

Why WebSphere Application Server on System z

Bill Jones
wgjones@us.ibm.com



© 2008 IBM Corporation

IBM Systems & Technology Group
IBM Software Group

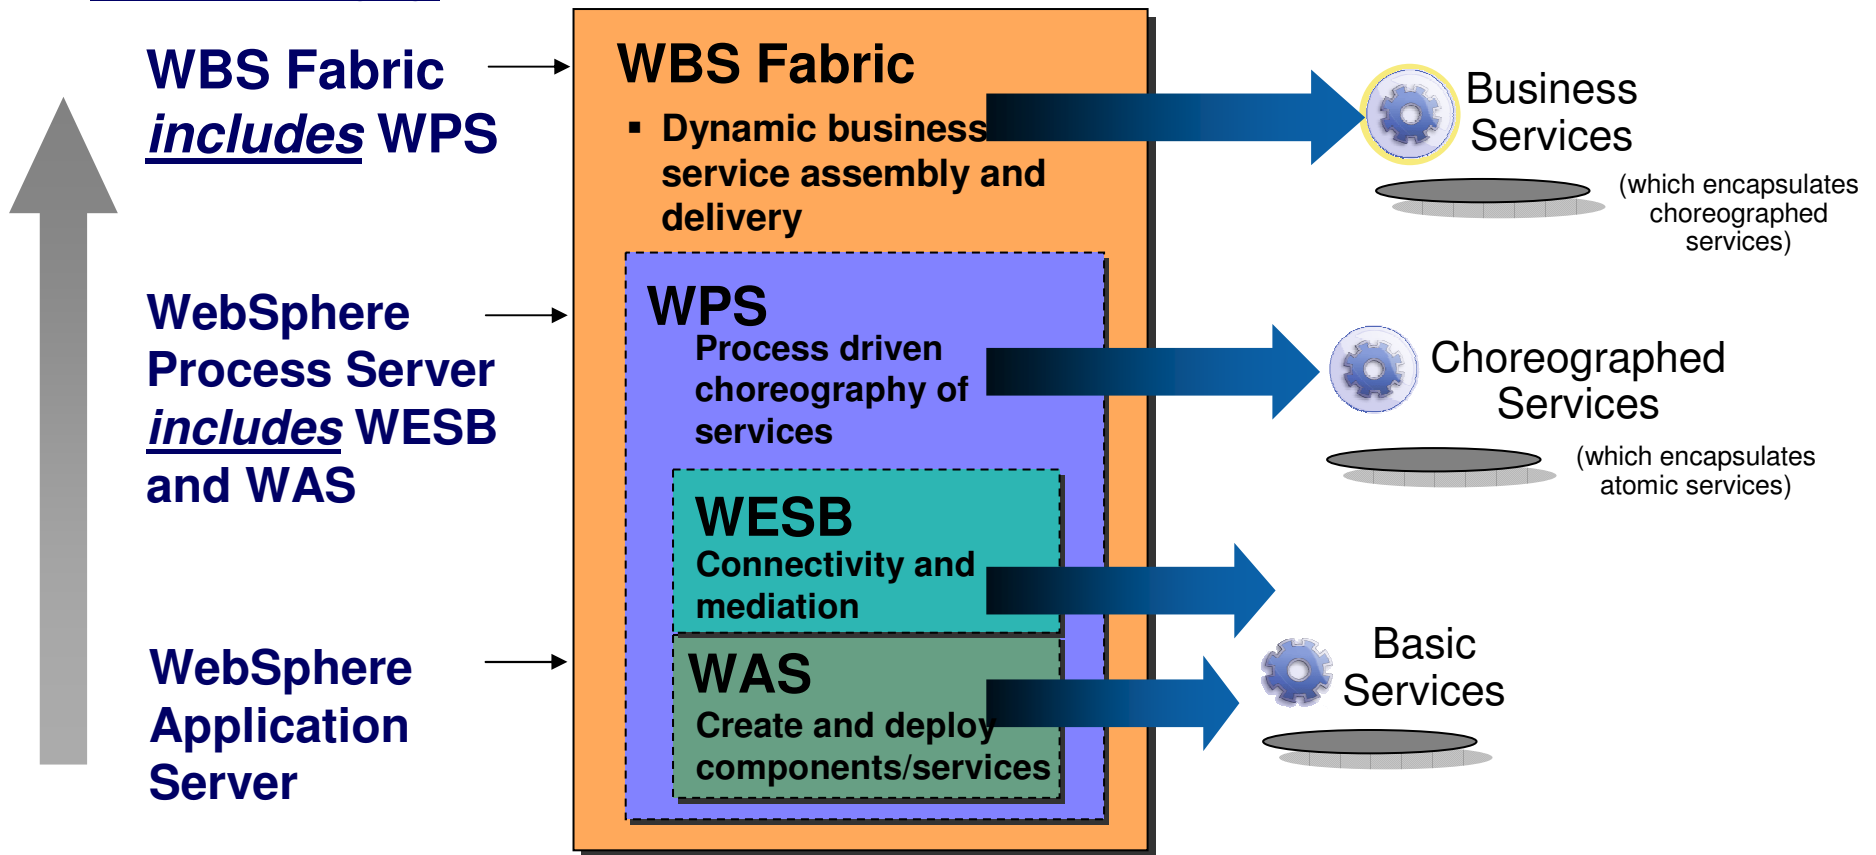
Topics for today's discussion

- **The WebSphere Application Server**
- **Why WAS on z/OS**
- **WAS on Linux for System z**
- **WebSphere Extended Deployment Overview**



A Complete WebSphere Business Services Solution Stack with a cumulative product packaging approach

Product Packaging:



SOA requires a robust, secure deployment environment



IBM WebSphere Application Server V6

The Industry's leading application server for building, running & managing business-critical application services

- 1 **Build** and **deploy** J2EE applications quickly and with ease, focusing on business logic while shielding applications from platform specifics
- 2 **Run** services in the most secure, scalable, highly available environment
- 3 **Reuse & create** Java assets and extend their reach to existing applications and data
- 4 **Manage** applications effortlessly
- 5 **Grow** as needs evolve, leveraging core skills and assets

System z is the fastest growing WebSphere platform for 2005



charles SCHWAB



WebSphere Application Server v6.1

Powering your SOA for the Ultimate in Business Flexibility



Standards Based and Open

- J2SE 5.0
- JDK 5.0
- Web services standards
 - WS-Interop Basic Security
 - WS-Notification
 - WS-Business Activity
 - WS-Resource Framework
- JSR168 Portlets
- JSR116 Session Initiation Protocol (SIP) Servlets

Consumability

- Application Server Toolkit, including automation tools and Command Assistance
- z Profile Management Tool
- Simplified Administration
- Simplified SSL Key/Certificate Management
- Security enhancements
- Virtual Member Manager
- IHS administration enhancements
- Integrated Support Assistant
- Platform Messaging Enhancements

Platform Capability

- Performance Enhancements
- Proxy Server Enhancements
- Integrated User Registry
- Government Standards



WebSphere Application Server V6.1

Concurrent delivery on zOS and Distributed

WAS v6.1 Feature Packs

Bringing State-of-the-Art Technology To Market Quickly!



Web Services

- 2Q07 GA, all platforms
- Delivery of Reliable Asynchronous Messaging Profile as optional component for WAS 6.1
- Provides secure, reliable business process integration with those of customers and suppliers allowing interactions to span long durations
- Java and SCA programming models
- Community Centric Profiles delivered in WAS 7.0

EJB 3.0

- 4Q07 GA, all platforms
- Provides support for EJB 3.0 and Java Persistence API (JPA) on WebSphere Application Server 6.1
- Aimed at simplifying EJB development
 - Annotations provide component metadata in code
 - No need to create XML EJB deployment descriptors
- JPA is a simple and powerful persistence framework
 - Relational databases represented using “plain old Java objects” (POJOs)
- EJB 3.0 provides a standards based POJO alternative to frameworks

Future WAS Feature Packs

Bringing State-of-the-Art Technology To Market Quickly!



Web 2.0

- 4Q07
- Web 2.0 to SOA Connectivity – For enabling connectivity from Ajax (Asynchronous JavaScript and XML) clients and mash-ups to external web services, internal SOA services, and JEE assets. Extends enterprise data to customers and partners through web feeds.
- Ajax Messaging – For connecting Ajax clients to real-time updated data like stock quotes or instant messaging.
- Ajax Development Toolkit – Best-in-class Ajax development toolkit for WebSphere Application Server based on Dojo (dojotoolkit.org) with IBM extensions.

