



WebSphere software

WebSphere Extended Deployment (XD)

Overview of Version 5.1 and Upcoming Version 6.0 Content



IBM Software Group

ON DEMAND BUSINESS™

WebSphere Software for On Demand Business: Integration and Infrastructure Software to Maximize Business Flexibility and Responsiveness

WebSphere software *We make **IT** work for business*

Innovative interactions

Maximize the use of your IT infrastructure to support new business models and reach users in new ways

Extend
your reach

Integrate
your people,
processes and
information

Optimize
your application
infrastructure

Improved flexibility and speed

Increase your responsiveness by connecting the right people with the right information at the right time

Operational excellence

Improve your reliability and performance with a proven, secure IT platform

Using WebSphere Extended Deployment to Optimize Your Application Infrastructure

Dynamic Operations

- Virtualized WebSphere Environment
- Goals-Directed Infrastructure
- Autonomic Management

Extended Manageability

- Visual Operational Monitoring
- Health Monitoring

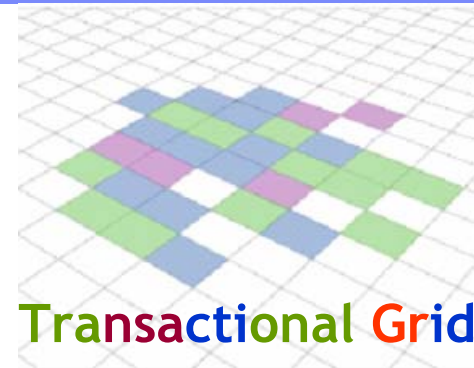


High Performance Computing

- Highly Scalable, Partitioned Applications



Dynamic Operations

- **Virtualized WebSphere Environment**
 - ▶ Resources are virtualized in a common pool so that they can be shared amongst multiple, transactional applications
- **Goals-Directed Infrastructure**
 - ▶ Workloads are prioritized, queued and routed according to established business goals and relative application importance
 - Service Policies
 - Goal-based Workload and Resource Management
 - Differentiated Workloads
 - Health Policies for early detection and automated correction of system problems
- **Autonomic Management**
 - ▶ Application resources can be dynamically adjusted based on actual demand using autonomic managers
 - ▶ Manual, Supervised and Autonomic “on demand” Modes





Static Clustered Environment Leads to Application "Silos"



