tablespaces across ESS Logical Control Units (LCU) to avoid contention on LCU
- Configure properly the number of channel available to DB2 subsystem
- Reorganize DB2 Tablespaces

Linux on zSeries for Business Intelligence Solutions



Benefits of z/OS for datawarehousing



- Single Operating system image (1 to N processors in a Sysplex)
- Easy Upgrade paths:
 - Upgrades from smallest-largest
 - Nondisruptive Upgrades(Capacity On Demand)
- Horizontal Memory Growth (Dataspaces & Hiperspaces)
- Full Connectivity to all disk
- Integrated LAN connectivity
- Superior Systems Management (Workload Manager, ...)
 - Prioritize work based on query execution time
 - Handle unpredictable peaks
- Unmatched Availability - Availability by design
 - Concurrent maintenance
 - Online DB reorganization
- Exclusive hardware compression
- z/OS - UNIX Branded Operating System
 - Highly accessible environment
 - Supports all end-user workstation tools
- Linux Operating System Support
- Warehouse/Mart Coexistence

Benefits of Linux on zSeries for BI



- Consolidation of Servers
 - Data Servers
 - Data Warehouse, ODS, Data marts
 - Application Servers
 - Off the shelf OLAP and analytics
 - Customized Roll Your Own
 - Middleware Servers
 - DB2 Connect
 - Web Servers
 - Web enablement of data marts and applications

Benefits of Linux on zSeries for BI

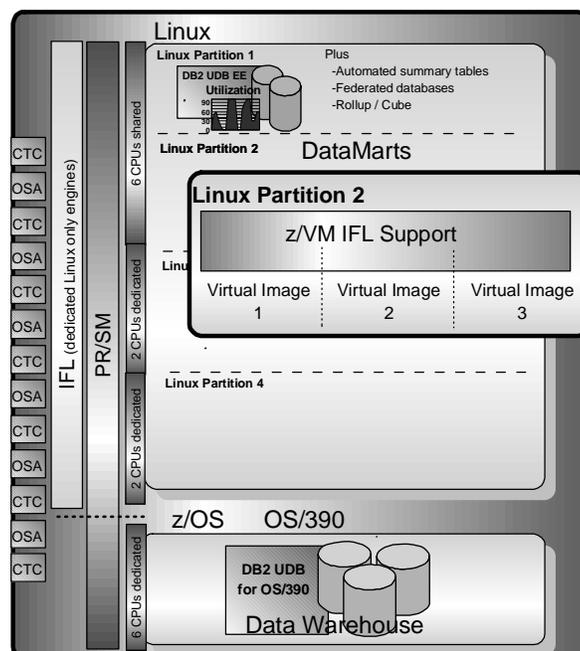


- Lower cost of implementing data marts and BI applications on zSeries
- Cheaper processors - Integrated Facility for Linux (IFL)
- OTC IBM / nonIBM software for Linux on zSeries
 - DB2 UDB
 - z/VM running on IFL
 - Linux distributions
 - Off the shelf applications
- Rapid deployment of data marts and applications
 - Simple recompile of applications - no porting!
 - Easy deployment on z/VM Linux guest systems
- Combined strengths of DB2's
- Better utilization of processing resources -> higher ROI
 - Resource sharing with application/data integrity
- High speed data transfers

Rapid Deployment of marts and applications



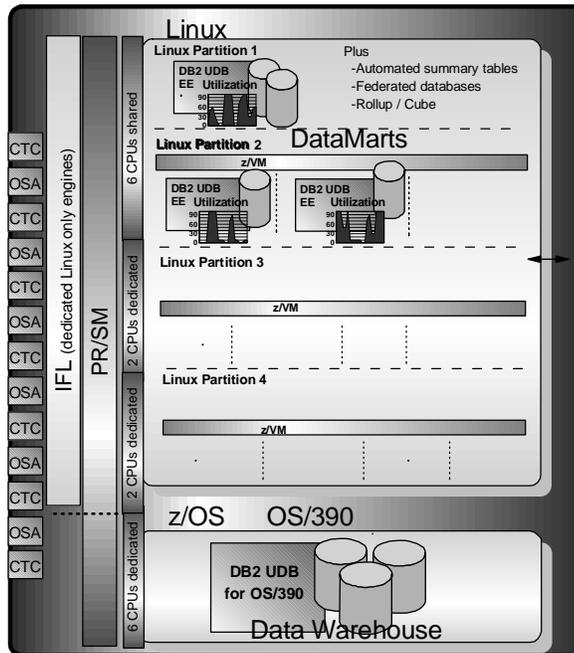
- Linux guests under z/VM
 - Dynamic additions of new Linux images
 - Fast image replication features
 - Dynamic easy assignment of shared resource
 - No z/VM skill requirements



Combined strength of DB2s



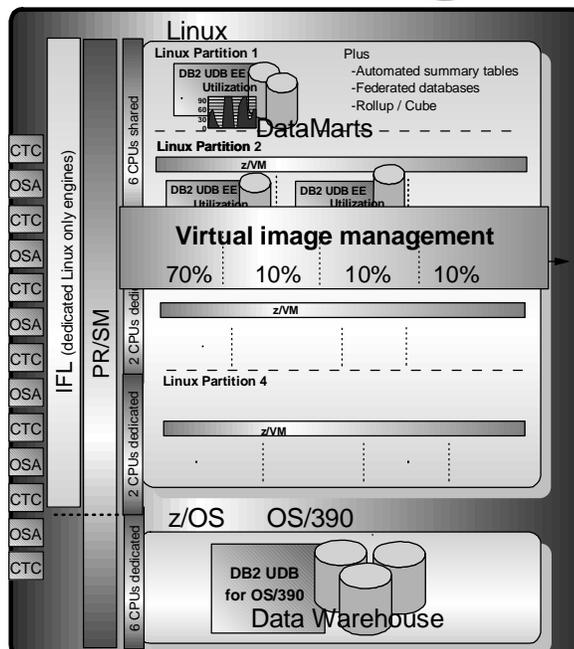
- The strengths of DB2 on z/OS on your Warehouse
 - Dynamic Workload Prioritization
 - High Availability
- Combined with rich end user functionality of DB2 UDB EE on Linux datamarts:
 - Automated summary tables
 - Federated databases
 - Rollup/Cube



High utilization with intelligent resource sharing



- Utilize processor "White Space"
 - Not possible with separate physical servers
- Dynamic resource prioritization
 - Based on business needs
- Maintain complete application isolation and data integrity





Linux on zSeries

Early Experiences on BI Solutions



Scalability and performance tests for

- Data population
- Queries

Data population scalability & performance

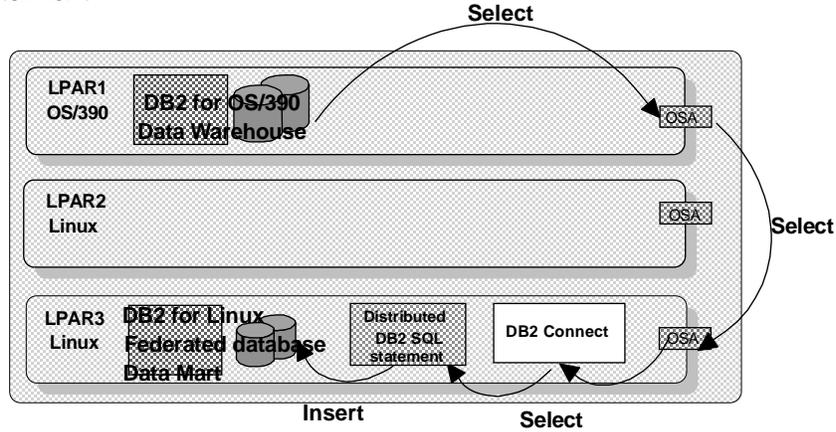


- SQL-based scenarios
 - Single distributed SQL statement
 - Select/Insert SQL statements in a C program
- DB2 utilities-based scenarios
 - DB2 DSNTIAUL, FTP, and DB2 Load
 - DB2 High Performance Unload, FTP and Load
 - DB2 Export and Import
 - DB2 Export and Load
- Scope of tests
 - Atomic-level analysis of simple data extraction and insertion techniques
 - Small volumes of data transferred
 - 1, 5, 100 MB and 1 GB of data
 - 1 CP - no indexes

Single distributed SQL

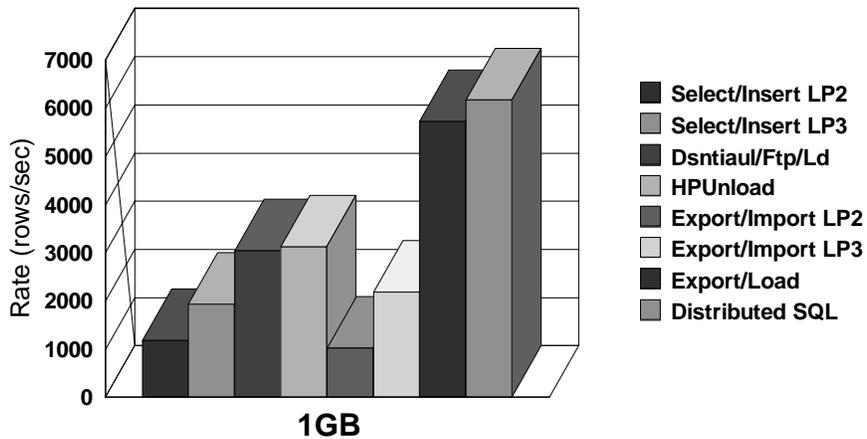


- DB2 for Linux on zSeries
 - Federated database
 - Datamart

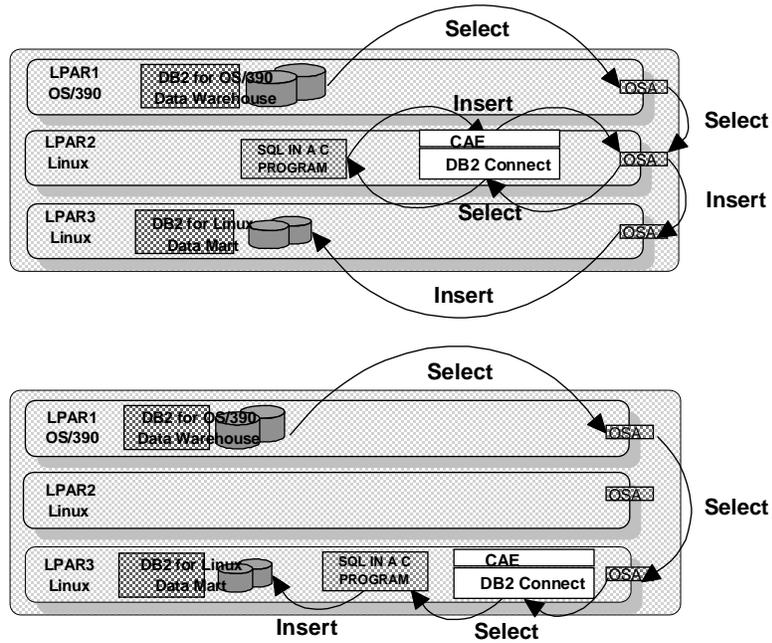


insert into DMtable1 select * from DWtable1 where orderkey < 123456

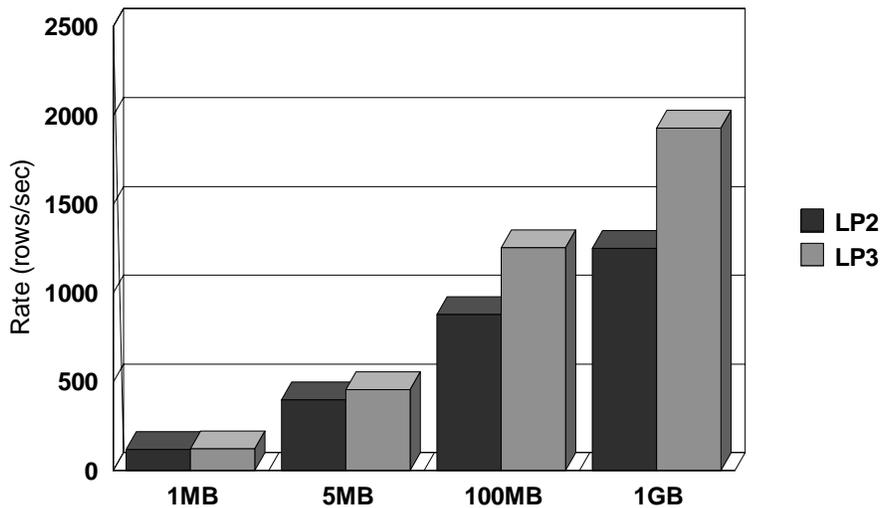
Single distributed SQL - rate summary



Select/Insert in a C program



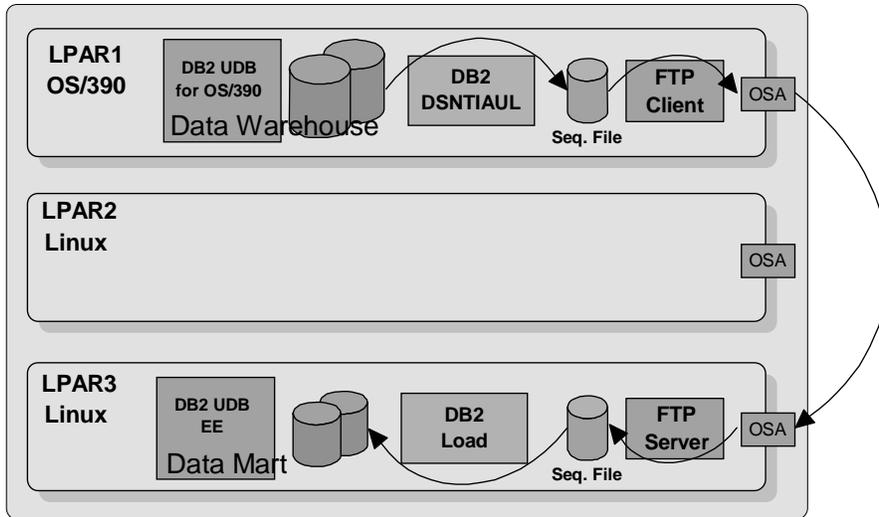
Select/Insert in a C program - rate



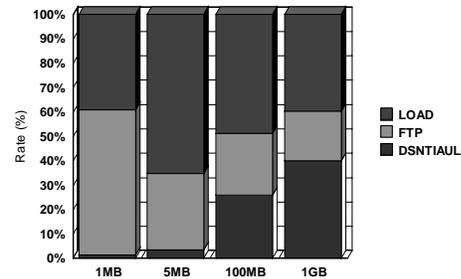
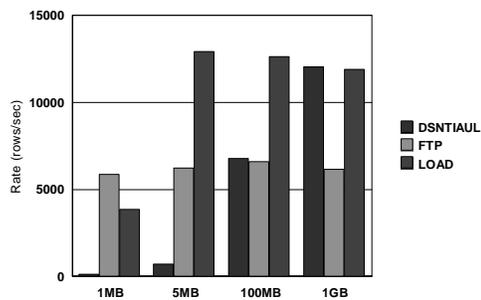
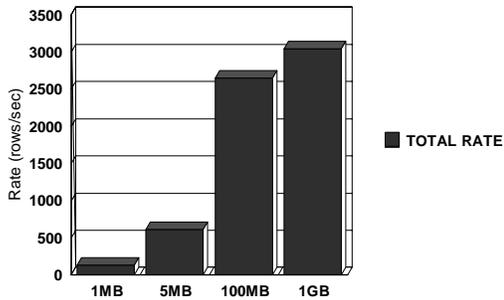
DSNTIAUL, FTP, LOAD