SOA

- 2H08 GA, all platforms
- Apache Tuscany code base on WAS release after 6.1
- SCA
 - POJO (Java Object) service component implementations including support for annotations
 - Intra and Inter-composite asynchronous capability
 - Recursive composition model support
 - SCA bindings include the Web Services binding, SCA default binding, JMS Binding, and EJB binding
 - HTTPSession Scoped components
 - Sample applications
 - Exploitation of Web Services feature pack functionality
- SDO
 - Dynamic data API support
 - Static data API support, including code generation capabilities
 - Some Helper classes implemented (XMLHelper, XSDHelper, DataFactory, CopyHelper, EqualityHelper, DataHelper, TypeHelper)
 - Partial ChangeSummary support
 - Java serialization of DataObject
 - SDO metadata configuration support
- SCA and SDO samples
- Whitepapers, BLOGs – to detail IBM's plans for SCA and SDO within an SOA environment
- Includes SOA Core from WPS 6.0 (for interoperability)

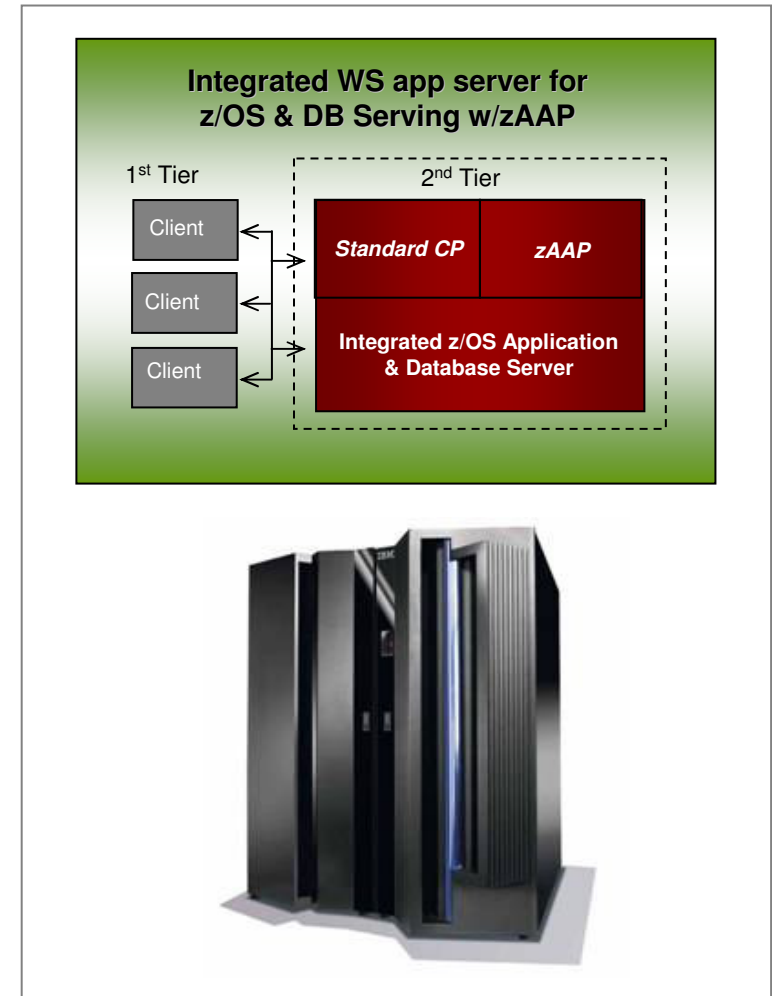
What's New for WebSphere Application Server for z/OS

■ Reduced cost of ownership

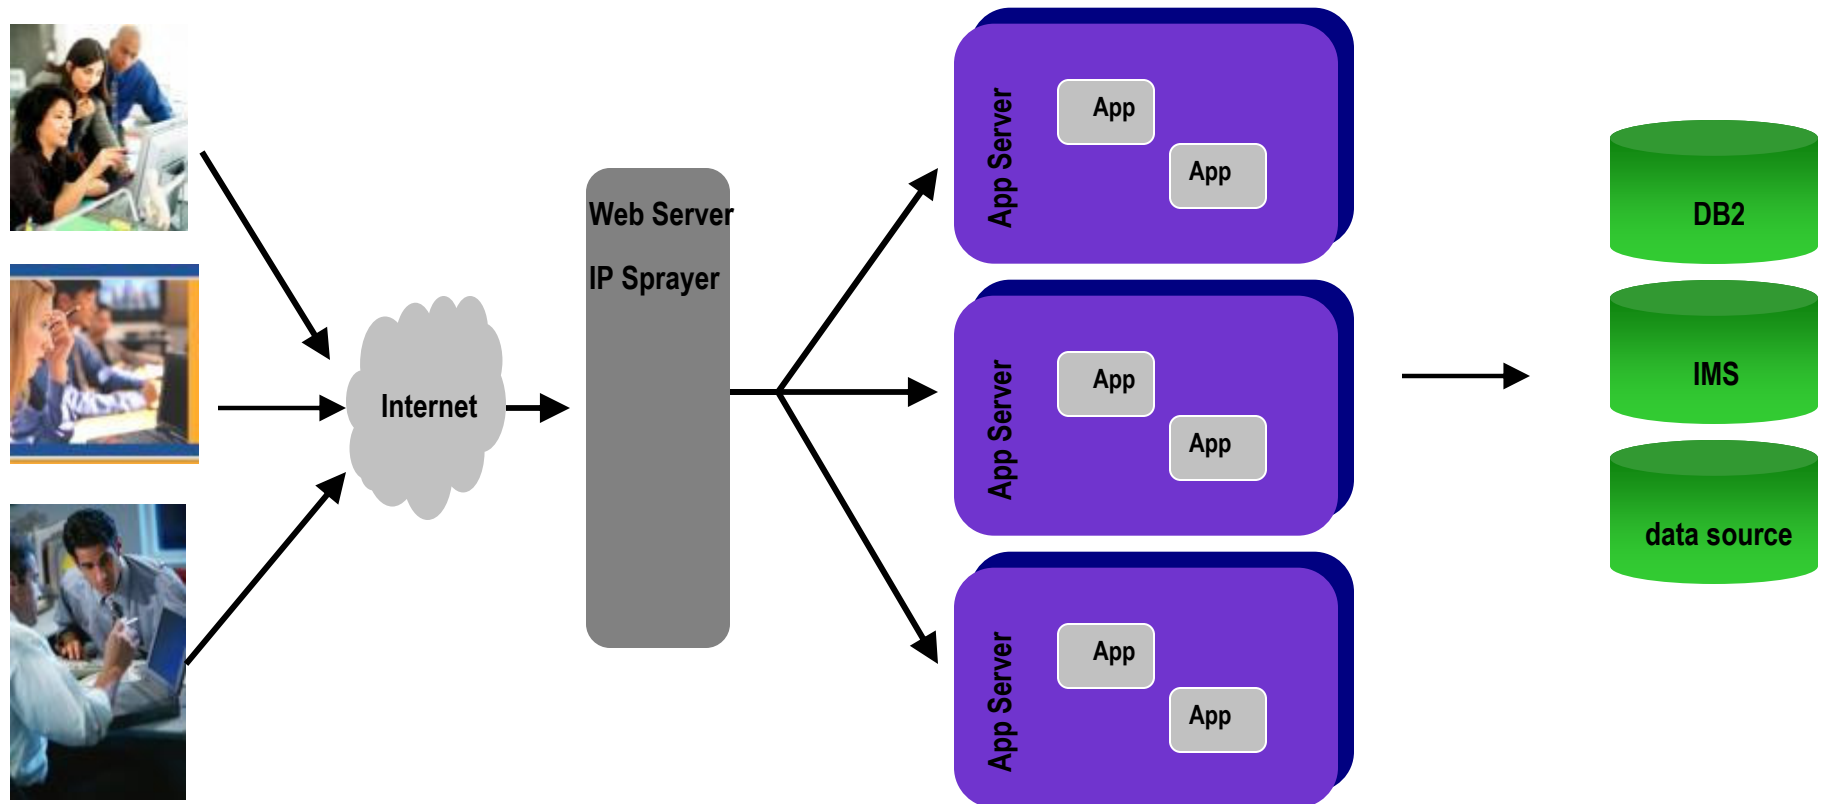
- Architected on SOA infrastructure & principles
 - Fully J2EE 1.4 platform certified
 - Leading Web Services support
 - WebSphere Rapid Development & Deployment
- zAAP enabled (System z9, z990, z890)
 - Run Java applications next to mission critical data
 - Lower the cost of computing for WebSphere Application Server (and all z/OS based Java applications)

■ Platform optimization / brand dedication

- Common code infrastructure
 - Administration skills shared between platforms
 - Develop anywhere, run on WebSphere Application Server for z/OS
 - Feature Packs for functional enhancements
 - Web Services, EJB 3.0, Web 2.0, SOA
- Native OS support; RRS, RACF, zWLM, Parallel Sysplex
- Optimization features designed to provide security and data interaction, including support for the traditional mainframe SW – CICS, IMS, DB2
- Enhanced QoS within the product, complementary to QoS of the platform
 - High availability manager – ARM
- Other “Hidden Gems” documented in White Paper (WP101138) on Techdocs at www.ibm.com/support/techdocs