Stock Trading



Account Management



Financial Advice



Resource Virtualization Maximizes Utilization and Improves Responsiveness




Stock Trading

Customer Support


Account Management

Risk Management

Financial Advice



RESOURCE POOL



* Hypothetical, for illustrative purposes only

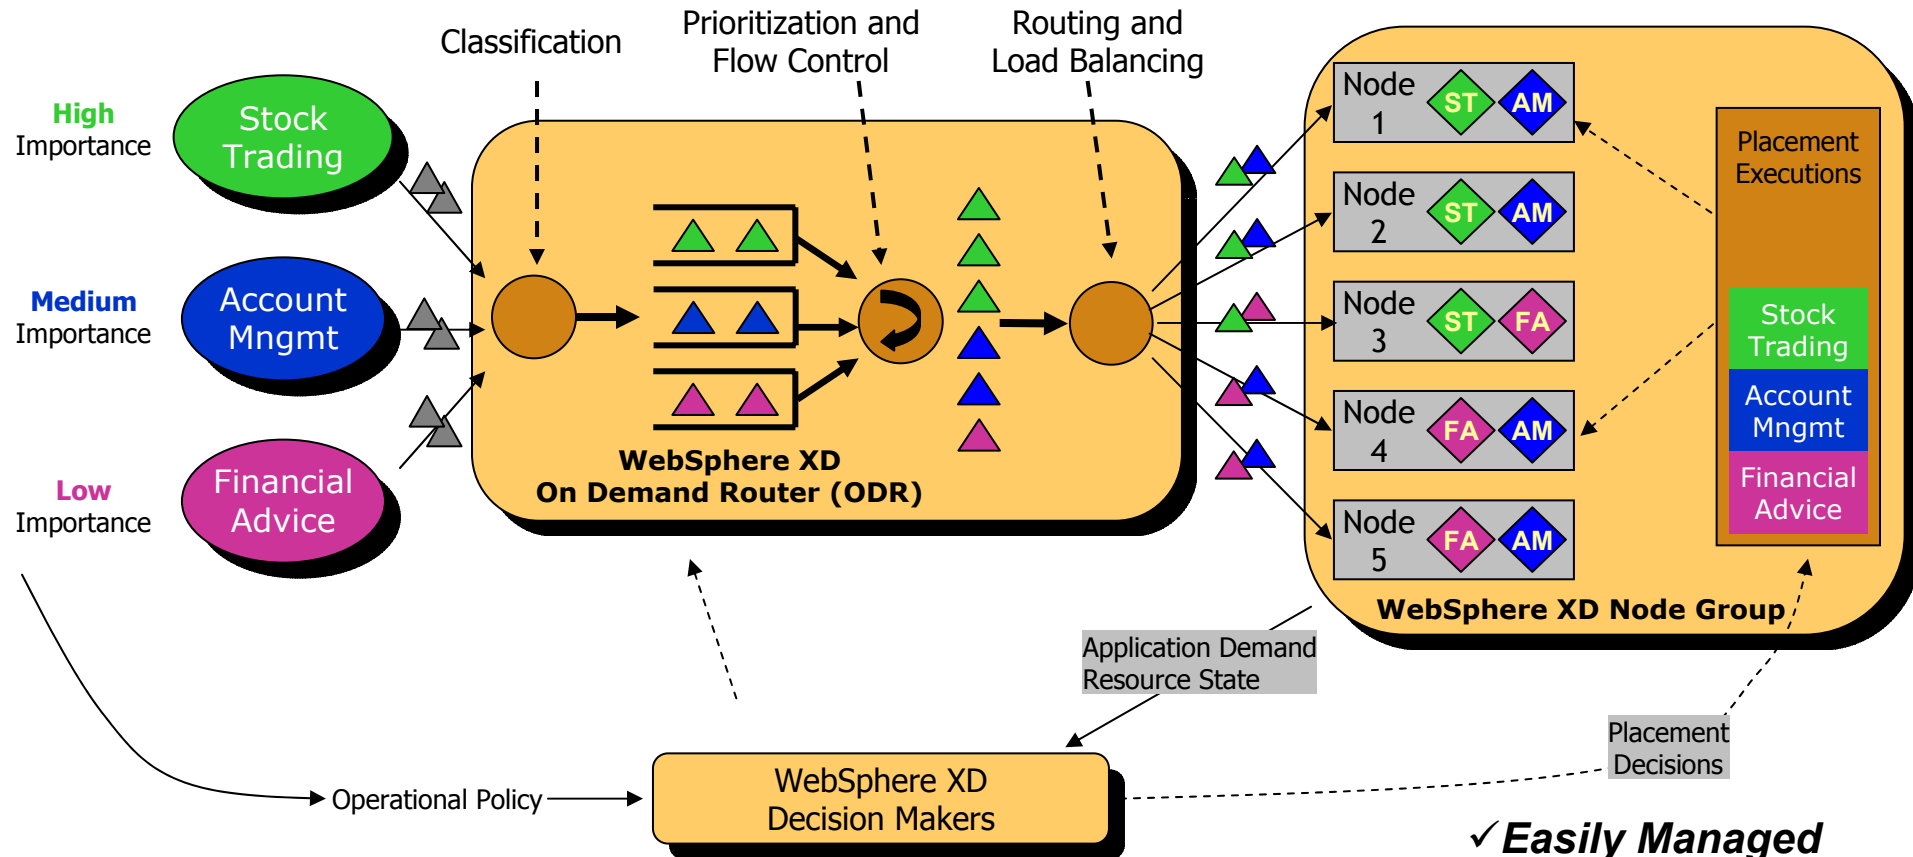
Key Features

- On Demand Router (ODR)
 - ▶ New Process built as a set of filters to the new WAS Proxy Server
 - ▶ Provides routing of HTTP traffic into a collection of WAS Cells
 - ▶ Provides three key On Demand Features
 - Flow Control/Queuing
 - Prioritization
 - Dynamic Workload Management
- Application Placement
 - ▶ Autonomic Placement of instances of applications on servers
 - ▶ Starts/Stops Preconfigured applications
 - ▶ At least one instance of each application running at all times.