z-Serie



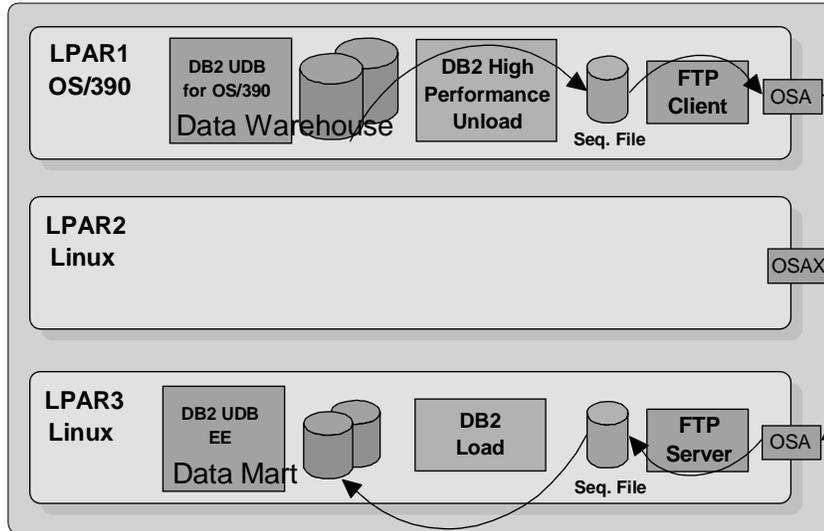
DSNTIAUL, FTP, LOAD - rate



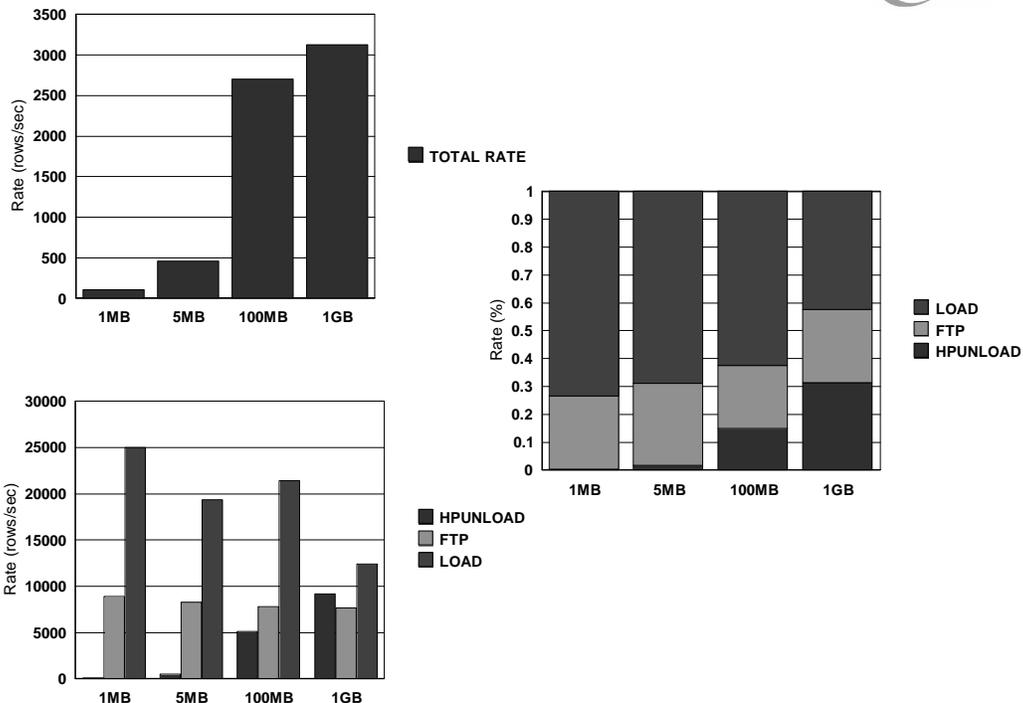
HPU, FTP, LOAD



z-Series



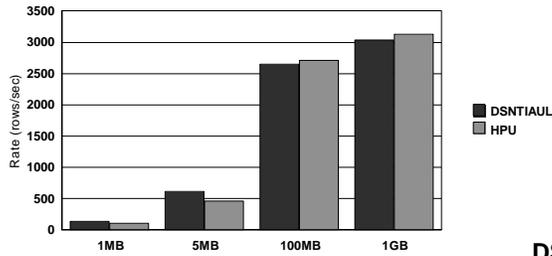
HPU, FTP, LOAD - rate



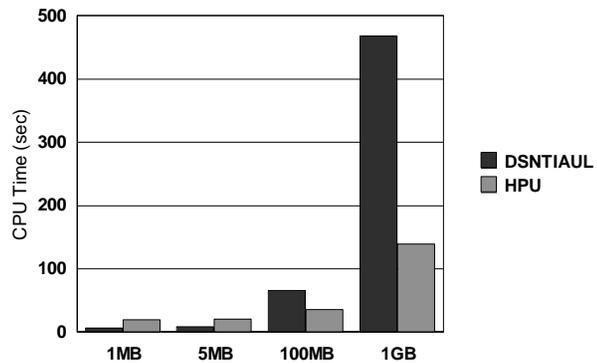
DSNTIAUL and HPU comparison



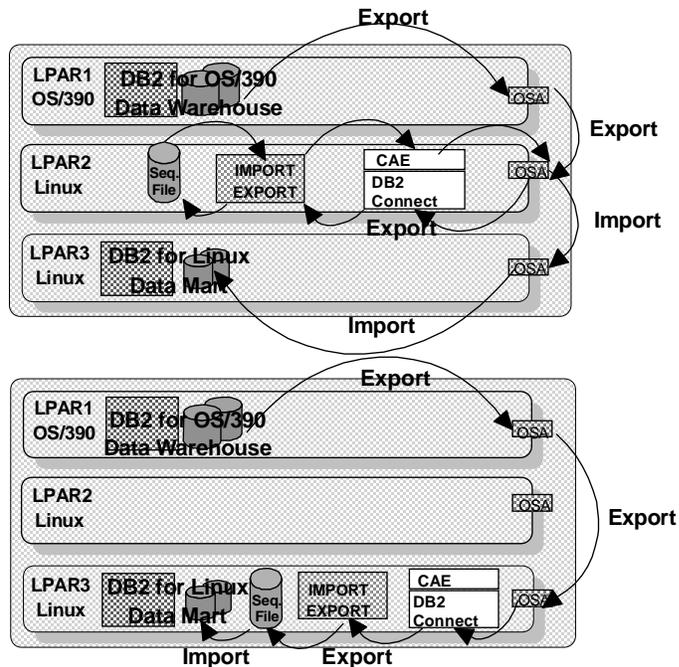
DSNTIAUL vs HPU - Rate



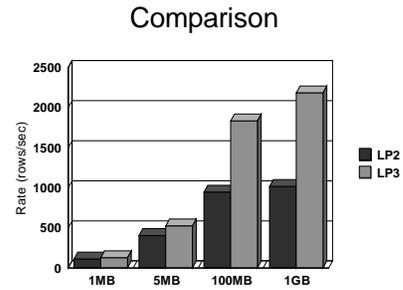
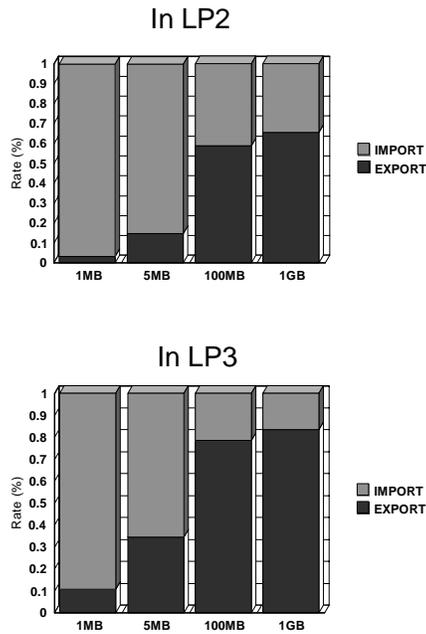
DSNTIAUL vs HPU - CPU Time



EXPORT, IMPORT



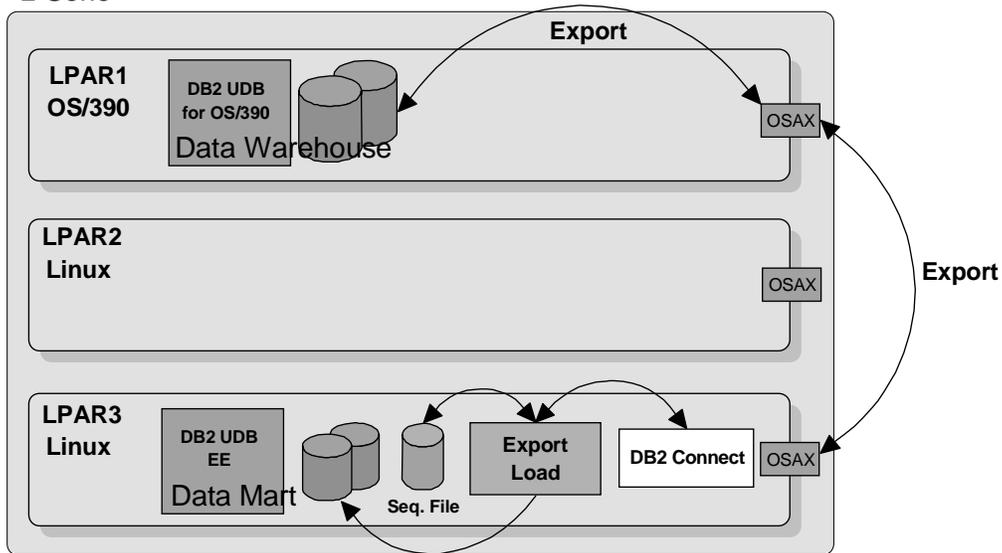
EXPORT, IMPORT - rate



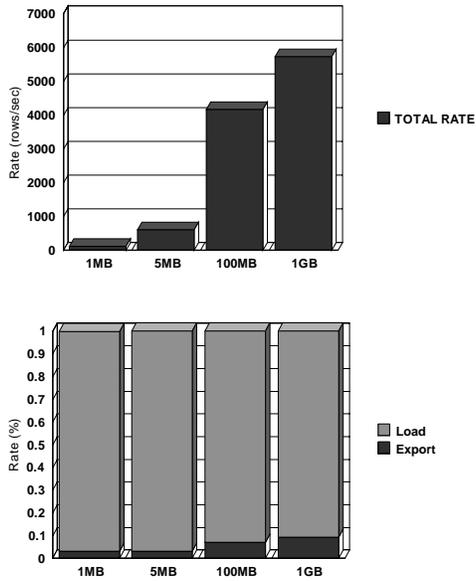
EXPORT, LOAD



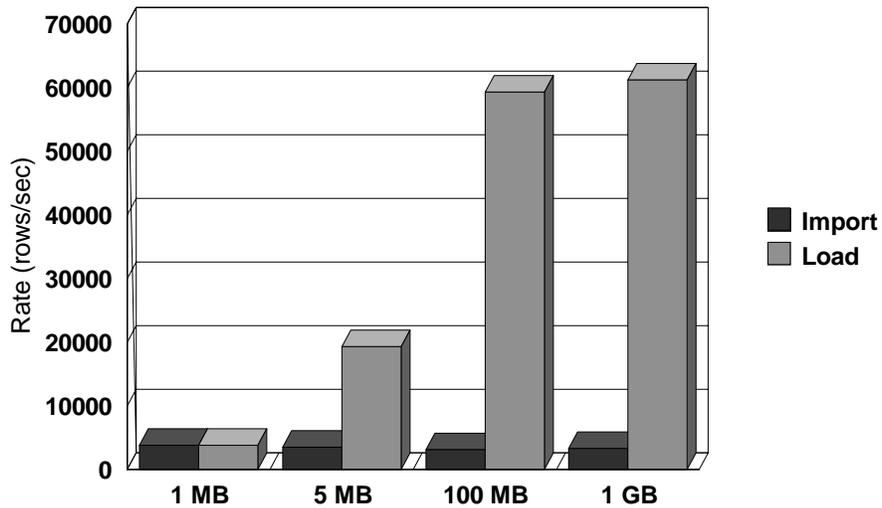
z-Serie



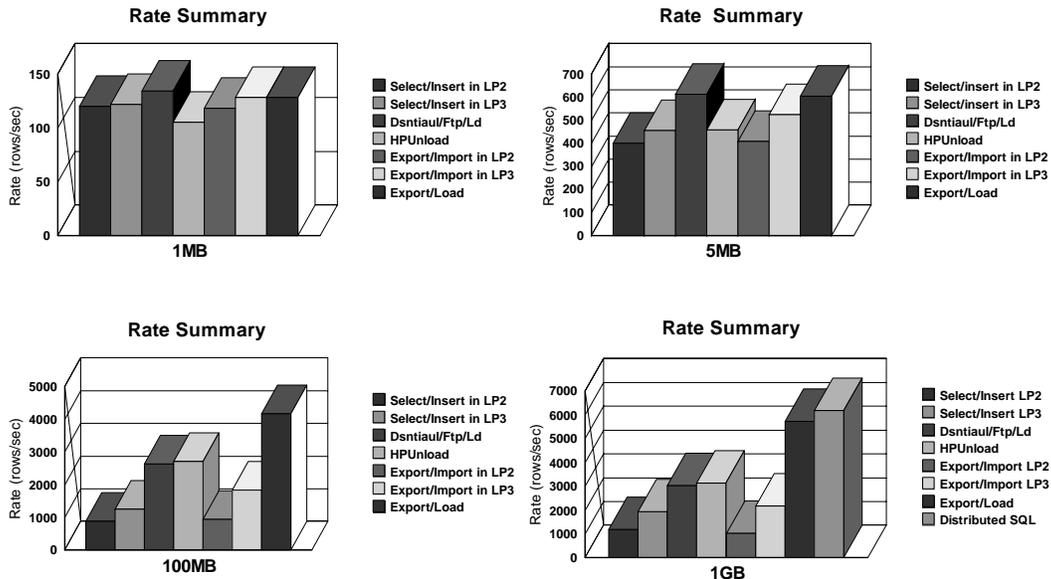
EXPORT, LOAD - rate



IMPORT and LOAD - rate comparison



Data population - performance summary



Data population tools summary

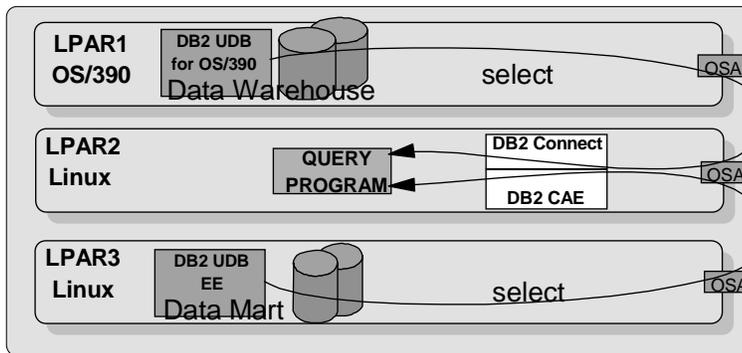


- What we tested
 - SQL statements
 - dynamic
 - embedded in a program
 - DSNTIAUL/LOAD
 - HPU/LOAD
 - EXPORT/IMPORT
 - EXPORT/LOAD
- Other DB2 tools
 - Data Propagator
 - Data Warehouse Manager
 - Data Joiner or Relational Connect
 - DB2 AutoLoader
 - DB2 for z/OS Unload utility (V7)

Query scalability & performance



Query scenarios



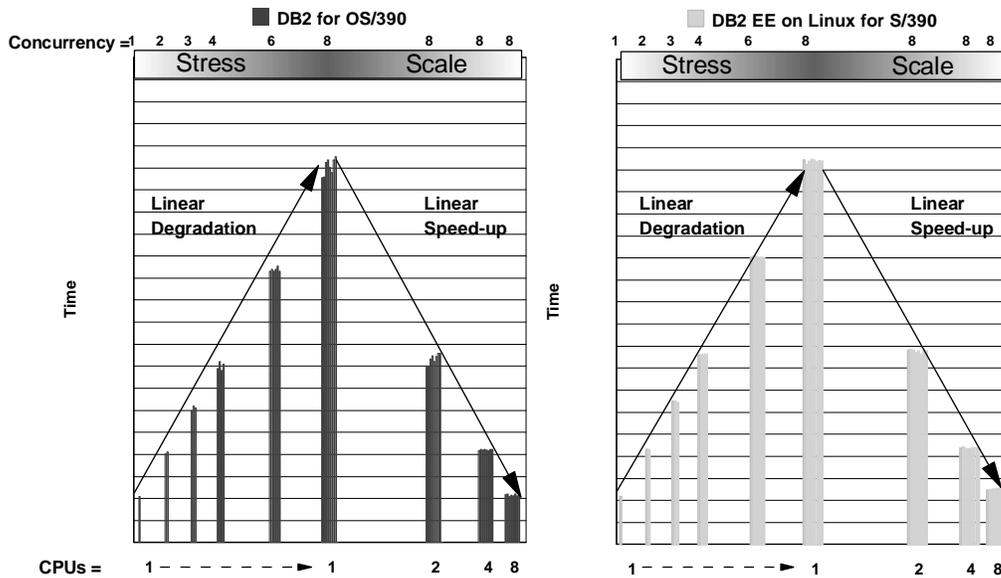
Scope of the tests

- Examine primitive attributes such as CPU, I/O, Memory

Workload objectives

- Processor intensive queries - Processor scalability
- I/O intensive queries - I/O, network, sort scalability
- Short queries - Memory scalability, buffer pool efficiency
- Trivial queries - What seems trivial on OS/390 is it also on Linux?

Processor-intensive queries



1, 2, 3, 8 concurrent queries run on 1, 2, 4, 8 processors
 → against DB2 for OS/390 and against DB2 for Linux on S/390

I/O intensive, short and trivial queries



- I/O intensive queries
 - 0 rows returned -- I/O scalability
 - 1,2,3,8 queries run each against a table on a different disk volume.
 - 2% and 4% ET elongation with 2 and 3/8 concurrent queries
 - Same behavior on both OS/390 and Linux
 - 40000 rows returned -- Network scalability
 - 7% ET elongation due to 40000 rows returned (network transfer cost) on OS/390
 - Cost is 8% on Linux
 - 40000 rows returned with Sort
 - Sort cost was 5% ET elongation on OS/390
 - 1% on Linux --- DB2 for Linux always performs a Sort!
- Short queries
 - Observe memory scalability and buffer pool efficiency
 - Read table from disk / from buffer pools (table is in memory)
 - Comparable ET on both OS/390 and Linux
- Trivial query
 - Observe if a trivial query on OS/390 is also trivial on Linux
 - Tablespace scan of small table (5 rows)
 - Subsecond ET on both OS/390 and Linux

Linux on zSeries More Experiences on BI Solutions

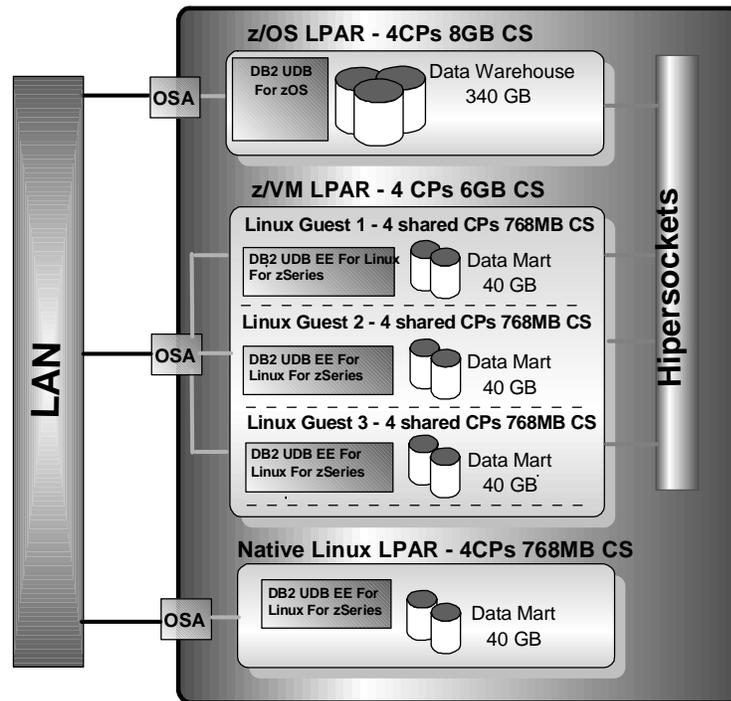


Performance tests for:

- Data population
- Data refresh
- Distributed Join
- z/VM resource allocation
- Dimensional Insight BI tool