A Typical Distributed Application Server Architecture



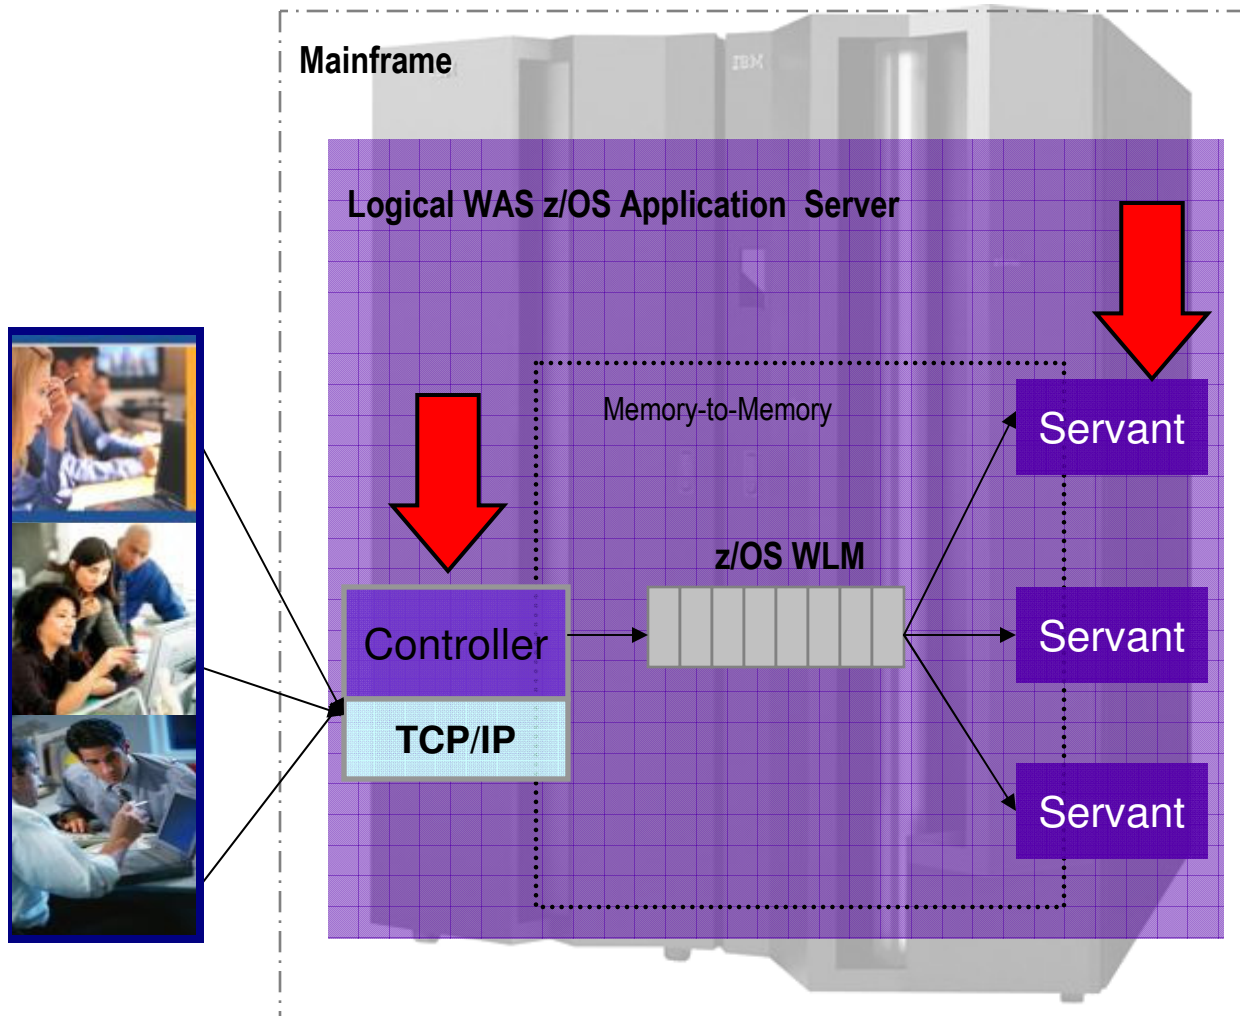
Typical Application Servers

- **How do you scale the Application?**
 - Add more boxes
- **How do you scale the Database?**
 - Add more boxes and replicate
- **How do you handle failover?**
 - Add more boxes
- **How is an application highly available?**
 - Install the application on many boxes
- **How are requests prioritized?**
 - Add more boxes to handle Gold Customers
- **How do I handle peak loads?**
 - Keep enough boxes available to handle the max possible load
- **What about when I don't have peak loads?**
 - The Servers will be underutilized
- **How do you route work evenly to multiple boxes?**
 - Round-robin inbound work
- **How do you ensure that your servers are evenly utilized?**
 - Best guess based on Application Knowledge

“Add More Boxes”

WebSphere Application Server for z/OS

Server Architecture



A WAS z/OS application server is comprised of multiple processes.

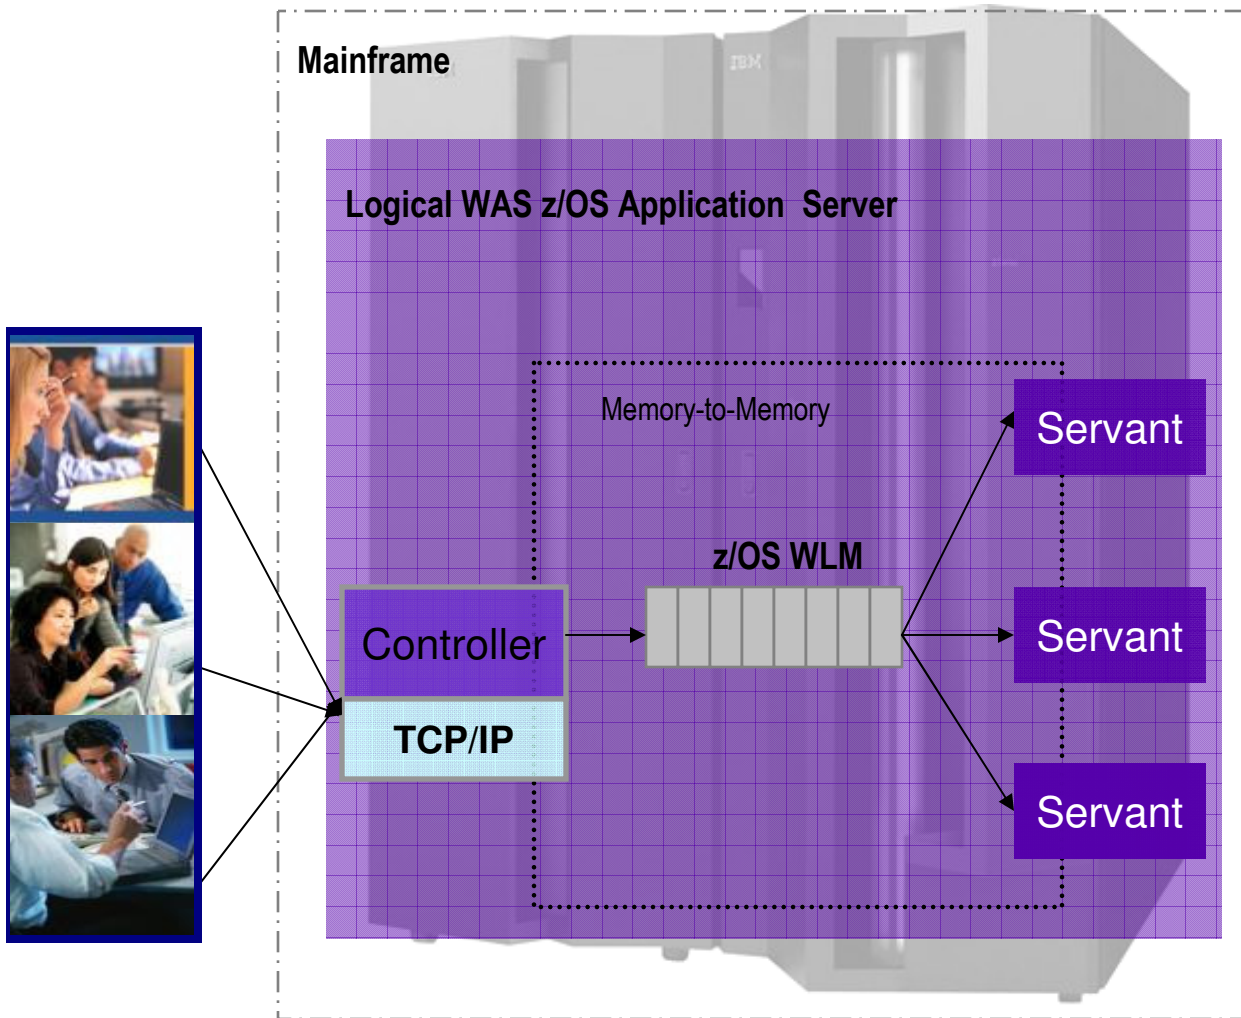
The Controller executes authorized, robust system code and acts as a control center for WAS

Application code is executed within the Servant Process.

- ✓ Theoretically, Servant processes are clones of each other.
- ✓ They can be terminated or created without interrupting the overall end-user experience
- ✓ Servant processes can be dynamically created or terminated by z/OS WLM (Workload Manager). The decisions can be based on user-defined service policies.