WebSphere Dynamic Operations Environment

✓ **Goals-Directed**

✓ **Virtualized**

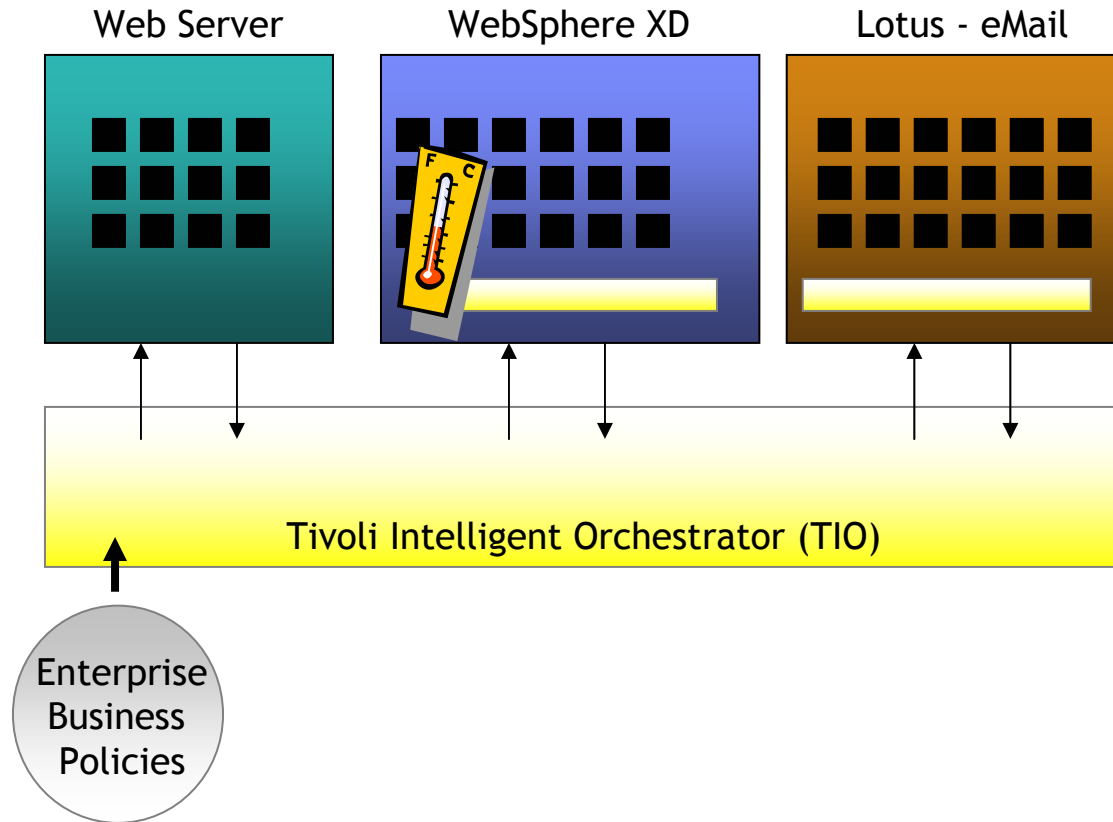


✓ **Easily Managed**


✓ **Autonomic**

✓ **Reliable, Scalable, High Performance**

Provisioning in XD using TIO




XD Topologies

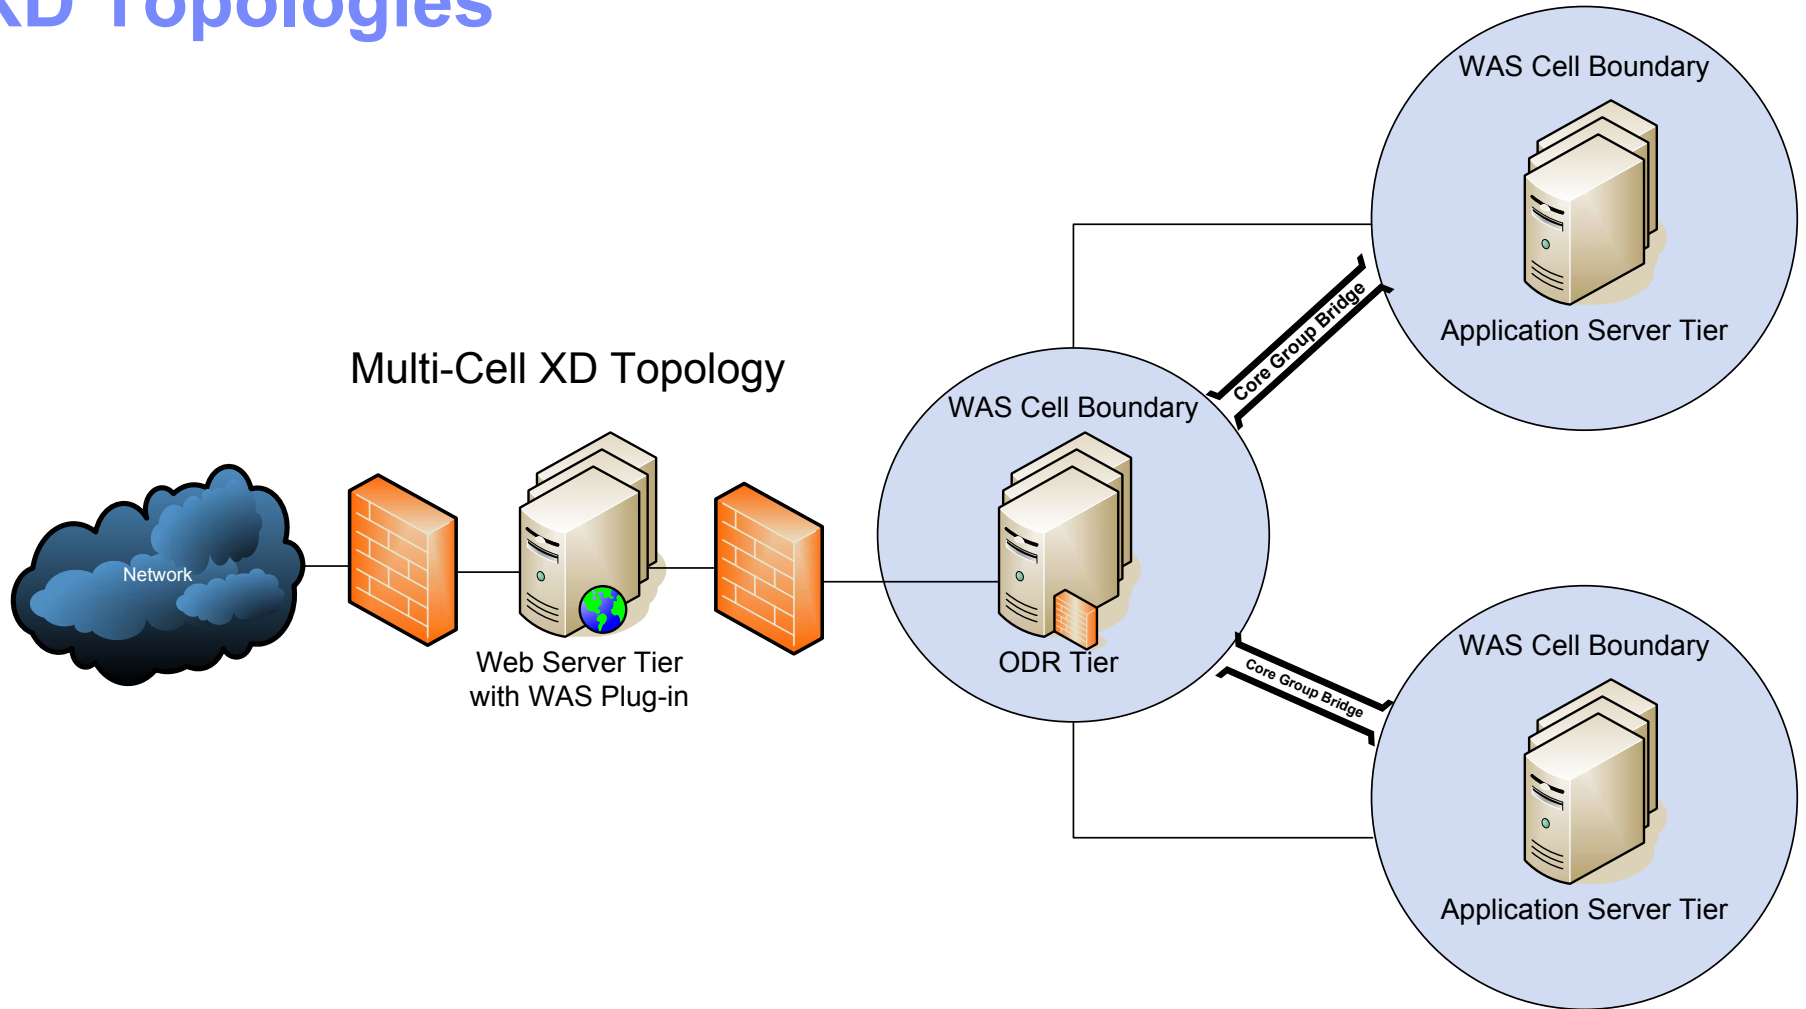


- Normal Topology would have multiple ODRs for both scale and HA reasons
- ODR will generate a new plugin-cfg.xml file for the WAS plugins
- ODR dynamically learns configuration of backend WAS Cell
- ODRs can be connected to multiple WAS cells simultaneously
- ODR support routing to any HTTP endpoint