System configuration

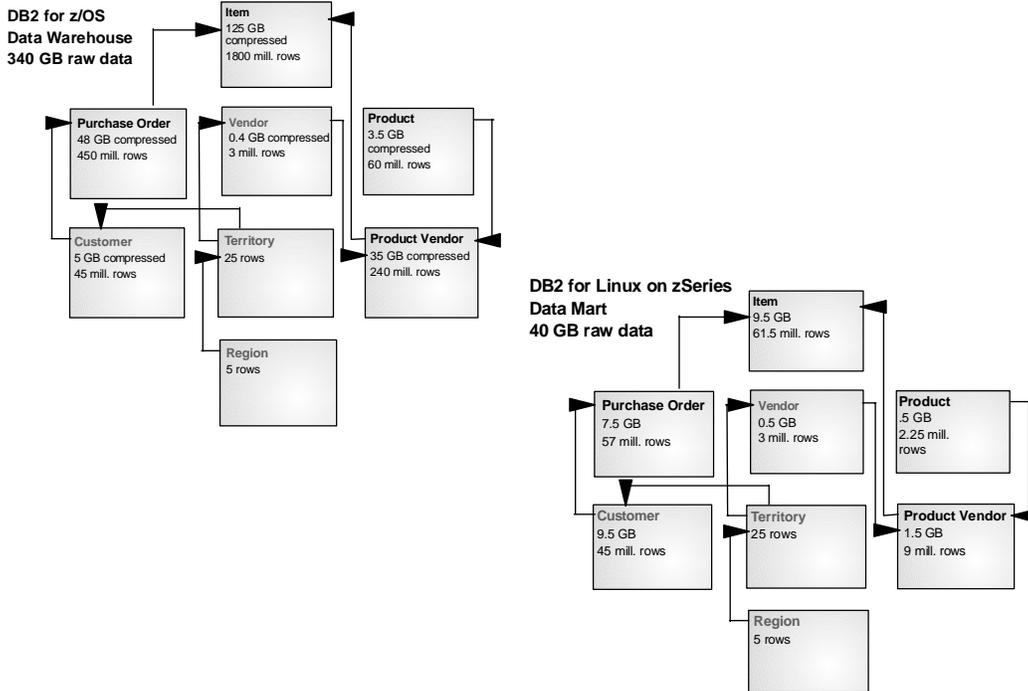
z900 2064-116



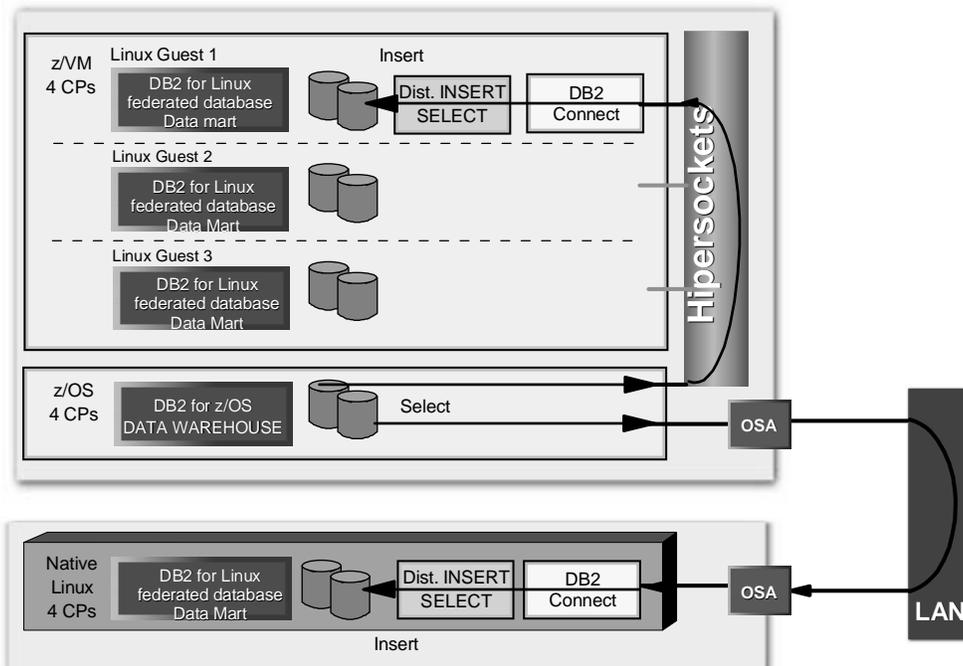
Configuration details

- H/W
 - 12 processors
 - 16 GB memory
 - 3 OSA Gb cards
 - 1.6 TB ESS box
- Images
 - z/OS 1.2
 - z/VM 4.2
 - 3 Linux guests 2.4 kernel
 - Native Linux 2.4 kernel
- Application Deployment
 - S/W
 - DB2 UDB for z/OS V7.1
 - DB2 UDB EE V7.2
 - Omegamon for DB2
 - DB2PM
 - Omegamon for MVS
 - RMF for z/OS
 - sar for zLinux
 - vmstat for zLinux
 - VMPRF
 - Dimensional Insight
 - Atlantis suite

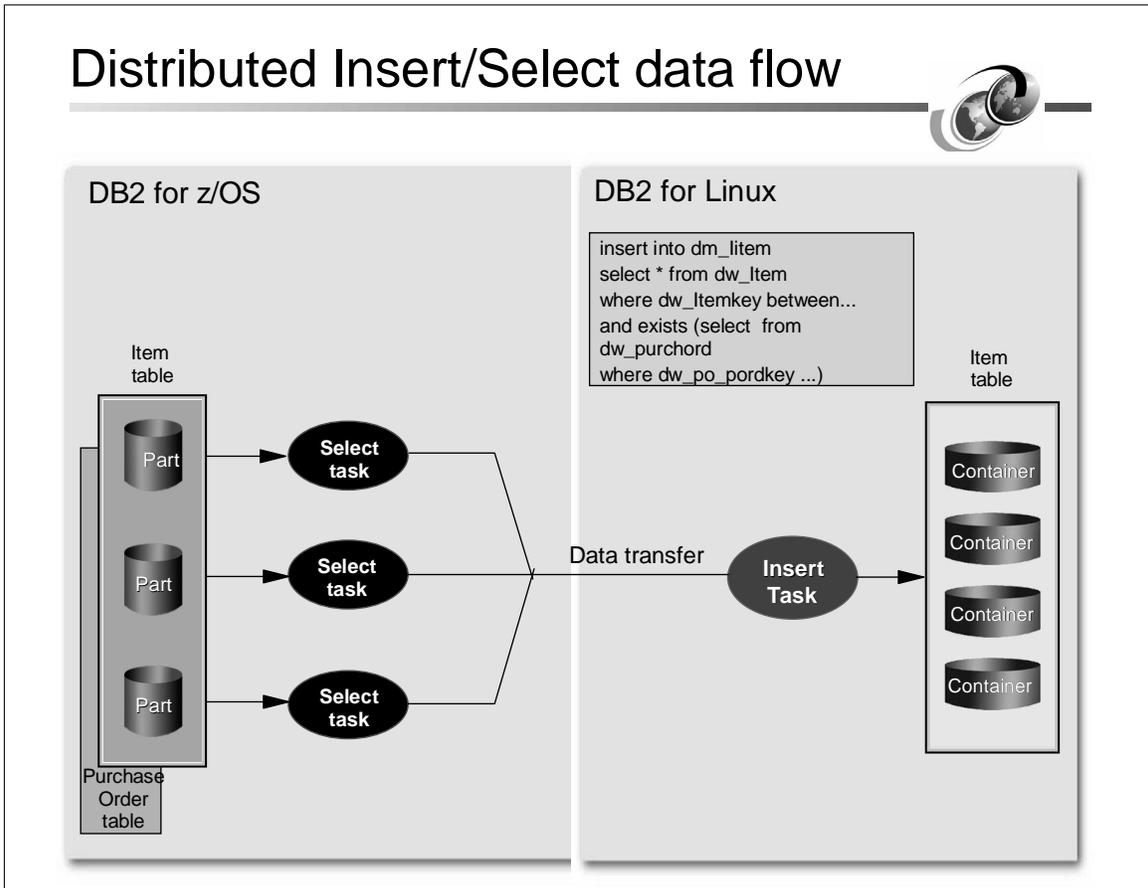
Data warehouse/mart data model



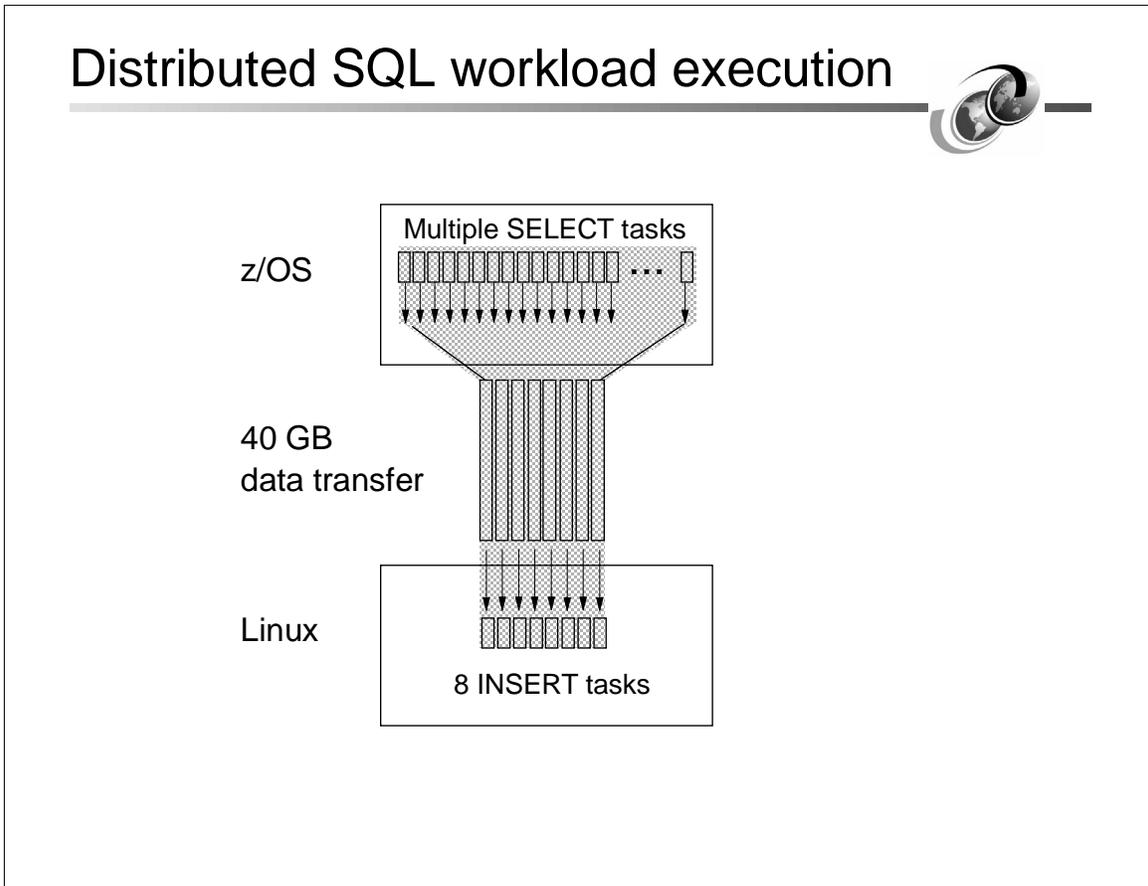
Data Population - distributed SQL scenario



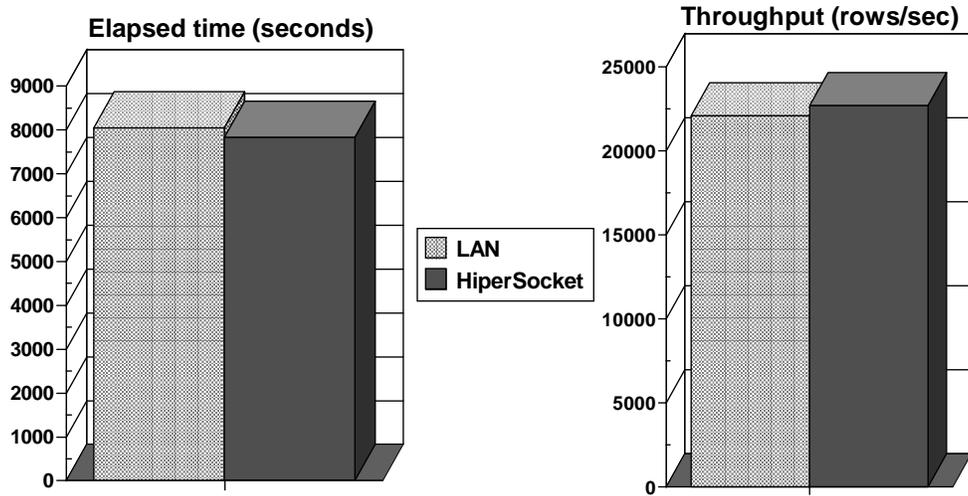
Distributed Insert/Select data flow



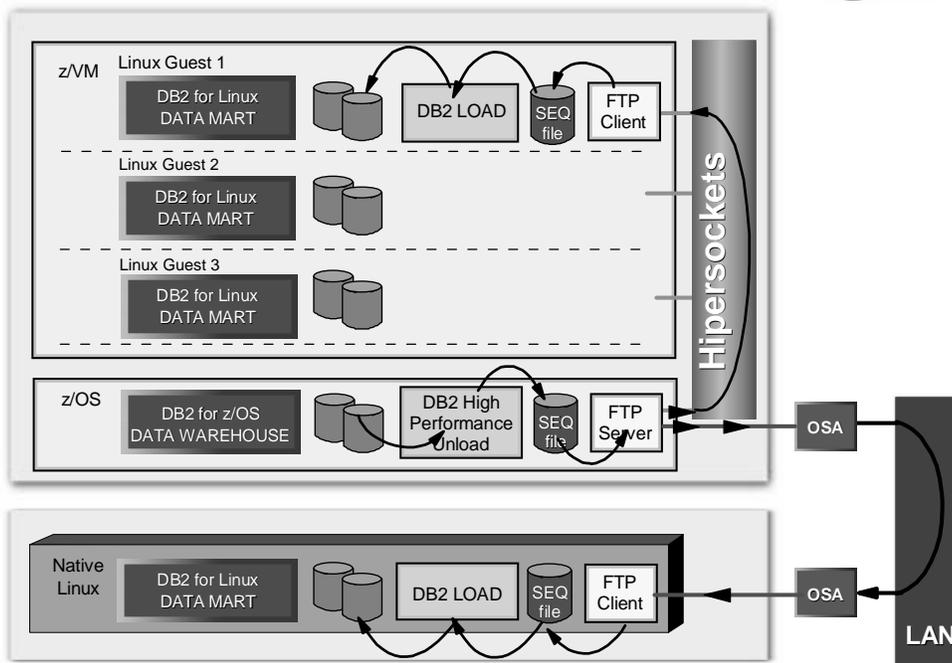
Distributed SQL workload execution



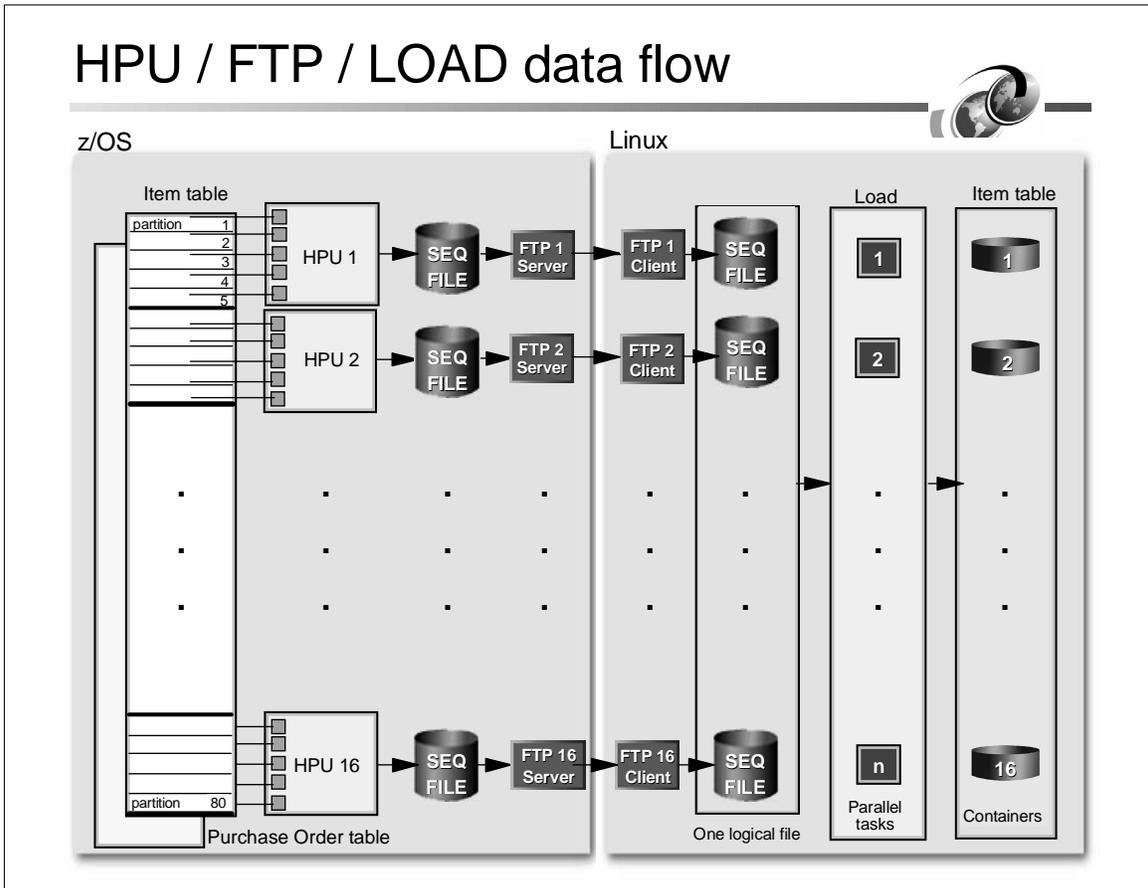
Distributed Insert/Select data transfer LAN vs HiperSockets