WebSphere Application Server for z/OS

Reliability

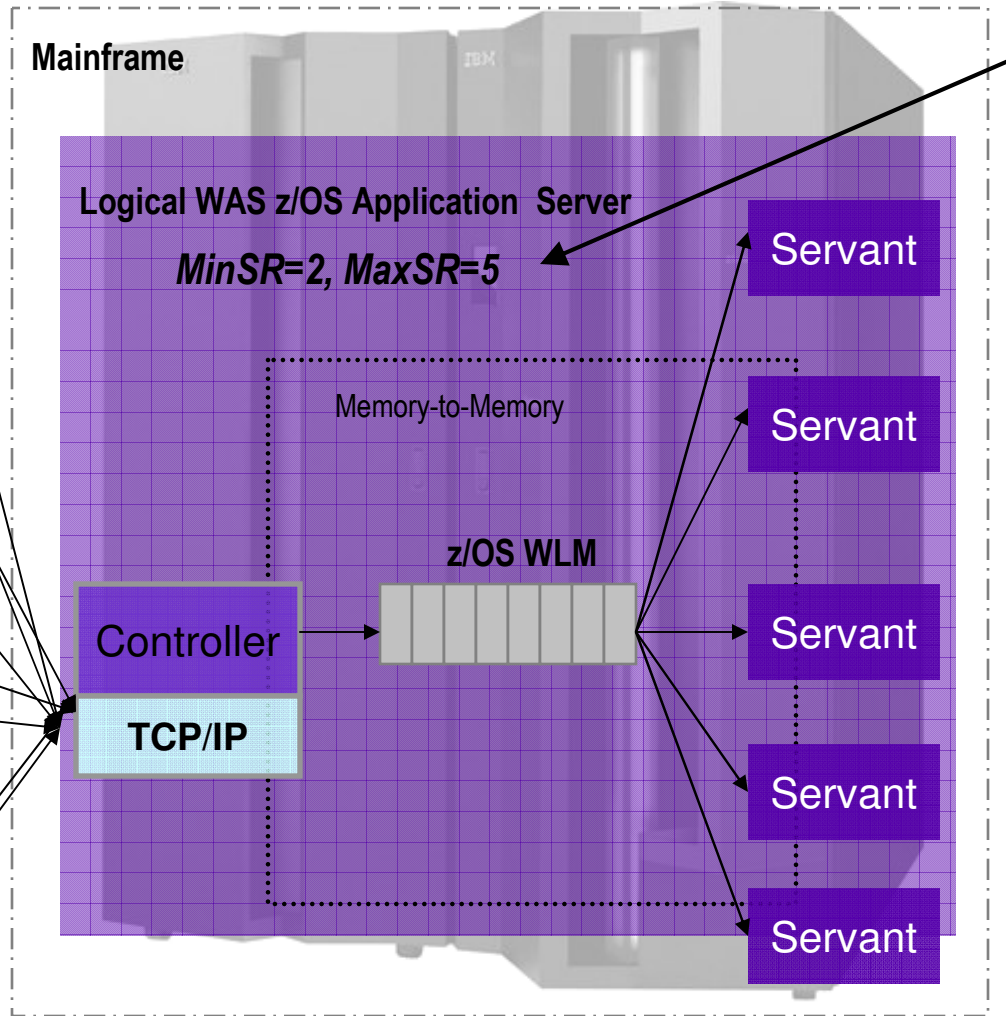


If a servant process executing application code is terminated, the application can still be served by any of the other servants within WAS

z/OS WLM detects the terminated servant and dynamically restarts it

WebSphere Application Server for z/OS

Scalability



WAS z/OS allows you to define a minimum and maximum number of servant processes

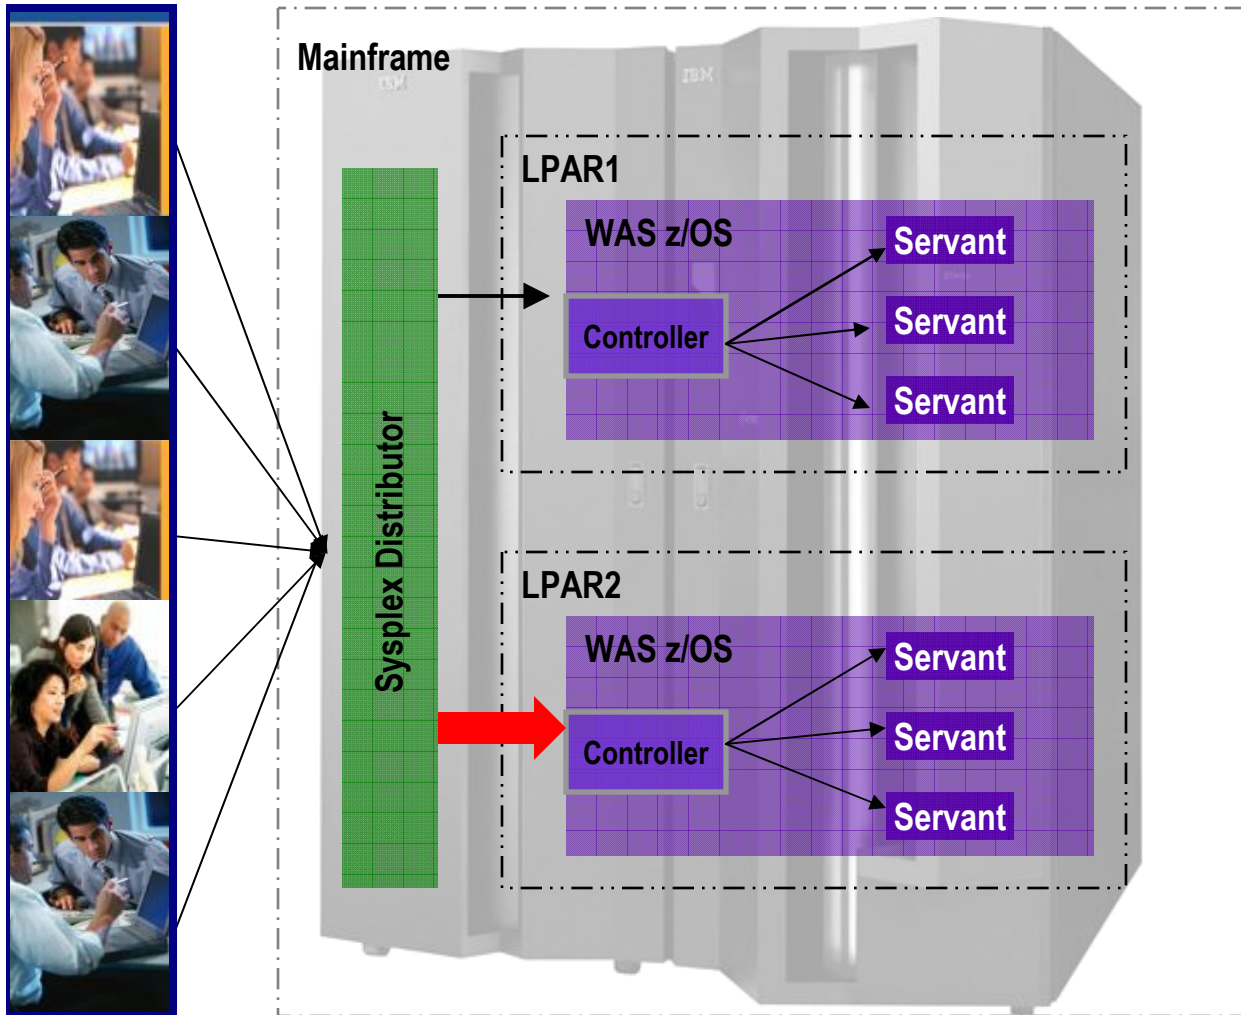
Workload can be unpredictable, and spikes in demand can demand application resources

When work begins to queue up waiting for application server resources, z/OS WLM takes notice and acts

Once the peak has subsided, z/OS WLM balances resources, so that it only maintains the servants needed to handle the incoming demand

WebSphere Application Server for z/OS

Availability



For true high availability, the Controllers need to be clustered

If a controller fails, the servants continue to process their in-flight work, but are not able to take additional work

The Sysplex Distributor reroutes work while **ARM** (Automatic Restart Manager) restarts the controller

All this is done without manual intervention

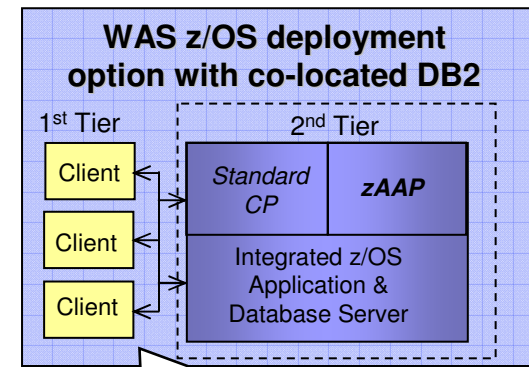
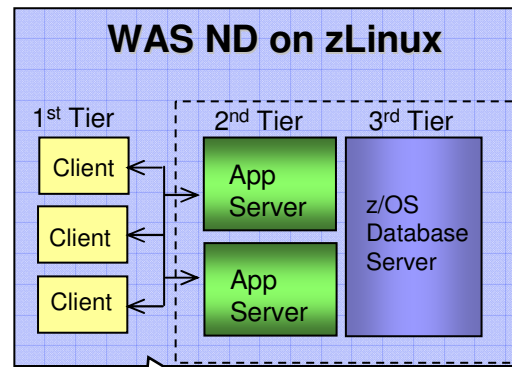