Data population - HPU/FTP/LOAD scenario

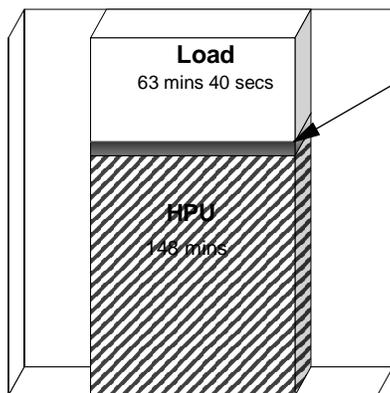


HPU / FTP / LOAD data flow

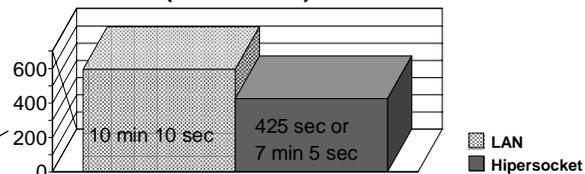


FTP data transfer performance

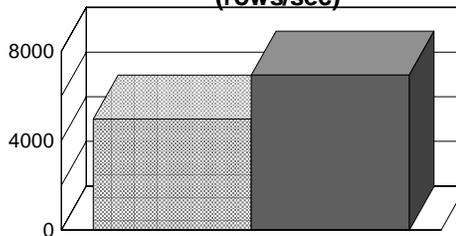
HPU/FTP/LOAD Elapsed Time



FTP Transfer Time (in seconds)

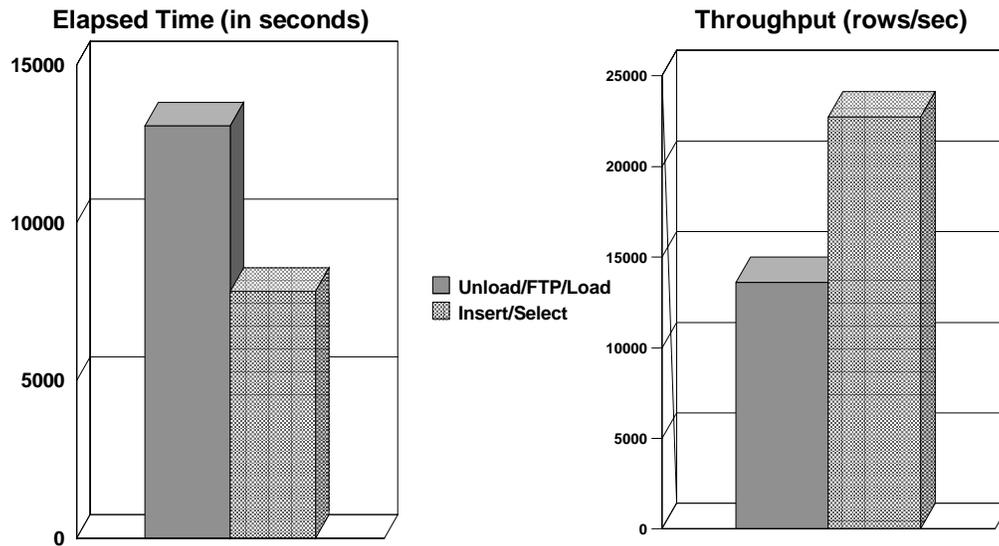


FTP Throughput (rows/sec)



- HiperSockets perform better
 - memory data transfer - 7 min 5 sec to transfer 40 GB of data
- LAN connection is slower
 - media speed, network traffic impact, device contention, cost of networking

HPU / FTP / LOAD vs distributed Insert/Select



- Insert/Select is 40% faster
- Populates 40 GB datamart in slightly more than 2 hours

Unload / FTP / Load vs. Insert/Select

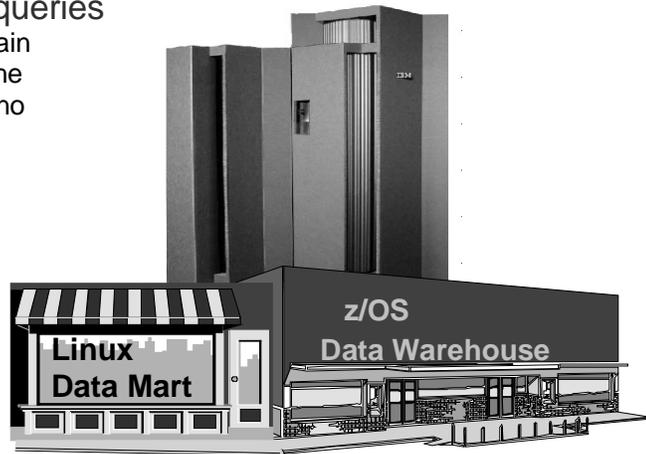


- Unload/FTP/Load
 - Can be broken down into 3 processes
 - Each can be performed at different times depending upon resource availability
 - Unload from image copies with the HPU utility to eliminate contention with online work
 - Spawn multiple tasks to unload a table in parallel
- Insert/Select
 - Online process executing in parallel in one step
 - Easy to implement
 - Turn off logging during data population
 - Use table-level locking

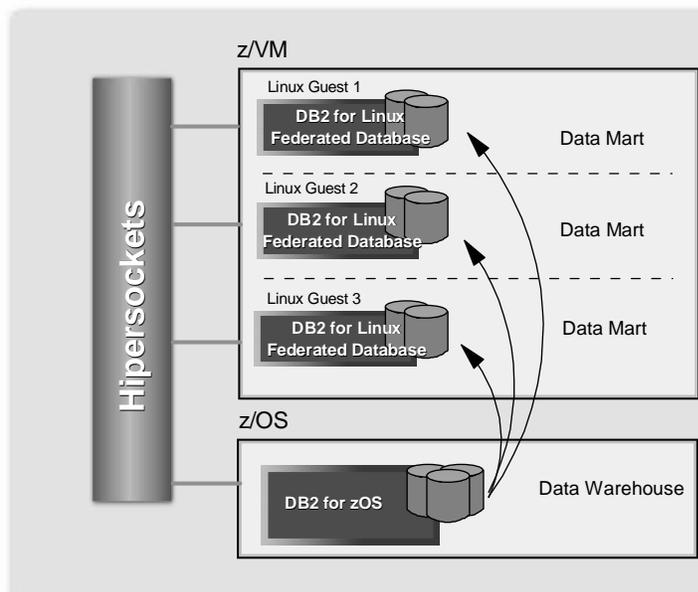
Data Refresh



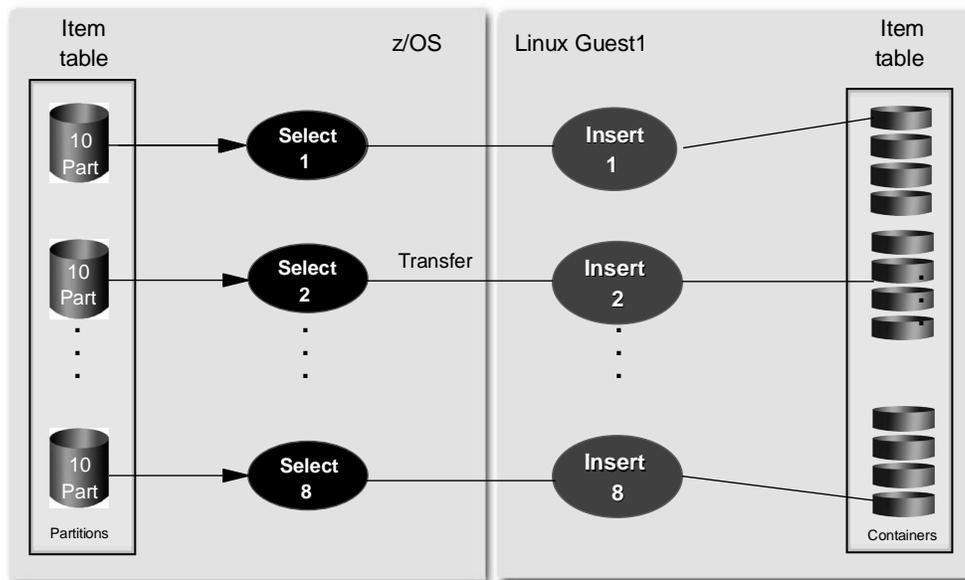
- Weekly refresh of 2 large tables in Linux data mart
- A total of 2 million rows (290 MB) selected from z/OS data warehouse and inserted into a Linux data mart on z/VM using HiperSockets
- Complex correlated subqueries used for insert-select to maintain relationship of data between the tables (SQL statements only, no utilities required)



Data Refresh - concurrent refresh scenarios



Data Refresh data flow



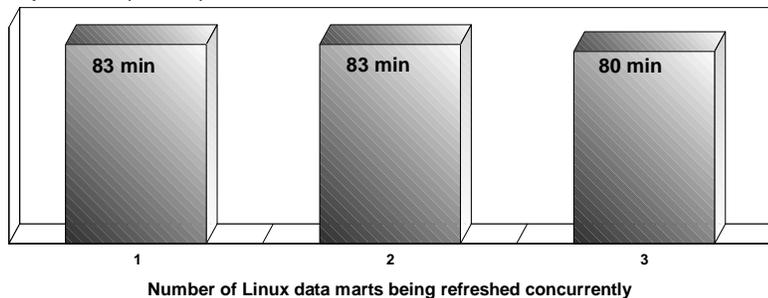
Data Refresh scalability



- Ability to support concurrent data refreshes of multiple Linux data marts
- High internal network bandwidth (HiperSockets)
- No elongation in elapsed time
- Maintenance window is shortened by refreshing all marts concurrently



Elapsed Time (minutes)

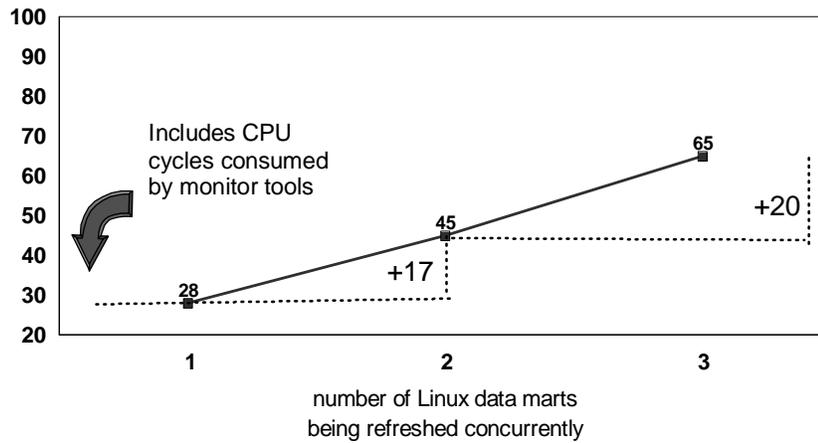


Data Refresh scalability...



concurrent data refreshes  of multiple Linux data marts

CPU Utilization % (z/OS)

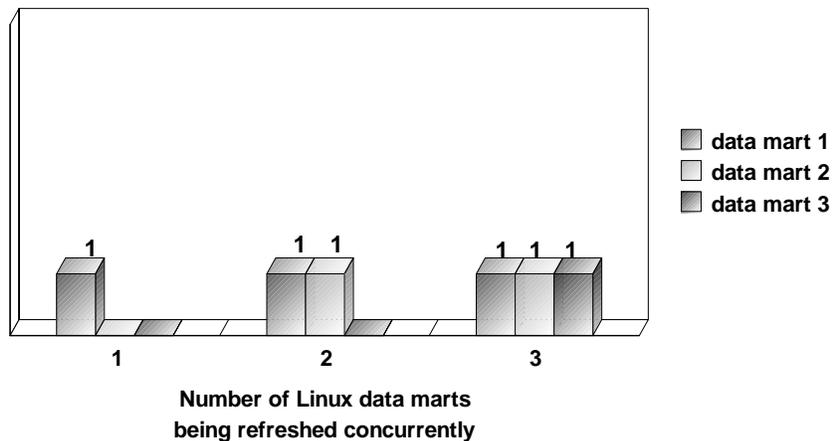


Data Refresh scalability...



concurrent data refreshes  of multiple Linux data marts

CPU Utilization % (Linux)



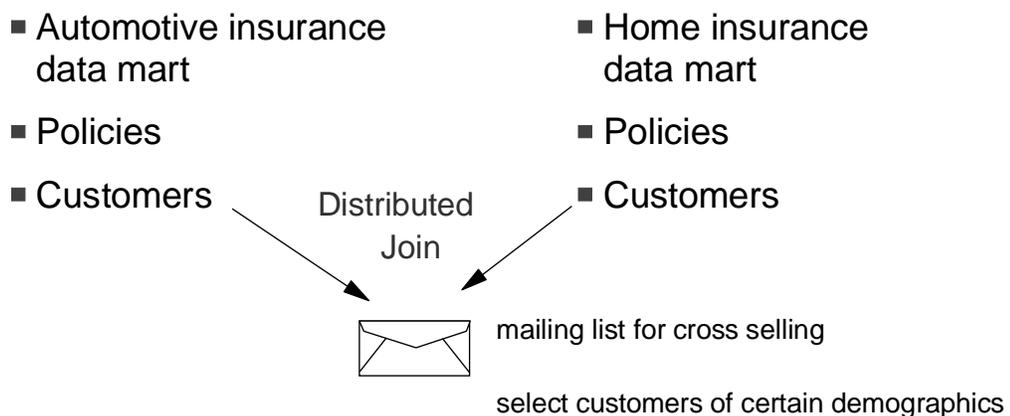
Only 1% CPU utilisation per Linux data mart during data refresh

Distributed Join



- DB2 federated database support provides transparent SQL access and join capabilities across heterogeneous data sources
- Using DB2 federated database support you can view your enterprise data as if it were local
- With a single SQL statement you can access and join tables located across multiple data sources without needing to know the source location
- No application coding is required for distributed joins

Insurance Policy scenario



Retail scenario



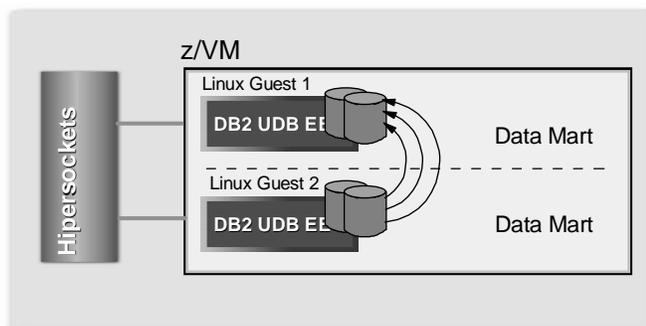
Warehouse



- Sales data mart
 - Transactions
 - Items
 - Quantities
 - Date
- Distributed Join
- Inventory data mart
 - Activities
 - Items
 - Quantities
 - Date

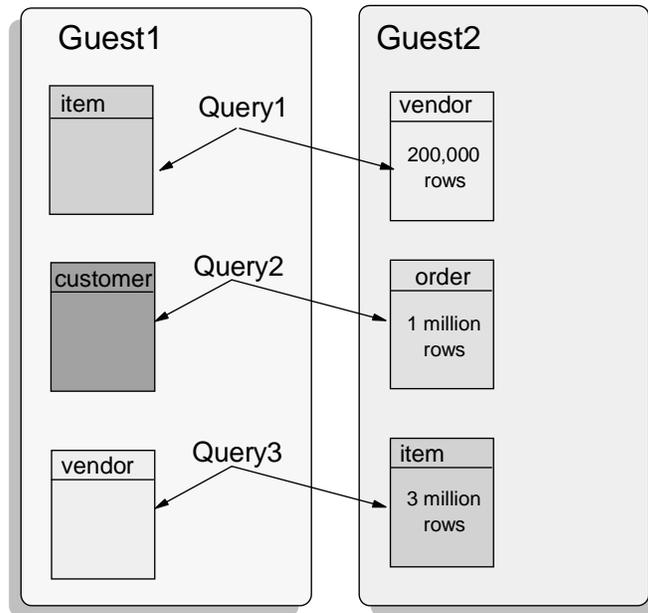
Potential loss of sales due to depletion of items
Excessive inventory driving up cost

Data movement

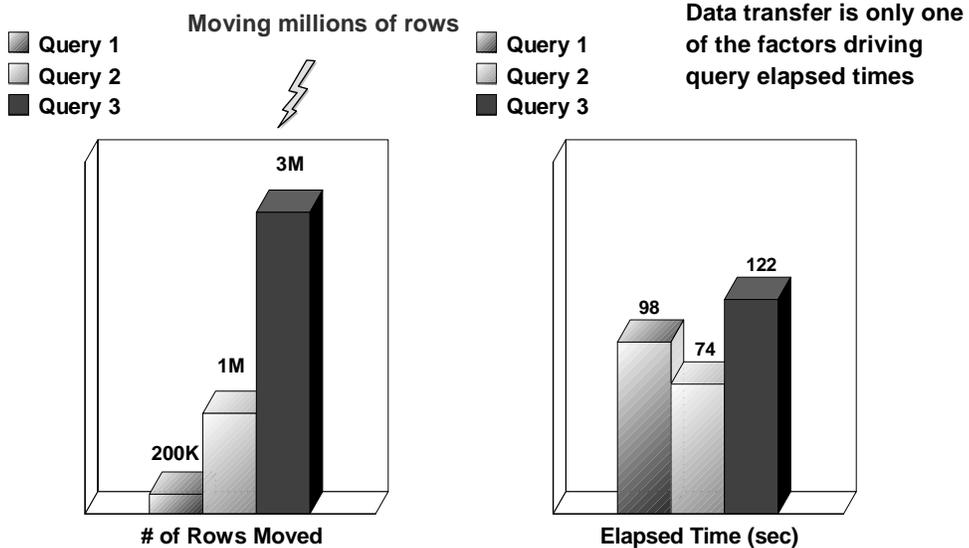


High speed
distributed joins
at internal
memory speeds!