WebSphere z/OS Application Server

- **How do you scale the Application?**
 - Increase the Max Servants per Server
- **How do you scale the Database?**
 - Use a Shared DB2 Database
- **How do you handle failover?**
 - Use ARM for Controllers. Servants are handled via WLM
- **How is an application highly available?**
 - Servants are restartable via WLM. Controllers are recoverable via ARM
- **How are requests prioritized?**
 - WLM
- **How do I handle peak loads?**
 - WLM spawns servants
- **What about when I don't have peak loads?**
 - WLM eliminates excess Servants, freeing up system resources for other work
- **How do you route work evenly to multiple Servers?**
 - Sysplex Distributor routes work to specific LPAR's based on WLM metrics for resource utilization
- **How do you ensure that your servers are evenly utilized?**
 - WLM Metrics are used to decide where to route work to

WAS integration with z/OS System Facilities provides a more robust, efficient, scalable environment

Platform Matters

WebSphere Deployment options



Typical distributed deployment with:

- Network layer between each tier
- Normal access to z/OS DB limited by the network overhead and processor speed

Typical distributed deployment on the mainframe with:

- Ability to use hipersockets for faster DB transactions
- Normal DB access

Unique configuration only for WAS z/OS with:

- Memory to memory transfer rates with DB (Type 2 connections) for high volume transfer rates
- Remove network layer and overhead

Application Characteristics for Platform Deployments: System p and System z

Compute Intensive

Small working set

Raw throughput

Predictable

Homogenous

Simple transactions

Heavy I/O – mixed workload

Large working set


High ROI

Unpredictable volumes

Access to core mainframe assets


Complex transactions

Business critical




pSeries/System p

Considering application characteristics is only one factor when deciding where to run your business applications, but knowing what your platform does best is *critical*



zSeries/System z

What happens when logic and data are separated?

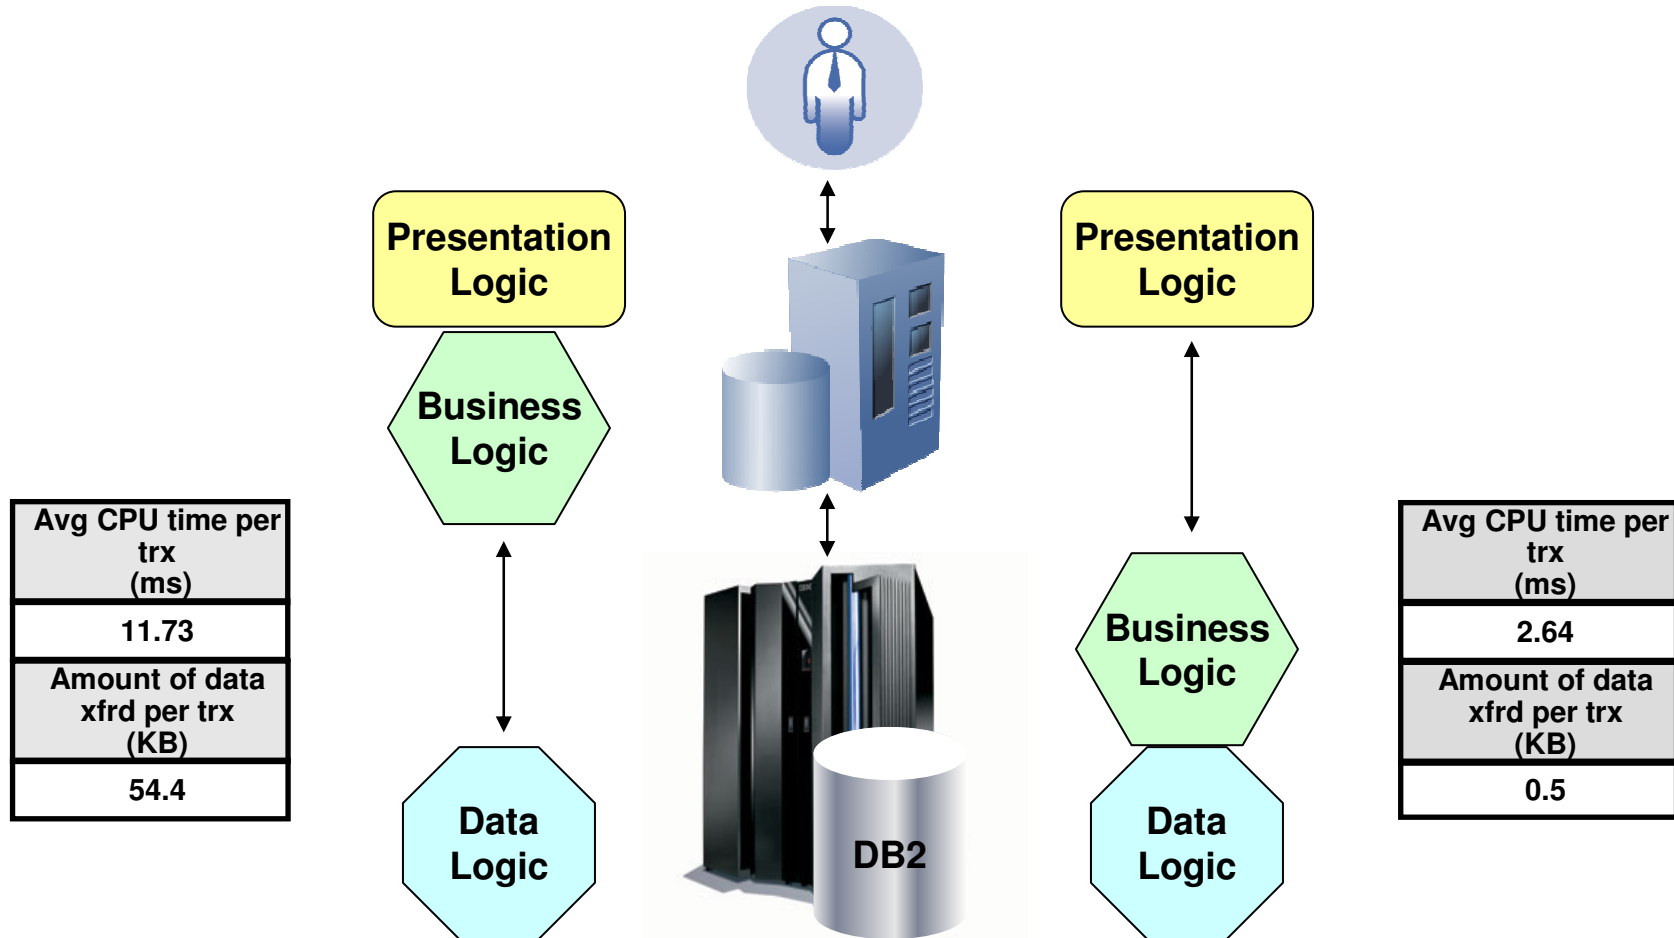


- Objects are converted into byte array at the requester (CPU, time)
- **Network latency is incurred (time)**
- More latency is incurred as service is dispatched (CPU, time)
- Objects are reconstructed at the server (CPU, time)
- Requested data is retrieved
- Objects are converted into byte array at the server (CPU, time)
- **Network latency is incurred (time)**
- Objects are reconstructed at the requester (CPU, time)

Some other considerations:

- Number of interactions between the tiers, volume of data passed
- No local optimizations of the access protocol
- Effect on server memory requirements due to locking

The value of proximity: transportation industry POC



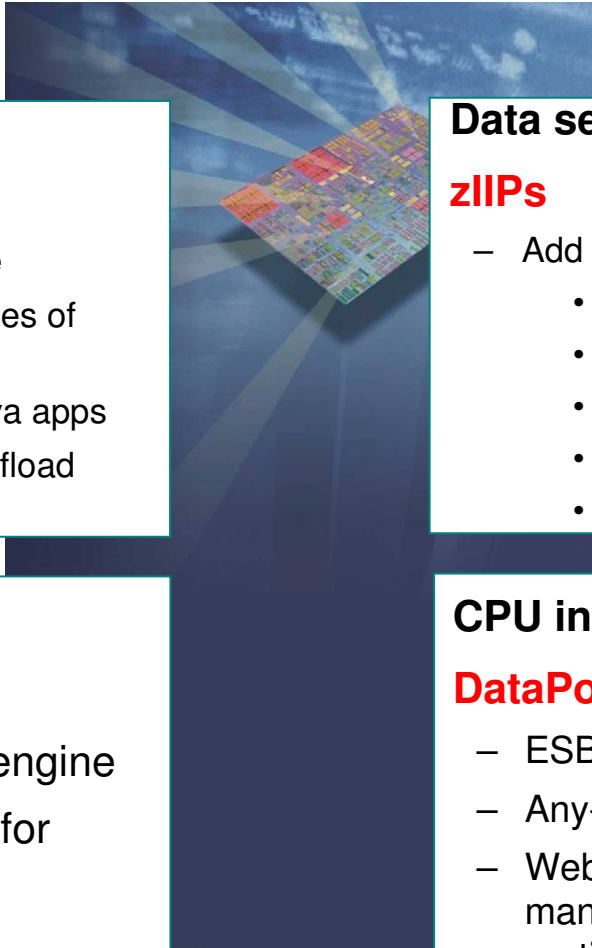
- **Effect of refactoring business logic to be co-resident with z/OS data:**

- Average CPU time per EJB transaction was reduced by over 77%
- Number of bytes of data transferred per EJB transaction was reduced by 99%

<http://www.ibm.com/support/techdocs>, Optimizing WebSphere Performance on DB2, WP100558

Specialized Hardware to Extend SOA Capabilities on System z

Customize your mainframe for your specific workload needs



Java on the mainframe

zAAPs

- Dedicated Java offload engine
- Average Java MIPs offload rates of 60%
- No development impact to Java apps
- z/OS XML System Services offload

Data serving on the mainframe

zIIPs

- Add value for DB2 workloads, including:
 - Enterprise applications
 - Data warehousing apps
 - DB2 z/OS V8 utilities
 - IPsec
 - z/OS XML System Services offload

Linux on the mainframe

IFLs

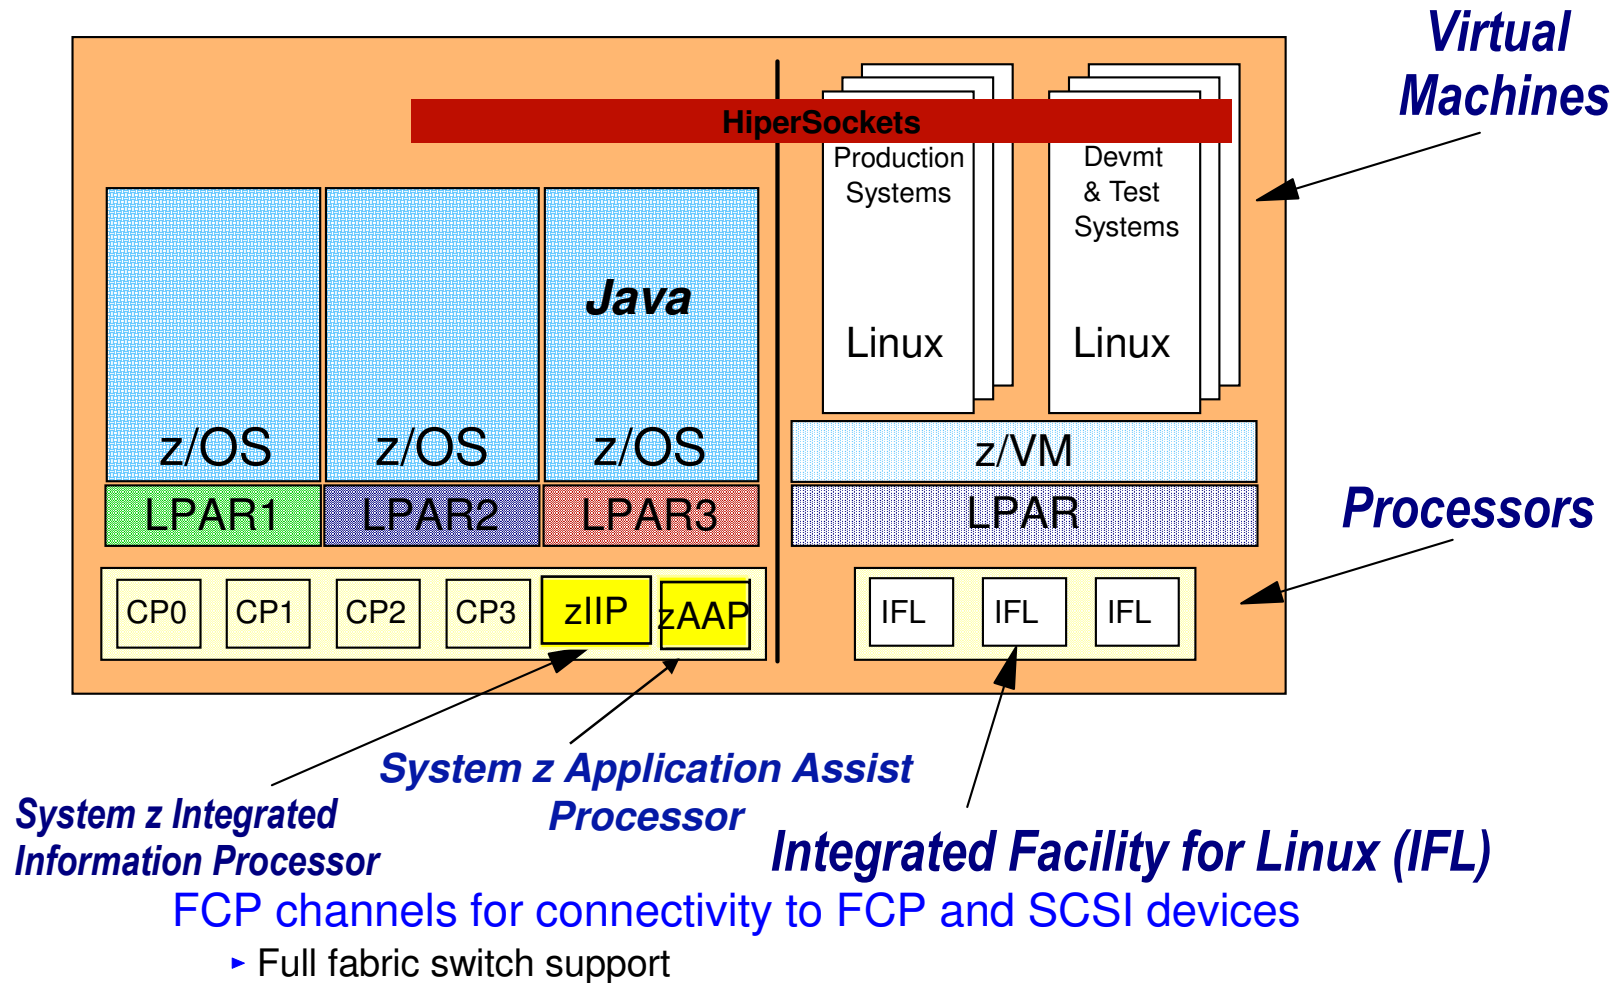
- Dedicated Linux offload engine
- Cost effective processor for Linux workloads
- LPAR mode specific

CPU intensive XML offload

DataPower

- ESB functionality in the hardware
- Any-to-any message transformation
- Web Service security and management, content-based routing, filtering

Sample IBM System z™ Configuration



WebSphere on System z Options

A self managing server environment with the versatility and power to help integrate your business

Linux deployment: Distributed Consolidation

- ✓ Large number of smaller apps from under utilized distributed servers
- ✓ Higher utilization than distributed servers
- ✓ Implement multi-tier applications in a single System z for better data proximity exploiting hipersockets
- ✓ Lower TCO with IFLs
- ✓ Speedy deployment – cloning/server provisioning
- ✓ Higher QoS than distributed
- ✓ Less stringent requirements than z/OS deployment
- ✓ Alignment with distributed WebSphere family
- ✓ Unrivaled virtualization with centralized management
- ✓ No z/OS Skills
- ✓ Web Serving infrastructure consolidation
- ✓ Presentation Services
- ✓ Flexible, virtualized Test/Migration/Prototyping Platform
- ✓ J2EE/ISV products not available on z/OS

Perfect for the System z customer requiring speedy deployment with less stringent QoS/integration requirements

z/OS deployment: Integration Option

- ✓ Highest QoS production environment
- ✓ Lower TCO with zAAPs
- ✓ Full exploitation of System z and z/OS
- ✓ Tight integration with DB2, CICS, IMS for chatty applications to eliminate network latency
- ✓ “Spikey”, unpredictable workloads
- ✓ Service level agreement management
- ✓ Dynamic load balancing, prioritization
- ✓ Strict security requirements
- ✓ Highest availability, reliability, scalability
- ✓ Disaster recovery and autonomic function
- ✓ Dynamic I/O configuration
- ✓ Storage management
- ✓ Capability/tools to modernize and integrate existing System z applications
- ✓ Migrate applications from another platform that require additional scalability and integration