Distributed SQL queries



Query performance



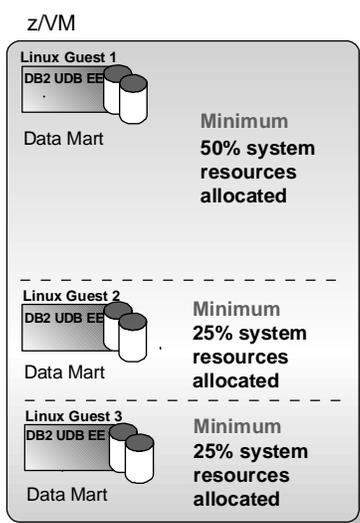
No special DB2 tuning - DB2 optimizer tuned to handle distributed queries
 Majority of Elapsed Time spent in query processing, not data movement (Q2)

Intelligent resource allocation

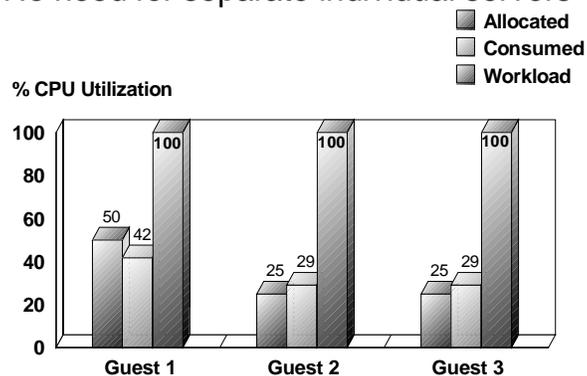


- Resource allocation based on individual business unit needs
 - Multiple applications on a single physical server
 - Resource sharing among applications
 - Full application isolation and data integrity
- Reduce costs by utilizing "white space"
 - Low cost of ownership due to high resource utilization
 - Provide capacity when/where it is needed
 - Dynamic workload prioritization
- Manage individual business unit fixed budgets
 - Resource capping to prevent budget overrun

Resource allocation - individual business unit needs



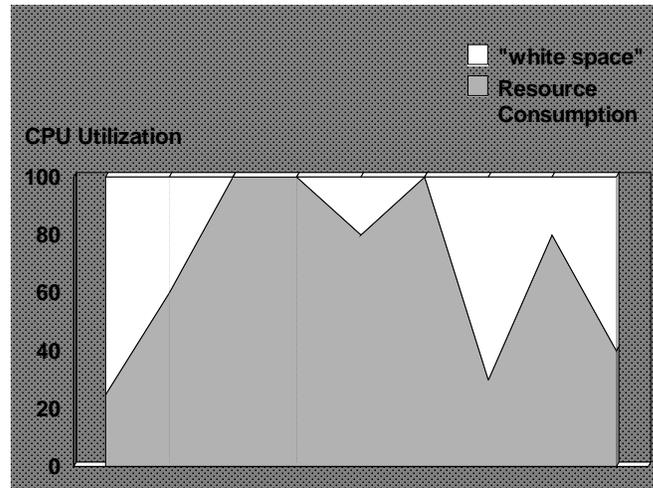
- Resource needs vary between data marts
- Define resources for each Linux Guest
- z/VM manages resources across Linux Guests
- No need for separate individual servers



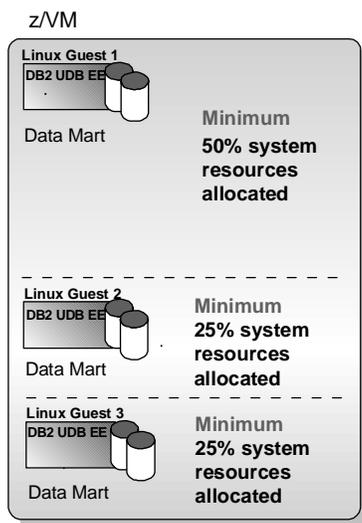
"White Space"



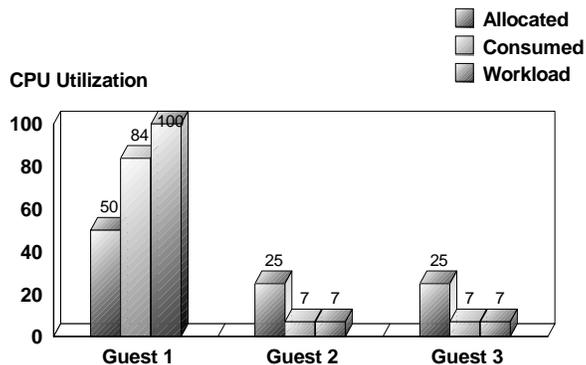
Under utilized (wasted) system resources



Reduce cost by utilizing "White Space"



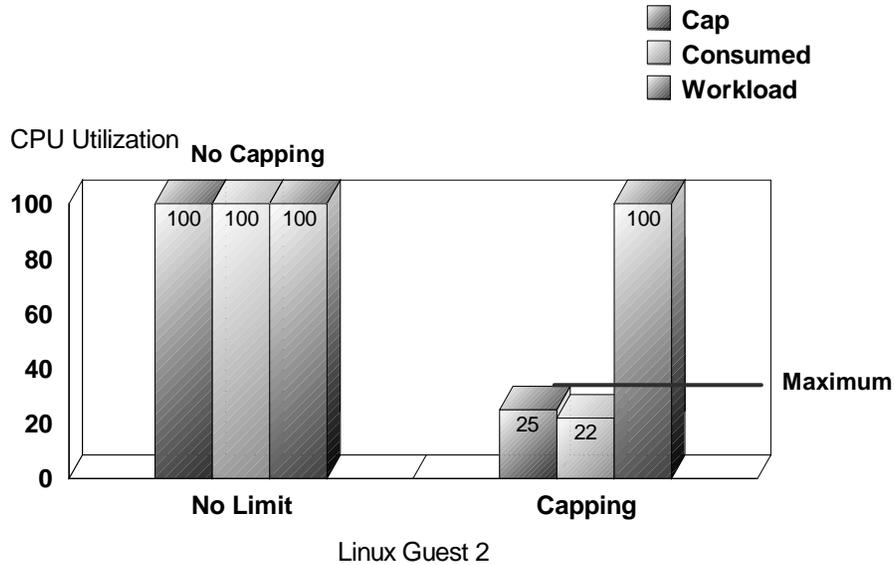
- Capacity when and where it is needed
- Dynamic workload prioritization
- High resource utilization across enterprise
- Lower costs across enterprise



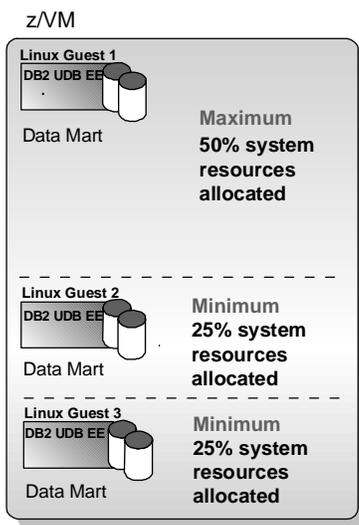
Resource capping



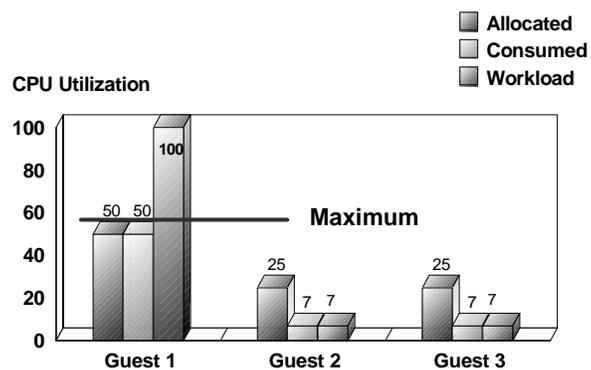
Ability to limit the amount of resources consumed by an individual Linux Guest under zVM



Manage individual business units fixed budgets



- Cap resources to control business unit costs
- Service bureau resource capping
- z/VM limits resources across Linux Guests
- No need for separate individual servers



Application testing



- Dimensional Insight
- DI-Atlantis suite of products
- Web-enabled OLAP solution (MOLAP-based)
- Developed on Linux for eight years
- Migrated to Linux on zSeries with a simple recompile
 - Less than 10 minutes!



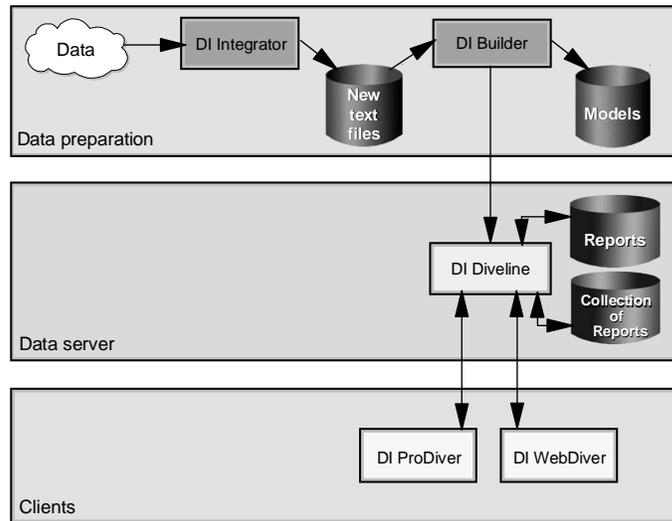
Application testing



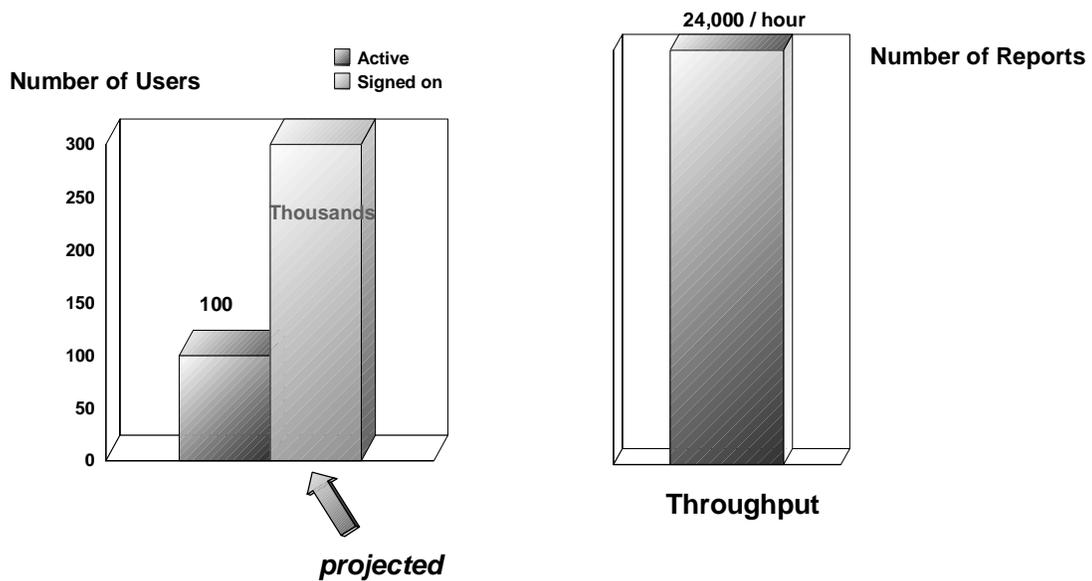
- Features Included:
 - OLAP Multidimensional Business Analysis and Reporting Solution
 - Over 800 Installations
 - Vertical Applications and ERP Integration Templates for
 - Healthcare / Insurance
 - Manufacturing
 - Distribution / Media
 - Models can be built or served from any platform
 - Fully web-enabled across all IBM Platforms
 - Rapid Implementation
 - 30 days from project start to finished datamart
- Platform Support
 - All IBM Servers
 - Remarketing agreement
- Target Audience
 - Small and Medium Customers
 - Low Priced/ High Performance OLAP
- For more information: www.dimins.com



Dimensional Insight product architecture



High concurrency and throughput



Summary of Teraplex studies



- Scalable query performance
- Data mart consolidation feasibility
- Intelligent resource allocation
- High speed data transfer
- Ease of data population and refresh
- Multi-database business analysis with no application coding
- Availability of robust BI applications supporting high concurrency and high performance

Summary of benefits of Linux on zSeries for BI



- Tight integration of data stores: ODS, EDW and data marts all under the same roof
- Enriched BI functionality on zSeries
- Multiple Business Intelligence applications on a single physical server
- Efficiencies of shared resources across multiple application servers
- Rapid deployment of data marts
- True end-to-end BI solution architecture on a single server



Linux on zSeries for SAP Solutions

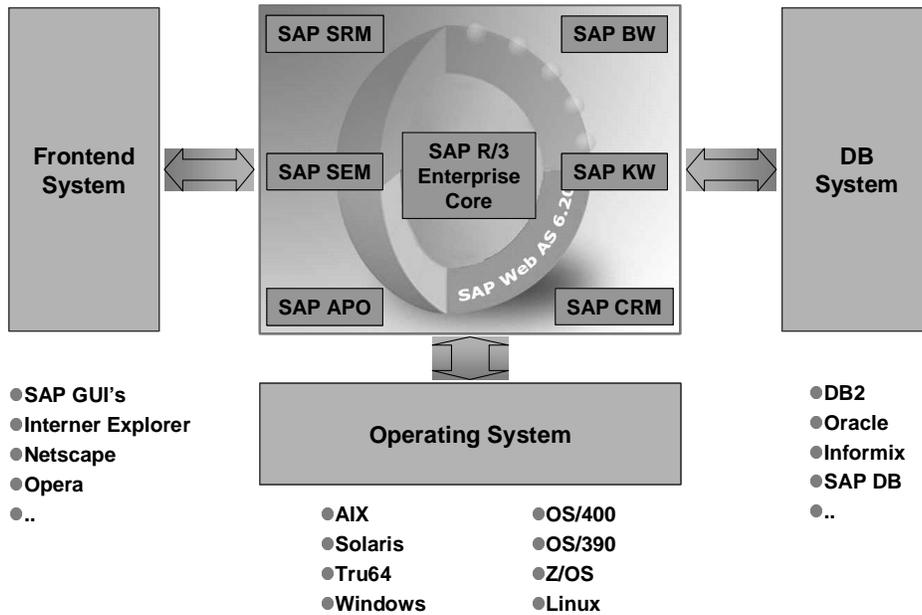


About SAP

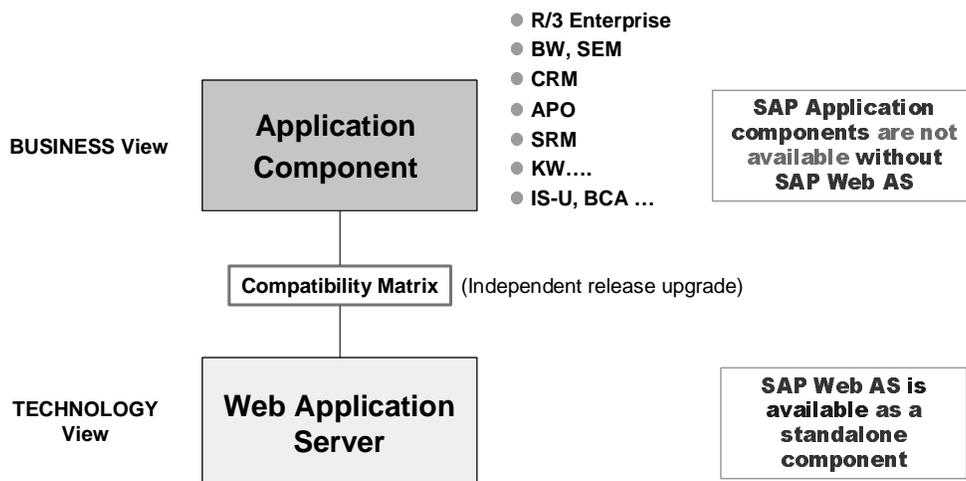


- SAP – 30 Years in the Business of e-Business.
- 12 Million Users. 60,000 Installations.
- 1,500 Partners. 23 Industry Solutions.
- Founded in 1972.
- Headquartered in Walldorf, Germany, SAP is the world's largest inter-enterprise software company, and the world's third-largest independent software supplier overall. SAP employs over 28,900 people in more than 50 countries.
- SAP Business Suite is open and flexible, supporting databases, applications, operating systems, and hardware from almost every major vendor.

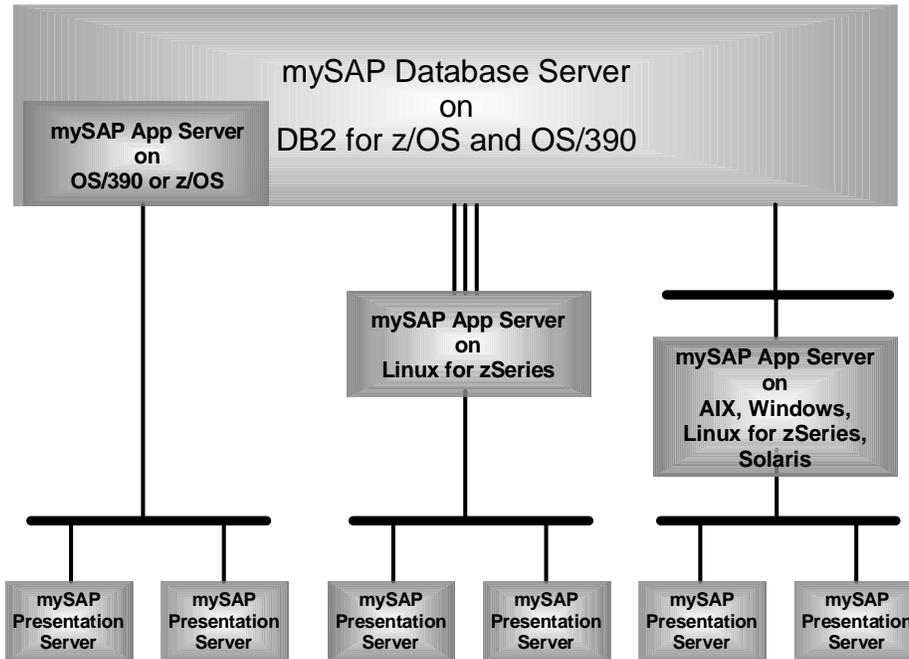
SAP WAS and platform compatibility



SAP Web Application Server



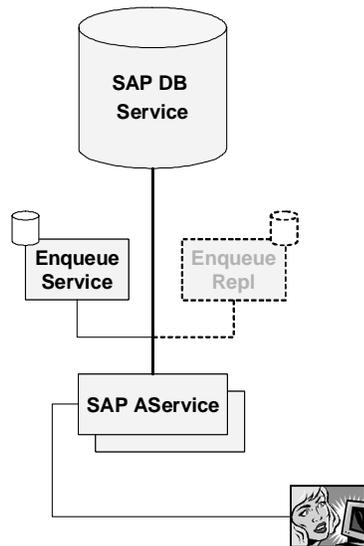
SAP platforms with zSeries



Understand SAP system architecture

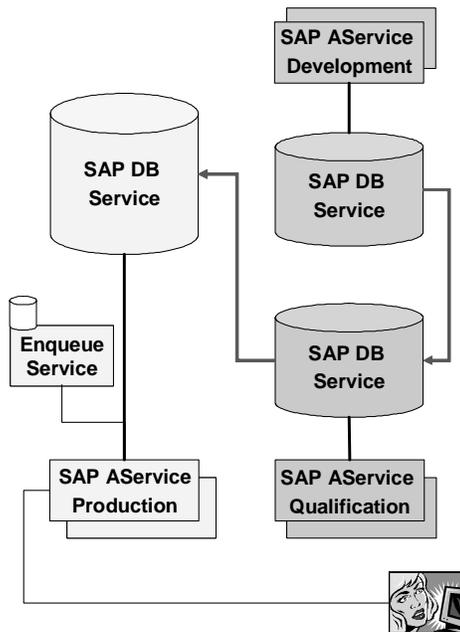


SAP DB Service = RDBMS + Storage Devices



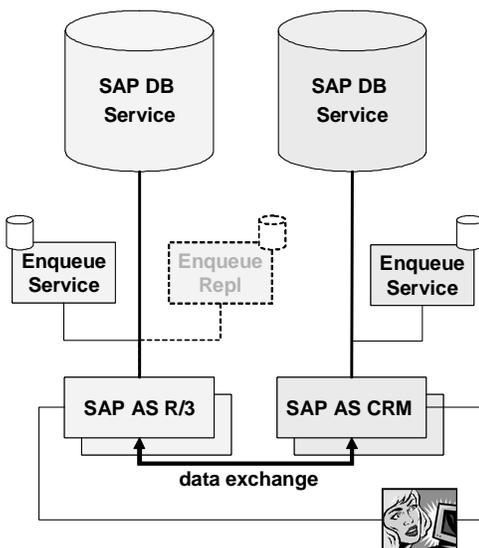
- **SAB DB Service** is single point of contention (No partition / No duplication)
 - Performance (RDBMS, I/O to storage)
 - Database Consistency (BU, Restore)
 - Availability (Outage, Admin)
 - Disaster recovery
- **Enqueue service** is single point of failure
 - Restart capability
- **Application Service dependencies** are customer business dependant
 - Non critical Work Processes (Dialog)
 - Multiple WP for scalability & availability
 - SAP Load balancing, Auto reconnect
 - Critical WP (batch)
 - Scheduling, Reliability
 - High DB Throughput WP (Batch, Update)
 - Network Dependencies

Understand SAP system implementation



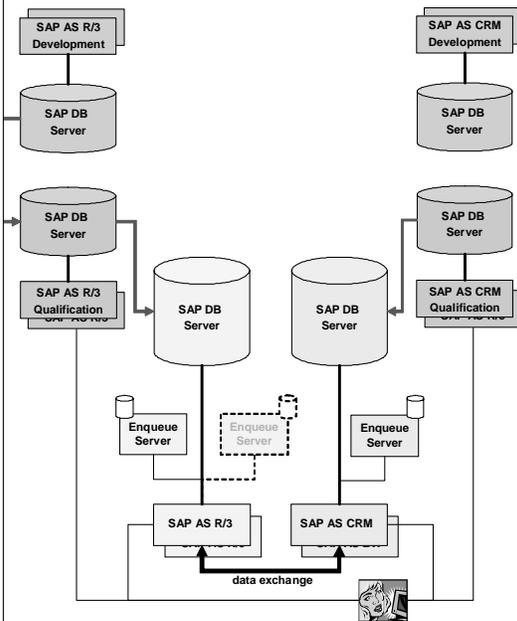
- Several SAB DB Services
 - Service level requirements significantly different
 - Dedicated servers
 - Administration, Unused Capacity
 - Consolidation
 - Isolation (availability)
- Several Application Services
 - Workload profiles significantly different
 - Dedicated servers
 - Unused Capacity, Administration
 - Consolidation
 - Resource management

Understand SAP business solution



- Global SAB DB Services
 - Several single points of contingency to manage concurrently & consistently
 - Performance, data consistency, availability, disaster recovery
 - Dedicated Servers
 - Administration, Unused capacity
 - Consolidation
 - Resource management (R/3, CRM, BW), isolation
- Several Application Services
 - Different workload profiles
 - Dedicated servers
 - Unused Capacity, Administration
 - Consolidation
 - Resource management (Batch, Dialog, R/3, CRM)

Understand SAP infrastructure challenges



- Every SAP Servers do not require the same level of availability, performance, disaster protection & recovery (High to Low)
- Every SAP servers are not loaded equally over the periods (High to Low)
- Every SAP Servers are not relying equally on administration performance & quality (High to Low)
- SAP solution landscapes are generally not homogeneous in term of Infrastructure dependencies.
- Heterogeneous infrastructures are likely valuable in improving TCO of large SAP landscape
- zSeries platform (z/OS, DB2, z/VM, Linux) is the more flexible and open platform for “n tiered” solutions

z/OS benefits - SAP data base server



- **Reliability and Availability**
 - Most reliable database
 - Online backup, reorg
 - Failover capabilities in Parallel Sysplex
 - Disaster recovery/site failover with GDPS
- **Server and Business Consolidation**
 - Multiple components in one database (MCOB)
 - Multiple DB2 subsystems in one LPAR
- **Scalability**
 - 64-bit with DB2 V8 for z/OS
 - Data sharing in a Parallel Sysplex provides best scalability
- **Autonomic Computing**
 - System Automation on z/OS
 - GDPS

z/OS benefits - SAP application server



- High Availability, Automation, Integration, Consolidation
- Autonomic Computing
 - Automation of entire SAP solution - across platforms
 - driven by Tivoli System Automation
- Recommended for SAP Central Services
 - Central Instance / Standalone Enqueue Server
- Web application server also used for ...
 - Business integration with IMS, CICS, WebSphere, ...
 - Host printing

However:

SAP applications limited to 2 GB (31-bit)

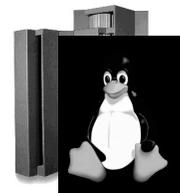
No support of SAP J2EE applications

No Unicode support

Linux for zSeries benefits



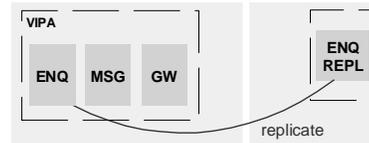
- 64-bit SAP application server on zSeries
- Server and business consolidation
 - LPARs, z/VM
 - Lower TCO (total cost of ownership)
- Autonomic computing
 - Linux benefits from zSeries reliability and hardware characteristics
 - High availability with System Automation for Linux
- HiperSockets
 - Absolutely reliable and fast network between LPARs and VM guests
- Linux is IBM's and SAP's strategic platform
 - SAP's first 64-bit Linux platform has been zSeries
 - Components of NetWeaver developed by SAP on Linux for zSeries (because of 64-bit)
 - All new technologies will be available on Linux for zSeries (J2EE, Unicode, ...)
- Support contract with Linux provider or IBM SupportLine necessary



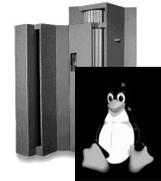
Achievements in 2002



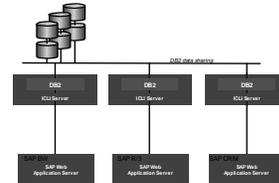
- SAP on Linux 64-bit for zSeries
 - Announced at CeBIT in 03/2002
 - GA in 05/2002



- High Availability solution based on Tivoli System Automation
 - Influenced SAP's architecture towards an open HA solution with z/OS and Linux
 - Enhanced IBM System Automation product
 - Deutsche Postbank in production since 10/2002



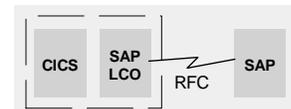
- Manageability Improvements for SAP on zSeries
 - Multiple SAP Components in One Database (MCOB)
 - Enhancements in SAP Control Center,
 - Real Time Statistics, monitoring, tuning, CC/390
 - z/OS Workload Manager further exploitation



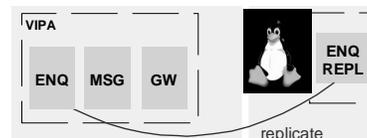
Achievements in 2003



- CICS - SAP connector
 - Requirement from banking customers
 - SAP LCO (Legacy connector) in production at one customer
 - Solution available as service offering by SAP



- SAP Java Connector
 - SAP JCo allows RFC/BAPI calls from/to Java or WebSphere programs
- Extension of HA/System Automation solution
 - Solution extended to Linux using Tivoli System Automation for Linux
 - Installation at first customer in progress



Recent deliveries



| Component | Release | SAP Technology | Shipment Date |
|----------------|-------------|----------------|---------------|
| Web AS | 6.20 | 6.20 | 7/2002 |
| BW | 3.0B | 6.20 | 8/2002 |
| R/3 Enterprise | 4.7 | 6.20 | 9/2002 |
| EBP / SRM | 3.5 / 2.0 | 6.20 | 10/2002 |
| CRM | 3.1 | 6.20 | 11/2002 |
| BW | 3.1C | 6.20 | 12/2002 |
| EBP | 3.5 SR1 | 6.20 | 02/2003 |
| R/3 Enterprise | 4.7 Ext 1.0 | 6.20 | 03/2003 |
| EBP / CRM | 4.0 | 6.20 | 06/2003 |
| SCM | 4.0 | 6.20 | 06/2003 |

- High Availability and Automation Solution on z/OS:
 - GA October 2002 (see SAP note 569996, IBM Redbook SG24-6836-00)
- High Availability and Automation Solution on Linux for zSeries:
 - First customer test ongoing

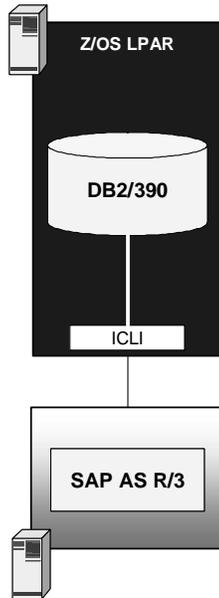
Release matrix



| SAP Technology | DB2 | OS/390, z/OS | Linux on zSeries | AIX | Windows | Solaris |
|----------------|------------|---------------|------------------------------|------------------|-----------------------------|----------------|
| 4.6D | 6.1 7.1 | 2.10+ | SLES 7 SLES 8 (64 bit) | 4.3.2+ | NT 4.0 Win2000 + | 8+ (64 bit) |
| 6.20 | 6.1 7.1 | z/OS 1.2 + | SLES 7 SLES 8 | 5.1+ (64 bit) | Win2000 + | 8+ (64 bit) |
| 6.30 | 7.1 | z/OS 1.4 + | SLES 8 (64 bit) | 5.1+ (64 bit) | Win2000 + | 8+ (64 bit) |
| 6.xx | 8.1 | z/OS 1.4 | SLES 8 (64 bit) | 5.2 (64 bit) | Win2003 (32 + 64 bit) | 9 (64 bit) |

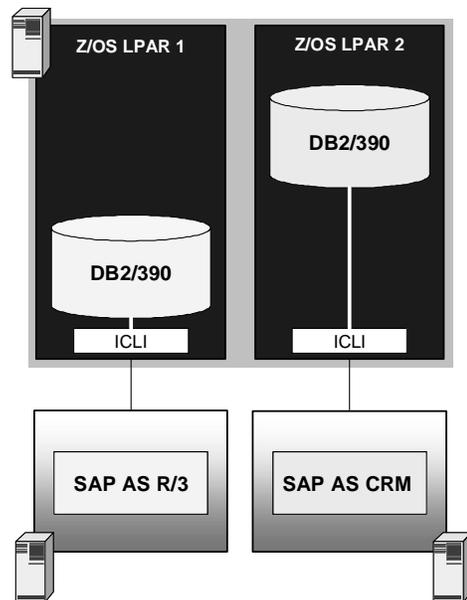
- Min. S/390 hardware: G5, minimum for Linux: zSeries z800, z900, z990
- Min. S/390 hardware for DB2 V8: zSeries z800, z900, z990
- 6.30/6.xx J2EE: JDK 1.4.1 on z/OS and Linux; JDK 1.3.1 else

SAP DB Server on zSeries



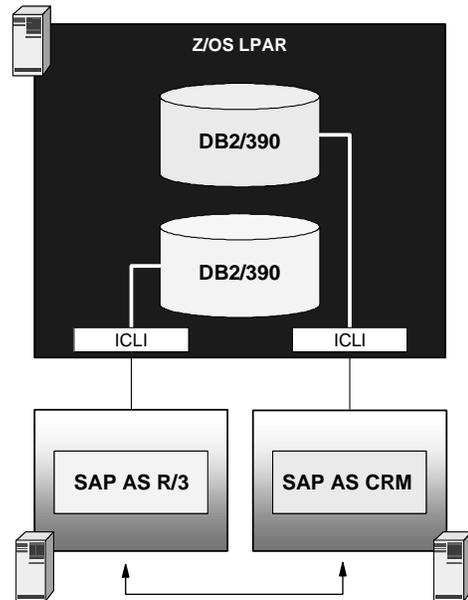
- Availability 1997 for SAP R/3
- DB2/390 only (5.1 to 7.1)
- BW, APO, CRM... all releases
- External SAP AS on Windows, AIX, Solaris, USS and Linux for zSeries.
- DB2 for z/OS outstanding strengths
- Single system availability (99,9 %)
- On Line Backup / Reorganisation
- DB2 for z/OS Hardware Data compression

SAP DB Server on zSeries



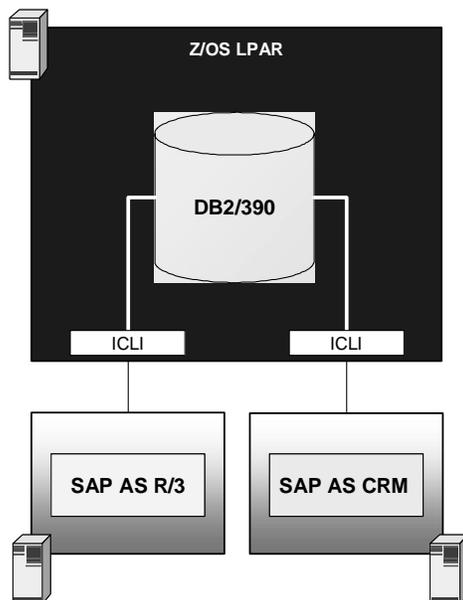
- 1 System = n SAP DB servers
- Hardware Resources Sharing
 - CPU,
 - I/O,
 - memory
 - WLM (scheduled performance goals according to business needs)
- Reduced operating cost
- LPAR Isolation

SAP DB Server within one LPAR



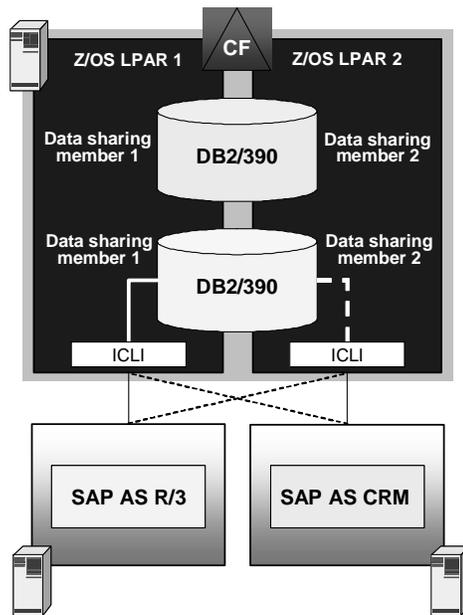
- 1 LPAR = n SAP DB servers
- No memory constraints with 64 bits real memory (zSeries)
- Hardware Resources Sharing
 - CPU,
 - I/O,
 - memory
 - WLM (scheduled performance goals according business needs)
- z/OS application Isolation
- Reduced operating cost
- Problem to keep data consistency across SAP components (... MCOD)

MCOD implementation on zSeries



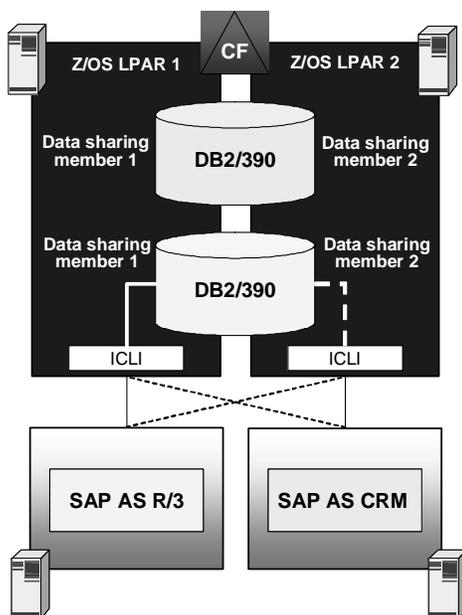
- MCOD (Multiple Components in One Database)
- Mini : SAP R/3 4.6C, CRM 2.0C, BW 3.0A
- 1 SAP DB for several SAP systems
- Separate Table sets
- Reduced disk space
- Reduced operating costs
- Synchronized Backup & Recovery
- Up to 50 (20 recommended) components in 1 MCOD landscape

SAP DB server HA with DB2 data sharing



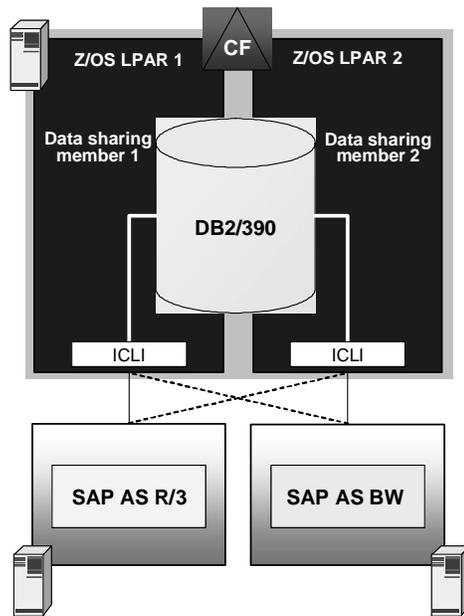
- Each SAP DB server run over > 1 LPAR within 1 system
- High availability for each SAP DB server (99,999 %)
- Each member run independantly
- Resources available to runing members (WLM / IRD)
- SAP AS auto reconnect to running DB member
- Automatic restart of failed members through z/OS automation
- Availability during DB2 or z/OS Upgrade

SAP DB server HA and scalability



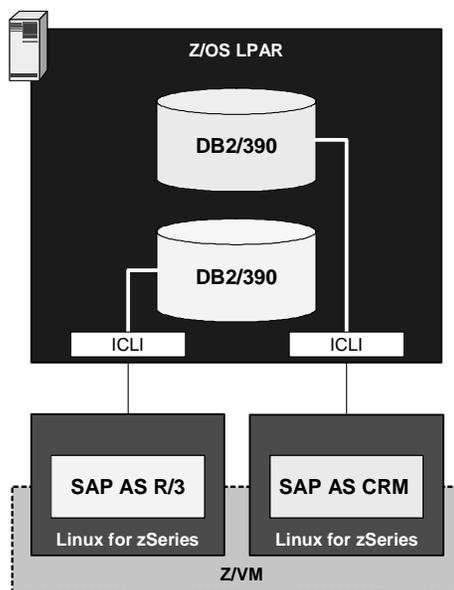
- Each SAP DB server run over several (< 32) systems
- Proved Linear Scalability
- // access to storage devices
- Very high Performance for storage intensive activities (Batches)
- High troughput capabilities for BackUp / Recovery
- Geographic dispersion of data sharing members (40 kms between CPU's)
- Service can be continued when one center is out of order
- High business continuity solution combined with data mirroring PPRC, XRC, ...)
- Many automation capabilities when supported by solution like GDPS.