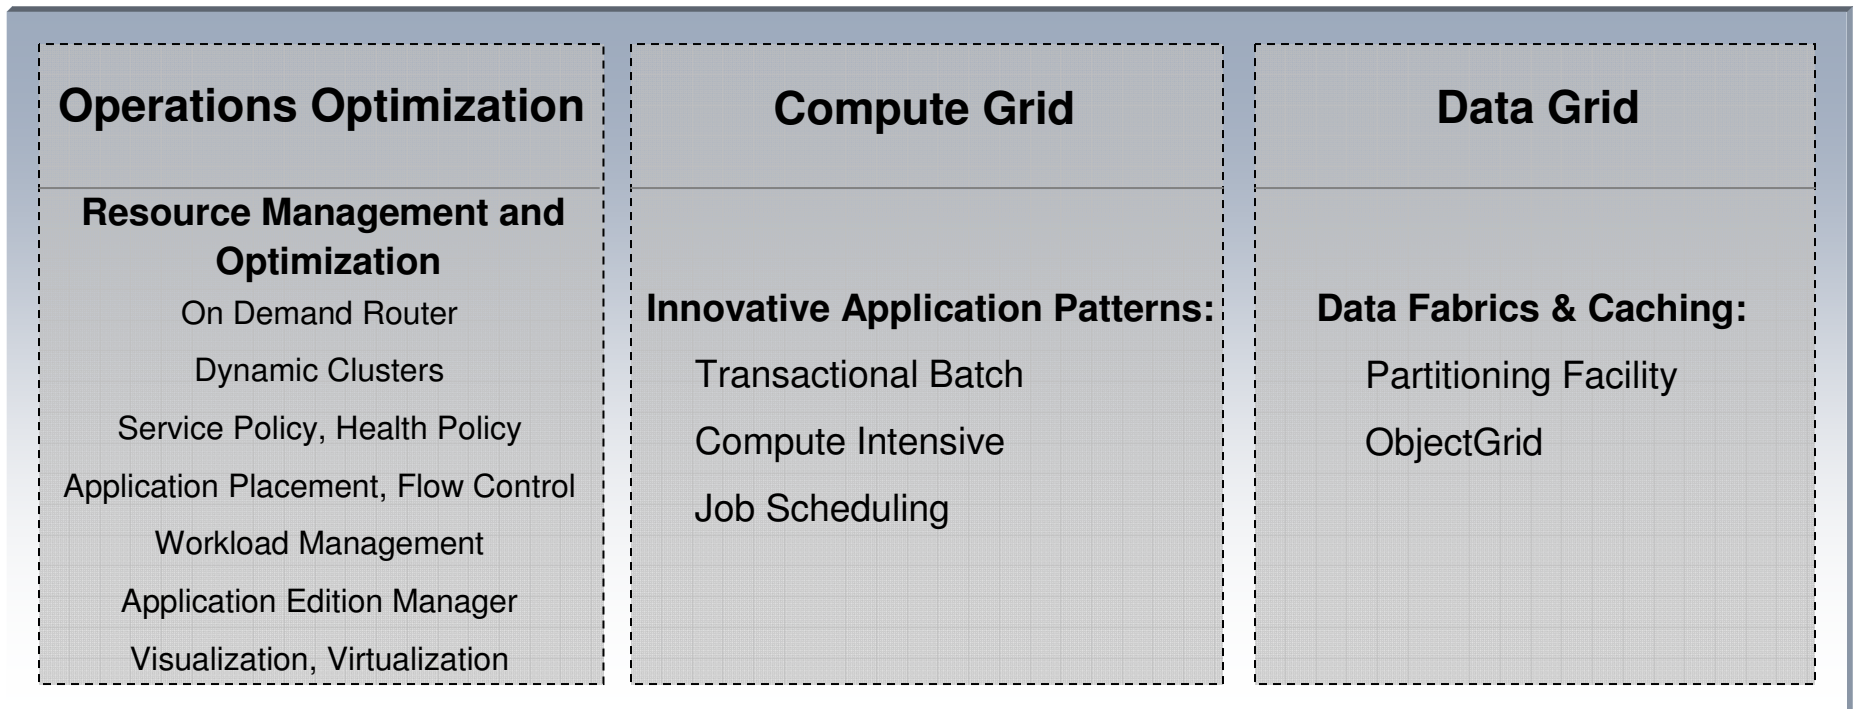
Perfect for the System z customer requiring high QoS and significant integration with CICS, IMS or DB2

Linux on System z: Linux is Linux... *but...*

- **Why do customers deploy to Linux on System z?**
 - Proximity to data
 - Operational simplification
 - Business resiliency
 - Security
- **What are they doing?**
 - **68%: Application Serving for z/OS (hosting multi-tier solutions on System z)**
 - z/OS as data serving back-end
 - Linux on System z hosts all other tiers (application servers, edge servers, etc.)
 - Other multi-tier solutions that benefit from System z QoS and/or virtualization
 - 10%: Data serving workloads not appropriate for DB2 on z/OS
 - 10%: Messaging, collaboration and groupware services
 - 10%: Consolidation of infrastructure and network edge services
 - 2%: Application development and deployment leveraging virtualization services
- **Utility Serving for z/OS**
 - Linux utility “appliance” hosting for z/OS, centrally provisioned and managed

WebSphere Extended Deployment for z/OS

- ✓ **Software to virtualize, control, and turbo-charge your application infrastructure**
- ✓ **Available as a single integrated package or as 3 individual components**



Customer Scenarios:

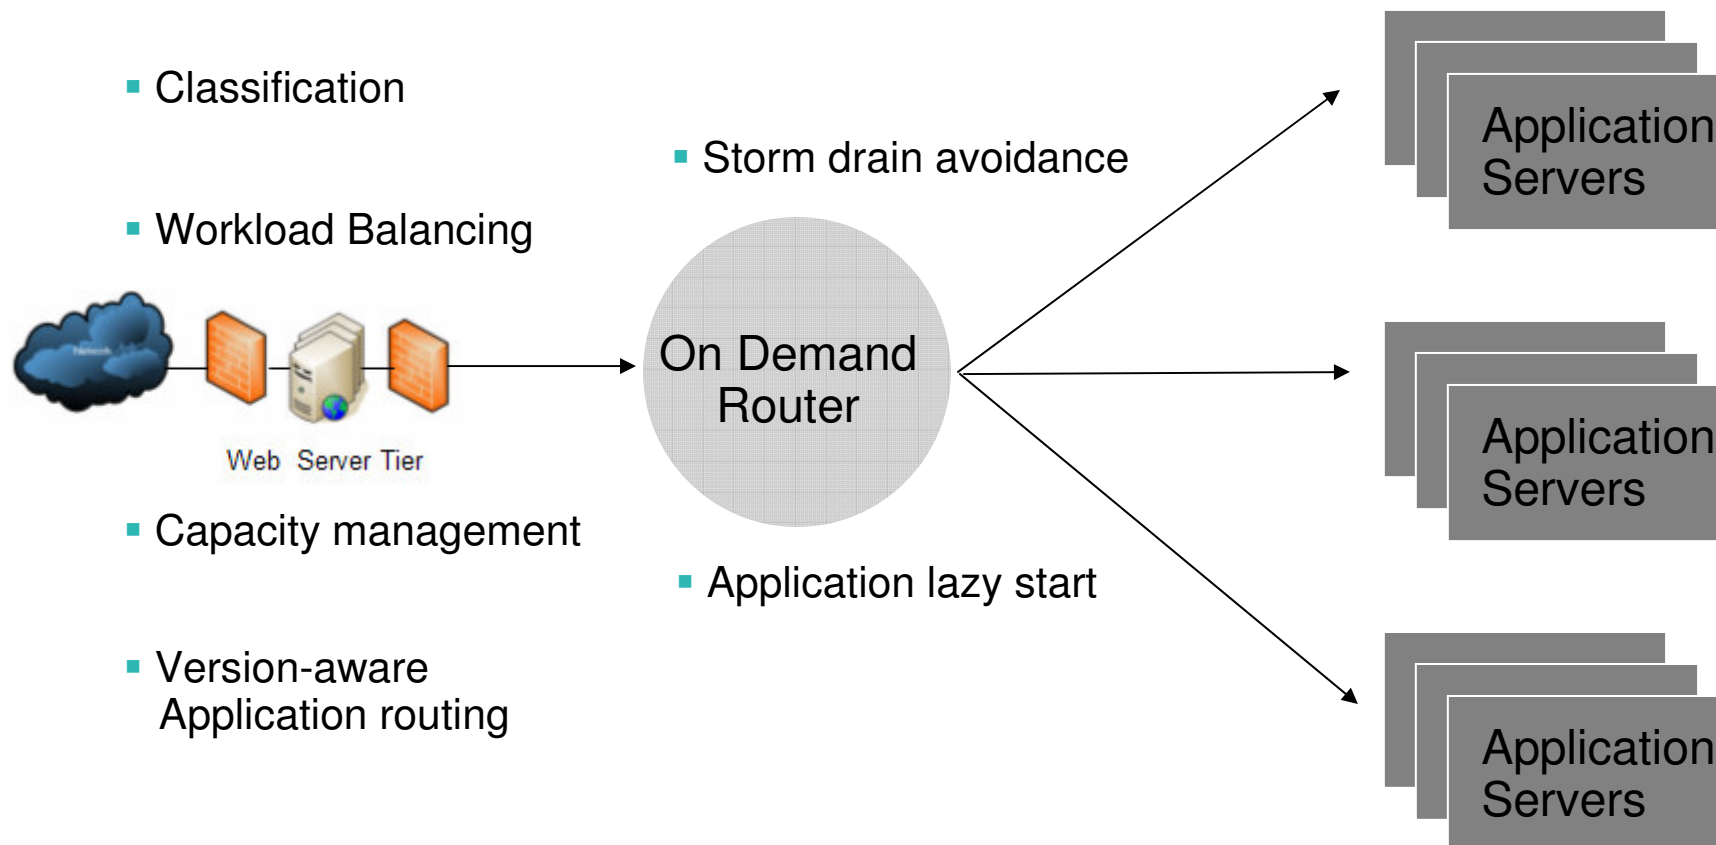
- Common code base for OLTP and Batch
- J2EE / WAS batch programming
- Integrated routing between z/OS and XD
- Application edition management for streamlined deployment and testing



WebSphere Extended Deployment Operations Optimization

On Demand Router (ODR)

- ✓ Provides an intelligent proxy for workload routing within and across LPARs



WebSphere XD Operations Optimization

Application Edition Manager

- ✓ Provides the facility for multiple versions of production applications, including operational support for piloting, staging, and rollback