MCOD implementation in Sysplex



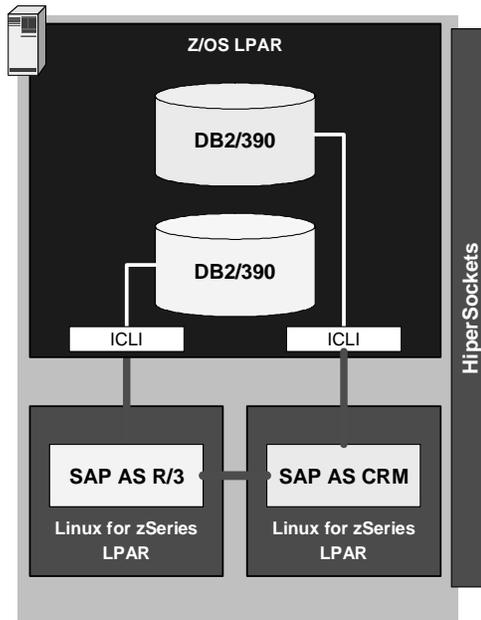
- OLTP & OLAP systems combination within a single MCOD
- Each DB members can have its own set of parameters and tuning
- A given data Sharing member can be primarily optimized for a given SAP system
- A given data sharing member is able to alternately serve any SAP system AS

SAP AS on Linux for zSeries



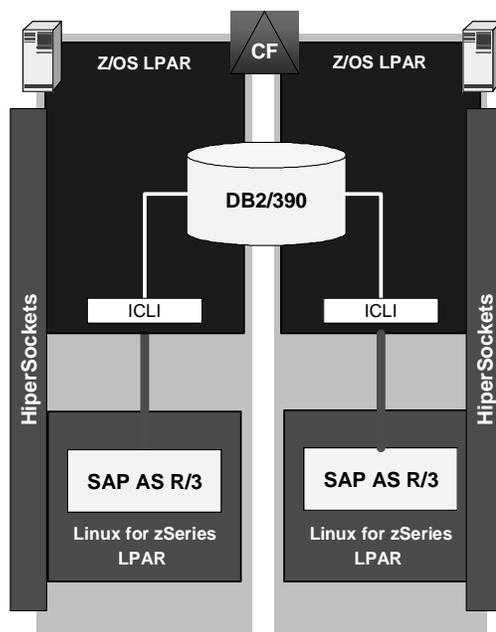
- AS on a native Linux or a Linux under z/VM
- SAP DB with DB2 for z/OS
- Use of ICLI (under USS + ICLI Client on Linux)
- Regular engine or IFL
- Availability May 2002 :
 - 4.6D SAP Kernel or higher
 - 4.6B SAP R/3 or higher
 - DB2/390 6.1 or higher
 - OS/390 2.8 or higher for DB server
 - z/VM 4.2 or higher
 - Linux 64 bits distribution with glibc 2.2.4 and kernel 2.4.x or higher (SuSE ES 7)
 - 64-bit z/Architecture available for AS

Linux consolidation



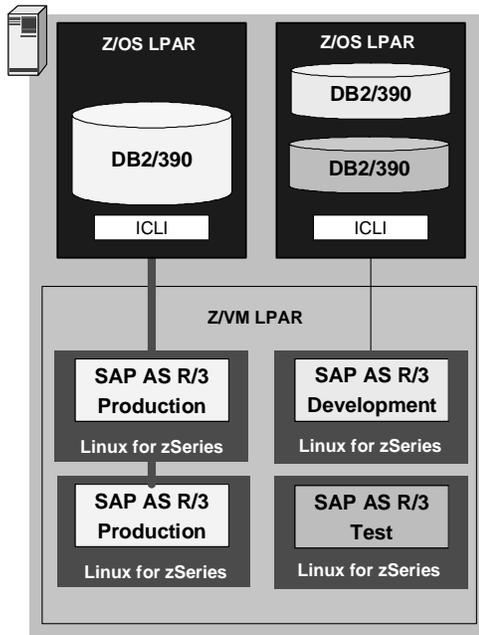
- SAP DB & AS Servers within 1 box
- SAP DB on z/OS or OS/390 LPAR and SAP AS on native Linux LPAR's
- LPAR's connected through Hipersockets (z/OS only)
- High throughput for I/O intensive activities between AS and DB (Batches, Updates)
- High throughput for data exchange between SAP systems
- Hardware Resources Sharing (CPU, I/O) with IRD (no mix regular engines / IFL)
- Reduced operating cost (connectivity)
- Network security inside zSeries
- LPAR Isolation

Linux consolidation



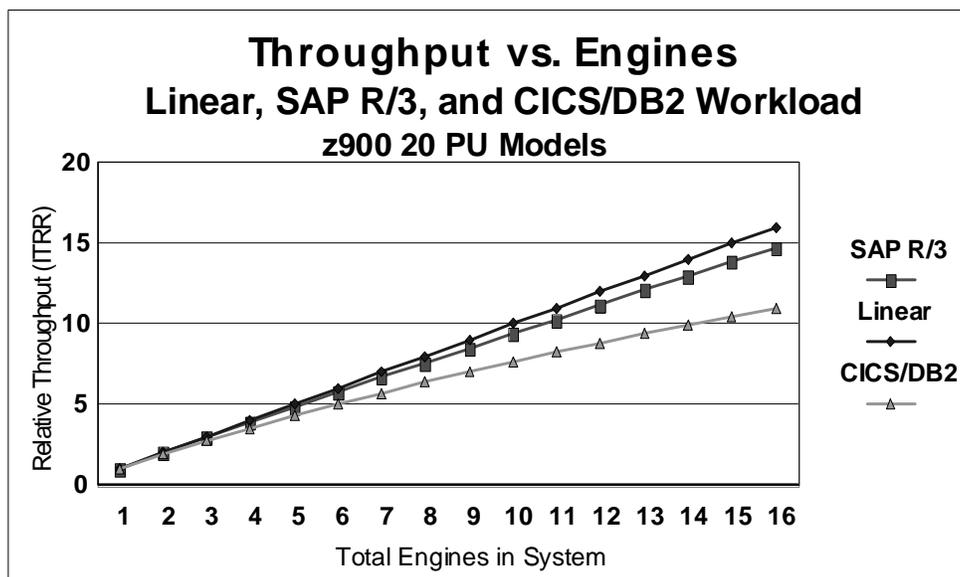
- 1 SAP system (DB + AS) over > 1 (< 32) systems
- High throughput performance between DB & AS within each box (Hipersockets)
- High DB performance with data sharing and ICF (for Batch work processes)

Linux consolidation



- Several Linux guests with SAP AS Servers within 1 z/VM LPAR
- z/VM not recommended for production systems
 - Production under z/VM requires preliminary analysis
- Reduced operation cost (central point of control)
- Quick guest and SAP AS creation / delete

SAP R/3 scalability on zSeries



SAP R/3 has a good scalability

Upcoming platform highlights



- DB2 V8 for z/OS exploitation
- SAP WebAS > 6.30 (SAP NetWeaver)
 - Requiring DB2 V8 with DB2 Connect V8, replacing ICLI
 - SAP NetWeaver: new component for SAP business solutions integration (Portal,)
- SAP WebAS on z/OS
 - Exploitation of Enhanced ASCII feature with file tagging
 - XPLINK provides performance gain

Introduction of DB2 Connect (DRDA)

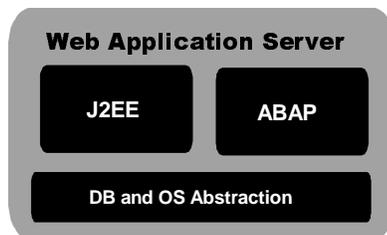


- DB2 Connect replacing ICLI
- Reasons and benefits
 - DB2 Connect
 - Supports Unicode, JDBC, XA
 - Is the IBM standard product for remote database connection
 - Is widely spread and tested
 - Is also used by SAP for DB2 on UNIX / Windows / Linux
- SAP Web AS 6.30 requires DB2 JCC
 - DB2 JCC (J2C)
 - Is the common Java database connector (JDBC)
 - Is packaged with DB2Connect
 - Consists of platform independent Java code
- SAP Web AS > 6.30
 - Requires installation of DB2 V8 and DB2 Connect as ICLI replacement

SAP NetWeaver requirements



- Platform requirements
 - Unicode
 - Supported by database
 - Supported by db connectivity
 - JDBC
 - Supported by db connectivity
 - 64-bit addressability for SAP components, middleware
- SAP WebAS 6.30
 - ABAP stack 6.20 patch level
 - new J2EE stack
 - requires JDK 1.3.1 or higher
- SAP WebAS >6.30
 - ABAP and J2EE stack updated
 - requires JDK 1.4.1



- Consequences on zSeries
 - DB2 V8
 - z/OS 1.4
 - Linux on zSeries SLES 8
 - DB2 Connect V8
 - JDK 1.4.1

Summary



- Large configuration choices available on zSeries
- Support of all SAP solutions
 - SAP Enterprise Portals and
 - Exchange Infrastructure together with SAP NetWeaver™
- Automation/HA solution utilizing enqueue replication
 - Driven by Tivoli System Automation
 - General available on z/OS
 - First customer on Linux for zSeries
- Remote DB2 connectivity switched to standard IBM product
 - ICLL to be replaced by DB2 Connect (DRDA)
- Full compliance with the SAP NetWeaver™ strategy
 - DB2 V8, Unicode, JAVA, ...
 - Linux as strategic application server platform

Linux on zSeries for Oracle Solutions



IBM/Oracle business relationship



- The global business and technology relationship between IBM and Oracle leverages the products and services of both companies to provide mutual customers with a unique value
- Proposition :
 - Oracle is a primary provider of database, tools and application software on UNIX platforms
 - IBM is a primary provider of Technology platforms.
 - IBM and Oracle are committed to support and delight their mutual customers.
- The critical measure of success: Customer Satisfaction

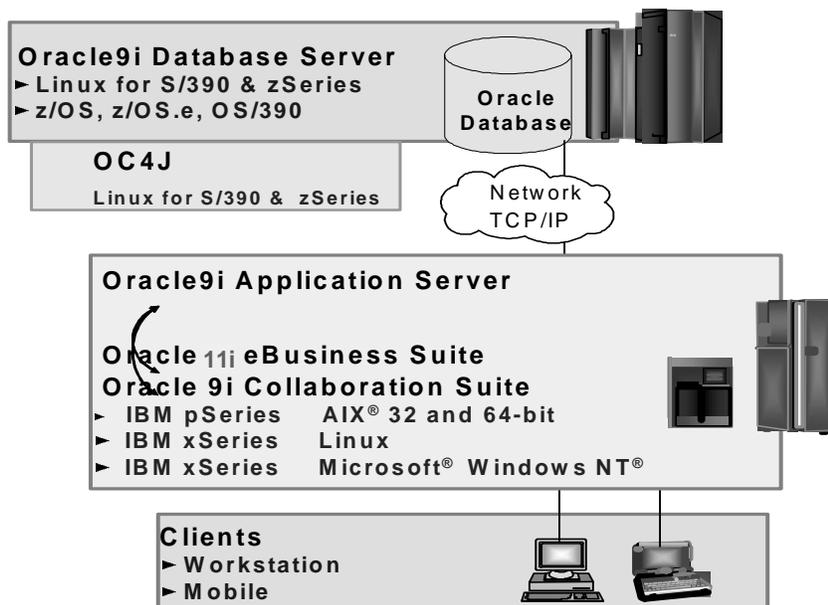
EMEA IBM / Oracle JSC



EMEA IBM/Oracle Joint Solutions Center
(JSC) (Opened in 1998)



Oracle9i architecture for zSeries



History of Oracle9i for Linux for S/390



- August 2001 - Oracle announced plans
 - Developer's Release of Oracle9i Database for Linux/390
- November 15, 2001
 - Oracle9i Database for Linux/390 (31-bit)
 - Download available on Oracle Technology Network
<http://otn.oracle.com>
- January 15, 2002 – New downloads available on OTN
 - Oracle9i Database for Linux/390 (31-bit) (update)
 - Oracle9i Database for z/Linux (64-bit)
- March 18, 2002
 - Oracle announced Oracle9i Database for Linux/390
 - Availability was announced as initially planned for summer 2002

History of Oracle9i for Linux for S/390



- August 28, 2002
 - Product available for download for early adopters
- September 2002
 - Sept. 13: Product available from store.oracle.com
 - Sept. 17/18: Seminar for early adopters held at Oracle HQ
- January 24, 2003
 - Oracle9i Release 2 - 9.2.0.2 patchset available
- May 2003
 - Oracle9i Release 2 - 9.2.0.3 patchset available
- September 2003
 - Oracle9i Release 2 - 9.2.0.4 patchset available

Early customers interest



- Statistics from Oracle Technology Network
 - 48,386+ Unique visitors to Linux/390 page on OTN
 - 1,960+ Customers/prospects filled out the initial survey
 - 441 New members on OTN
 - 1,557 (79%) want a DB Server on Linux S/390
 - 833 (43%) want an Application Server on Linux S/390
 - 367 (19%) want ERP/CRM/CSM Applications on Linux S/390
- Developer's Release Downloads
 - 1,914 Oracle9i Database for Linux S/390 (31-bit)
 - 1,273+ Oracle9i Database for Linux zSeries (64-bit)
- Production Release Downloads
 - 5,243+ Oracle9i Database for Linux S/390 (31-bit)

What is Oracle9i database for Linux?



- Oracle9i Release 2 Database
 - Same as any other major UNIX port
 - Includes Real Application Clusters (RAC) feature
- Oracle9i Release 2 Database for Linux for S/390 (31-bit)
 - 31-bit database implementation
 - Short name: Oracle9i Database for Linux/390
 - Built on SuSE (SLES7) for IBM S/390 (31-bit) August 2002
 - Certified on SuSE SLES8 for IBM S/390 (31-bit) May 2003
 - 9.2.0.2 requires patch 2964102
 - 9.2.0.3 requires patch 2757568
 - 9.2.0.4
 - Runs on IBM S/390 (31-bit) and zSeries (64-bit) hardware

What is Oracle9i database for Linux?



- Oracle9i Release 1 Database (Developer's Release) for z/Linux (64-bit)
 - 64-bit database implementation
 - Short name: Oracle9i Database for z/Linux
 - Built on SuSE (SLES7) 64-bit for IBM zSeries 1/15/2002
 - Runs on IBM zSeries (64-bit) hardware
- Oracle Transparent Gateway for DRDA 9.0.1 and 9.2
 - Certified for SuSE SLES7 and server 8.1.7 – 9.0.1 and 9.2

Oracle DB for Linux for S/390 pre-requisites



- Oracle9i Database Enterprise Edition Release 2 for Linux for S/390 (31-bit)
 - SuSE Linux Enterprise Server 7 (SLES7) with patch CD1
 - 31-bit only
 - Patches for raw devices support with RAC
 - devs-2001.10.11-8.s390.rpm and rawio-2-33.s390.rpm
 - SuSE Linux Enterprise Server 8 (SLES8) with Service Pack 2
 - 31-bit only
 - RedHat 7.2 for Linux for S/390
 - 31-bit only
 - Patches for raw devices support with RAC
 - devs-2001.10.11-8.s390.rpm and rawio-2-33.s390.rpm
 - Disk space: 4800 cyl.
 - Oracle installation requires xWindow interface for Auto Installer
 - IBM JDK 1.3.1 (Java Development Kit)

Oracle Database support matrix



| Distribution | Kernel | Mode | Oracle 9i & 9i RAC 31-bit | Oracle 9i & 9iRAC 64-bit | Oracle 10g & 10gRAC 64-bit |
|----------------------|--------|--------|---------------------------|--------------------------|----------------------------|
| SuSE SLES7 (1) | 2.4.7 | 31-bit | Certified | - | - |
| SuSE SLES7 (1) | 2.4.17 | 64-bit | - | - | - |
| SuSE SLES8 (2) | 2.4.19 | 31-bit | Certified | - | - |
| SuSE SLES8 (2) | 2.4.19 | 64-bit | Planned | - | Planned |
| RedHat 7.1 | 2.4.9 | 64-bit | - | - | - |
| RedHat 7.2 | 2.4.9 | 31-bit | Certified | - | - |
| RedHat AS 3.0 (Beta) | 2.4.21 | 31-bit | - | - | - |
| RedHat AS 3.0 (Beta) | 2.4.21 | 64-bit | - | - | Planned |

(1) SuSE SLES7 Patch CD1 needed

(2) SuSE SLES8 Service Pack 2 needed - gcc 3.2 is standard on SLES8
Oracle9i certified with gcc 2.95 provided with SLES8 SP2

Oracle9i database key features



- Availability/Scalability
 - Real Application Clusters (RAC)
- Availability/disaster recovery
 - Data Guard - Logical or physical standby database
- XML database features
- Information integration (Streams)
 - Event capture/propagate/apply
 - Advanced Queuing, Advanced Replication
- Manageability
 - Enterprise Manager (OEM agent on Linux S/390)
- Security
 - LDAP support

Oracle9i database features availability



- MicroFocus COBOL compiler available
 - Pro*COBOL early release in customer test (Aug. 2003)
 - G/A with 9.2.0.4 patchset Sept. 2003
 - Only supported on SLES8 (31-bit)
- Oracle Enterprise Manager
 - Intelligent Agent for Linux S/390
 - All functions except Management Server supported
 - All functions with OEM 4.0
- INSO filters unavailable
 - Oracle Text/Ultra Search affected, can use remotely
- VeriSign, Entrust, CyberTrust not available
 - Authentication using other methods
- No Generic Connectivity
 - OLE, ODBC gateways

Backup / recovery and management



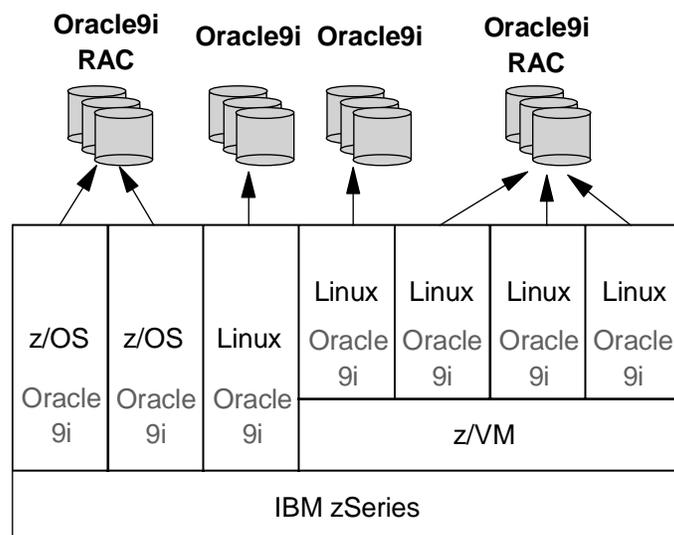
- Backup and Recovery
 - Recovery Manager (RMAN) provides extensive backup and restore/recovery tools
- Media management tools
 - Tivoli available (Tivoli Data Protector)
 - Legato Storage Manager - Available
 - Veritas Netbackup - Available
 - BMC SQL BackTrack - Available
 - FDR Upstream - Available
- Database Administration and Performance Tuning
 - Oracle Enterprise Manager (OEM)
 - DB Patrol

Oracle9i scalability



- Vertical – scale large processor configuration (SMP)
 - Increase LPAR weight
 - Increase number of real CPUs or IFLs
 - Multiple databases on single server
 - Single database on single server
- Horizontal – scale large number of virtual servers
 - Add new instances of Linux – a penguin colony
 - Multiple distinct databases across multiple OS images
 - Oracle9i Real Application Clusters (RAC)

Oracle9i deployment options on zSeries

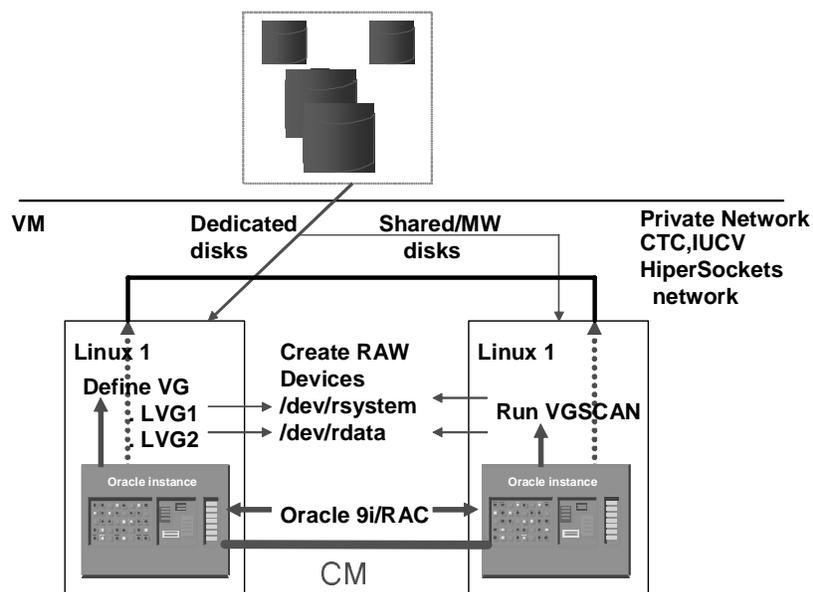


Oracle9i Real Application Clusters



- Oracle Real Application Clusters (RAC)
 - Increases availability and scalability
 - Uses raw device for database file sharing (patches needed)
 - Uses Oracle Cluster Manager
 - Single database across multiple virtual servers
 - LPARs or virtual system images
 - Single database across multiple physical servers

Oracle9i RAC for Linux S/390 & zSeries



J2EE applications for Linux zSeries



- Running a full java application on zSeries
 - 9i / 9iRAC for the database tier
 - OC4J engine for the application tier
 - Container for Servlets, JSPs and EJBs
 - IBM JDK 1.3.1 or 1.4.1
- Tools to develop and deploy the J2EE applications
 - Oracle Jdeveloper (9iDS)
 - Not available on Linux for zSeries
- Benefits of the Solution:
 - Developing, deploying full J2EE java applications connected to an Oracle database, and fully running on Linux zSeries platform

Early experience



- Oracle9i Database: same product across platforms: Oracle is Oracle
- Oracle9i Database for Linux/390
 - Built on SLES 7 (SuSE Linux Enterprise Server 7)
 - Work on Red Hat 7.2 (31-bit) 2.4.9-17 for S/390
 - Work on TurboLinux 7.0 alpha code for S/390
- Good quality and very stable
- Good performance
 - Oracle9i Database for Linux/390 performance comparable to Oracle9i Database for Linux/Intel
 - Based on IBM, Oracle and customer tests

Oracle products directions



- Oracle10G Database for z/Linux
 - 64-bit implementation only
 - Built on SuSE SLES8 for IBM zSeries (64-bit)
 - Certify on Red Hat Advanced Server 3.0 for IBM zSeries (64-bit)
- Split-Tier certifications Planned
 - E-Business Suite 11i
 - Oracle Collaboration Suite Release 2 (9.0.4)
 - Oracle Application Server
- Native ports (full implementation) Not in plan
- Oracle E-Business Suite 11i
 - Oracle Collaboration Suite
 - Oracle9i Application Server
 - Oracle9i R2 Database for Linux for zSeries (64-bit)

Oracle10g highlights



- Oracle 10g is ready for enterprise Grid Computing
 - Low cost, modular storage and servers
 - Highest Quality of Service
 - Capacity on demand
 - Load balancing
 - Share information throughout the Grid
 - Oracle Database 10g can federate all data sources
 - Access remote information wherever it is located
 - Web services consumer
- Oracle 10g delivers self-managing databases
 - Built-in Intelligent Infrastructure
 - Automatic Database Diagnostic Monitor
 - Automatic performance diagnostic engine in the database
 - Automatic Tuning Optimizer
 - SQL Profile to tune packaged applications

Why Oracle9i on Linux for zSeries?



- ✓ Total Cost of Ownership (TCO)
 - ✓ Server Consolidation Strategy
- ✓ Skills and Staffing
- ✓ Reliability, Availability and Serviceability
- ✓ Rapid deployment of Linux servers
 - ✓ Time to market competitive advantage
- ✓ Ease of management of servers
 - ✓ Cloning “Virtual Blades”
- ✓ Scalability
 - ✓ Ease of horizontal and vertical growth
- ✓ Portability and ease of migration

Oracle reference material



- Oracle Technology Network
 - All Oracle9i Generic Linux documentation
 - All Oracle9i Linux on Intel documentation
 - Oracle9i for Linux/390 Release notes
- White Paper
 - Oracle for Linux/390 Overview
- Oracle Linux Forum
 - <http://otn.oracle.com/tech/linux/content.html>
- Oracle for Linux/390 Information and Download
 - <http://www.oracle.com/go/?&Src=915044&Act=13>

Oracle9i installation process

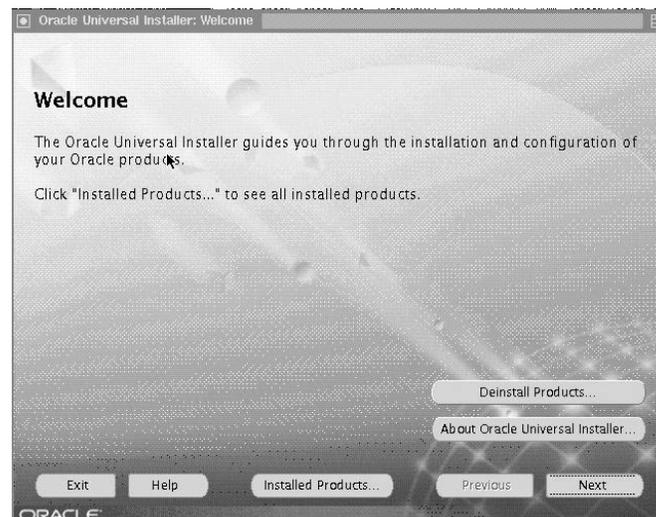


- Oracle Universal Installer requires an xWindows interface:
 - VNC, eXceed (Hummingbird)
 - AIX, Linux or NT client
- Installation
 - Documentation in several Oracle Linux books
 - Set up environment
 - Run the Auto installer : `./runinstaller`
 - Start the Oracle Cluster Manager on all nodes:
 - `$ export ORACLE_HOME=/oracle/product/9.2.0`
 - `$ORACLE_HOME/oracm/bin/ocmstart.sh`
 - Choose RAC option

Oracle9i installation



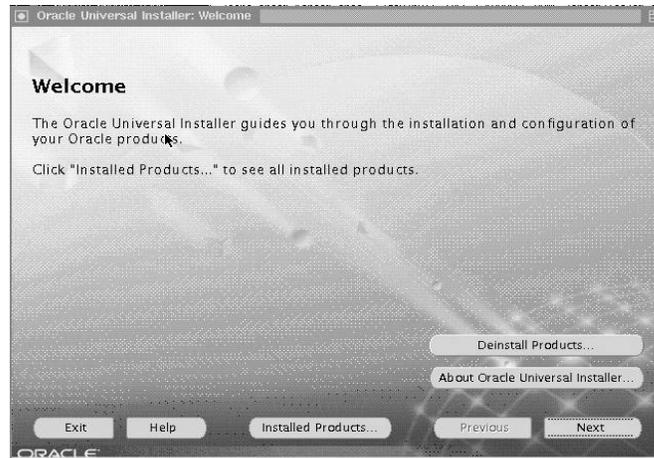
Installing Oracle9i on Linux for S/390



Oracle9i installation



RAC installation



Things to consider when moving a db



- There are several items to consider when moving Oracle
 - This will most likely be a migration
 - Transforming Oracle7 to Oracle8/8i and then 9i
 - Or and upgrade
 - Upgrading is moving (in place) from Oracle8/8i to 9i
- So, do you
 - Upgrade/migrate first?
 - Or do upgrade/migrate as part of the move?
- Keep in mind
 - There is no magic here
 - There aren't any really cool tricks - its moving a database
 - We don't know more than you or your DBAs

Before you start



- Memory concerns
 - Remember, this is a **31bit** operating system
 - It gives you 2GB of memory space
 - Oracle max SGA is 1GB (with changes to kernel parms)
 - SGA: Systems Global Area
 - This is not a performance concern unless you need a large SGA (Oracle memory footprint)
 - 31-bit and 64-bit Oracle servers have the same "on-disk" formats for database files
 - Any compiled objects like PL/SQL or nested tables will have different binaries between 31-bit, and 64-bit servers
 - Physical database size determines
 - Storage needs
 - Time and method to move the database

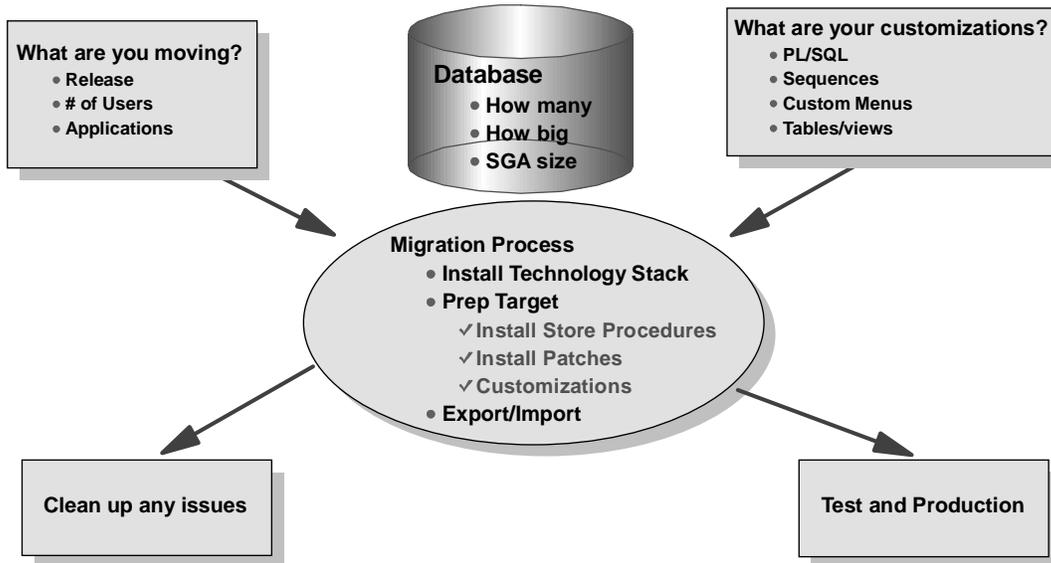


Short history on problems

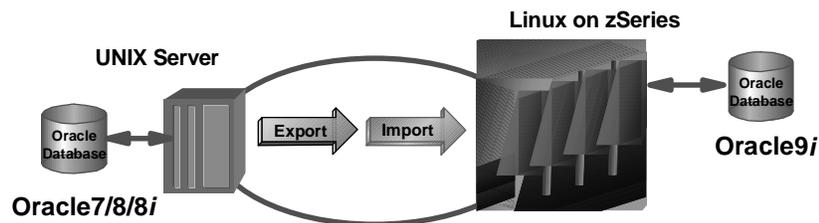


- genksms - allows 1GB SGA
 - Oracle patch needed (9.2.0.2)
- opatch - problem applying patches
 - Fixed at level 14
- Sharing Oracle binaries
 - Good idea but some problems
 - No recipe
 - No Oracle support
- COBOL compiler
 - Requires MicroFocus with Pro*COBOL
 - Available in SLES8

OK, how do I move it?



Ways to move the data

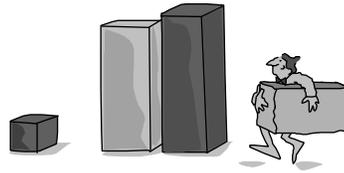


- More Choices
 - Upgrade as part of FULL Export / Import
 - Use export level of current DB level and import with 9i
 - SYS schema issues may occur upgrading levels
 - Or upgrade in place and then do full export
 - Make sure you have the capacity you need to upgrade
 - Or Unload / Load data using SQL*Loader
 - Or CTAS (create table as)

How long does a move take?



- It depends
 - Import is a single threaded job
 - What is the model of the zSeries?
- On a z900
 - About 3 GBytes/hour per CPU
 - z990 should be better but not yet tested
- On G5/G6 or z800
 - About 1 to 1.5 GBytes/hr per CPU
- Keep in mind that
 - Work in other images will decrease this number
 - Run multiple imports if possible
- These are estimates, your mileage may vary!!!



Application compatibility issues



- Applications written for Oracle7, Oracle8 and Oracle8*i* do not need to be rewritten
 - Unless they take advantage of new 9*i* function
- Oracle Forms written in Oracle7, Oracle8 and Oracle8*i* and Oracle9*i* run the same
 - Check to see if using new function would help improve the application
- Other things to check
 - Check code for new reserved words in 9*i*
 - Clients should be at SQL*Net V2 or Net8
 - Check issues for OCI
 - No need to precompile, compile apps if not using new function
 - Don't forget server manager went away
 - New commands for SQL*Plus

Database functional issues



- The 'compatible' parameter in init.ora
 - Sets compatibility with previous levels
 - Controls the database function and settings (init.ora)
 - Minimum level is 8.1.0 for 9.2. If set to 8.1.0
 - Can not use new function in 9.2 that will make the database incompatible
 - Can not use database setting that are new in 9.2
 - All database functions requiring 9.2 are listed in the Migration Guide

Summary for moving the database



- Plan the move carefully
- Select the best way
 - Migrate/upgrade then move
 - Migrate/upgrade as part of the move (Export/Import)
 - Load/unload data
- Plan storage and processor needs
- Practice, practice practice
- Test, test, test

Linux on zSeries sizing



- The Linux on zSeries is a two fold strategy
 - Proximity. Get the applications as close to the data as possible.
 - Move application servers that access legacy data into Linux on zSeries and access through HiperSockets (zSeries)
 - Server consolidation
 - Moving low utilization, non peaking servers to Linux on zSeries

- For Oracle that means
 - Large databases from 64-bit UNIX are not a good fit
 - Moderately sized production databases and development and test environments are good fits.
 - The virtualization benefit can be realized

Sizing vs. capacity plan

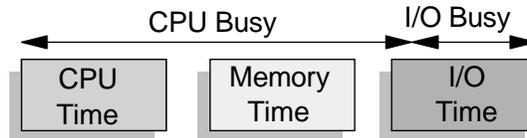


- Sizing is an estimate only
 - Has to be based on some similar test data results
 - Data is available for server consolidation
 - Understand the relationship of moving workloads from OS/390, Linux/Intel, UNIX to Linux on zSeries
 - IBM can offer sizing assistance for Server Consolidation
- 
- Capacity Planning
 - Ongoing process
 - Based on actual application with real users or pilot
 - Historical data tracked
 - This allows for reasonable extrapolation of future needs

Comparative performance



- Processor Design - Four components of capacity
 - CPU
 - Memory Hierarchy
 - I/O Structure
 - Management



- Processor designs vary greatly in the balance of capacity across components which results in a wide range of relative capacity.
- Processor, Memory and I/O time vary greatly with application

What you need to provide



- Performance Engineers in Gaithersburg need
 - Server profile
 - Make model and number of CPUs
 - Memory
 - Utilization profile
 - Average utilization and peaks
 - VMSTATS is a good thing - its better than guessing
 - Application characteristics
 - OLTP, BI, etc..
 - Number of users
- You don't need to remember this!
 - They have a questionnaire for you to fill out



That's all folks!

Thank you!