Visualization

- ✓ Customizable real-time charts display the success of your applications

Runtime maps enable at-a-glance assessments of the components of your application server cells

Health Policy

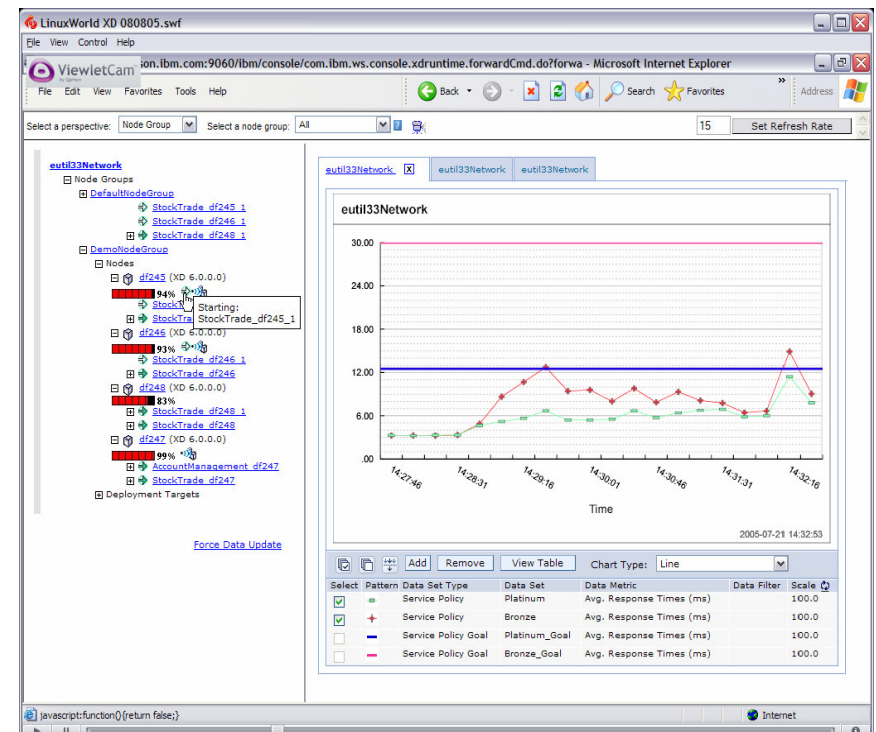
- ✓ Proactively seeks software maladies, such as hung servers, excessive memory consumption, storm drain situations, and addresses them

Service Policies

- ✓ Enables granular classification of application workloads which is addressed within the ODR
- ✓ Value is compounded by z/OS WLM which manages XD application workloads, as well as all System z workloads

Application Lazy Start

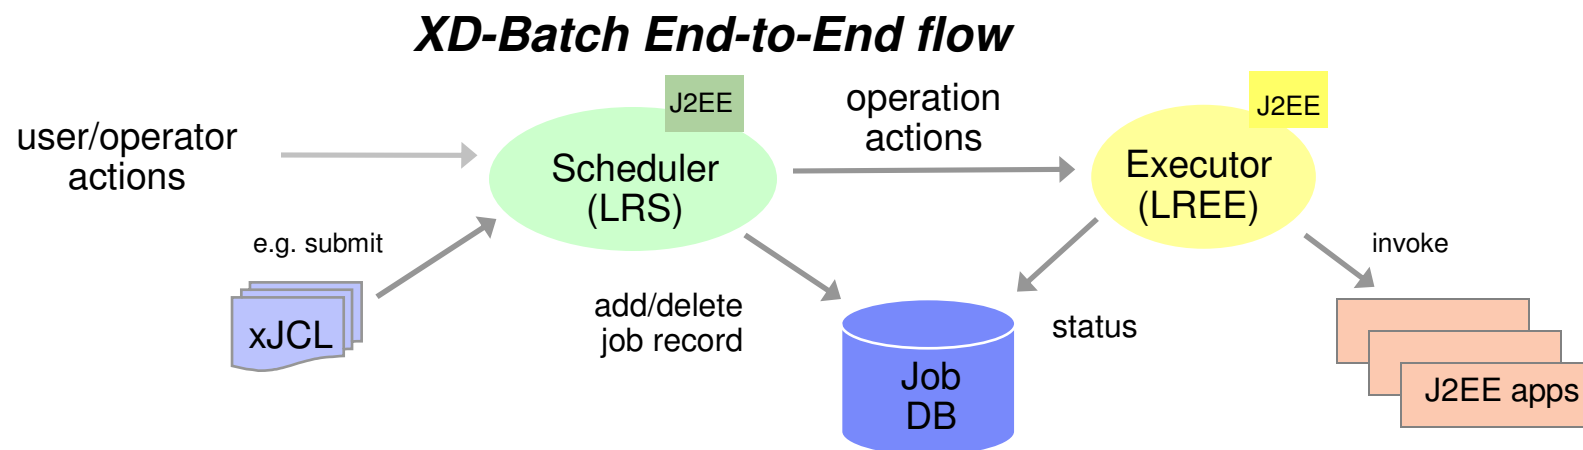
- ✓ Deactivated applications are lazily started when requests come in



WebSphere XD Compute Grid

WebSphere Batch environment

- ✓ Designed for structured J2EE batch workloads
- ✓ Scheduling agent to ensure batch workloads are disseminated to garner unused WebSphere resources
- ✓ Service policy support to differentiate workload importance





WebSphere XD Data Grid

Partitioning Facility

- ✓ Ability to break applications into partitions which enable smart, consistent caching at a partition level
- ✓ Highly available application partitions
- ✓ Provides reduced lock contention on shared DB2 resources

ObjectGrid


- ✓ Creates a caching fabric for shared object data which reduces unproductive backend datasource transactions



New and Enhanced!

And it's Not Just for WebSphere Application Server...

WebSphere XD virtualizes, optimizes and manages the most application servers & environments in the industry



WXD v6.1 Information Links:

Announcement IBM WebSphere Extended Deployment V6.1

IBM Offering Information (United States)

- IBM WebSphere Extended Deployment V6.1 Software Announcement 207-088 PDF File
- IBM WebSphere Extended Deployment for z/OS, V6.1 Software Announcement 207-087 PDF File

IBM WebSphere Extended Deployment V6.1 Product Information

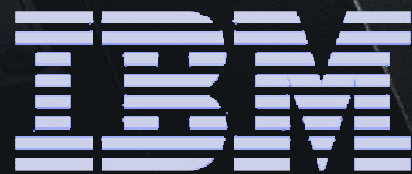
- IBM WebSphere Extended Deployment Product Overview
<http://www.ibm.com/software/webservers/appserv/extended/>
- IBM WebSphere Extended Deployment Features and Benefits
<http://www.ibm.com/software/webservers/appserv/extended/features/>
- IBM WebSphere Extended Deployment Product Documentation
<http://www.ibm.com/software/webservers/appserv/extended/library/>
- IBM WebSphere Extended Deployment InfoCenter
<http://publib.boulder.ibm.com/infocenter/wxdinfo/v6r1/index.jsp>
- Redbook: Best Practices for Implementing WebSphere Extended Deployment, SG24-7343
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247343.pdf>
- developerWorks: Extended Deployment
<http://www.ibm.com/developerworks/websphere/zones/xd/>
- IBM WebSphere Application Server Network Deployment Product Overview
<http://www.ibm.com/software/webservers/appserv/was/network/>
- The IBM Solution Assurance Library (Search for Extended Deployment)
<http://w3.ibm.com/support/assure/assur30i.nsf/Web/SA>

The IBM logo is visible on the top left of a server rack, rendered in a metallic, embossed style.

THANK YOU

The words 'THANK YOU' are written in large, 3D block letters. Each letter is filled with a different portrait of a diverse group of people, including men and women of various ethnicities and ages. The letters are set against a dark blue background with a perspective of server racks.

wgjones@us.ibm.com

The IBM logo is displayed at the bottom center of the image in a white, stylized font.

Trademarks

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

| | | | |
|----------------------------|---|-------------------------------------|-----------------------|
| AIX* | ES/9000* | Language Environment* | S/390* |
| AnyNet* | FICON* | Lotus* | Sysplex Timer* |
| CICS* | FlashCopy* | Multiprise* | System Storage |
| DB2* | GDPS* | Notes* | System z |
| DB2 Connect | Geographically Dispersed Parallel Sysplex | OMEGAMON* | System z9 |
| DB2 Universal Database | HiperSockets | On demand business logo | SystemPac* |
| developerWorks* | Hiperspace | OS/390* | Tivoli* |
| DFSMSdfp | HyperSwap | Parallel Sysplex* | TotalStorage* |
| DFSMSdss | IBM* | PR/SM | Virtualization Engine |
| DFSMSHsm | IBM eServer | Processor Resource/ Systems Manager | VTAM* |
| DFSMSrmm | IBM e(logo)server* | pSeries* | WebSphere* |
| DFSORT | IBM logo* | RACF* | z/Architecture |
| Domino* | IMS | RAMAC* | z/OS* |
| DRDA* | Infoprint* | Redbook | z/VM* |
| Enterprise Storage Server* | IP PrintWay | RMF | zSeries* |
| ESCON* | iSeries | | |

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of the Intel Corporation in the United States and other countries.

Linux is a trademark of Linux Torvalds in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Notes:

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

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

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

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

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

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

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

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.