 - ▶ Could be place in front of HTTP Servers

Standard XD Topology



XD Topologies



High Performance Computing with WebSphere Extended Deployment

Dynamic Operations

- Virtualized WebSphere Environment
- Goals-Directed Infrastructure
- Autonomic Management

Extended Manageability

- Visual Operational Monitoring
- Health Monitoring



High Performance Computing

- Highly Scalable, Partitioned Applications

High Performance Computing – Key Concepts

▪ Scalable

- ▶ Achieve near linear scalability for high write-rate OOLTP applications
- ▶ WebSphere Partitioning Facility (WPF) allows for the partitioning of applications and data, improving database as well as in-memory caching and workload management

▪ High Availability


- ▶ High reliability for high-end OLTP (class 5 availability)
- ▶ Quick application recovery in response to resource failures

▪ OLTP with less skill and lower costs

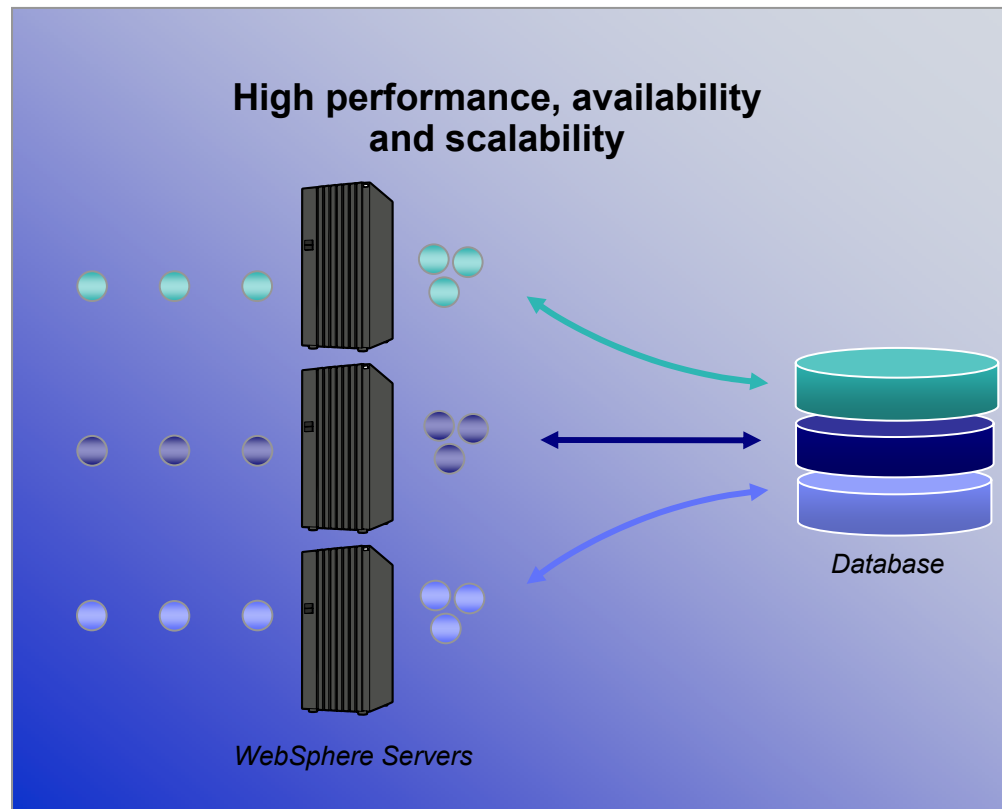
- ▶ OLTP on Java at lower cost and leveraging existing Java skills
- ▶ Unified administration and management environment



Transaction Throughput and Linear Scalability Can Be Limited by Database Access Requirements



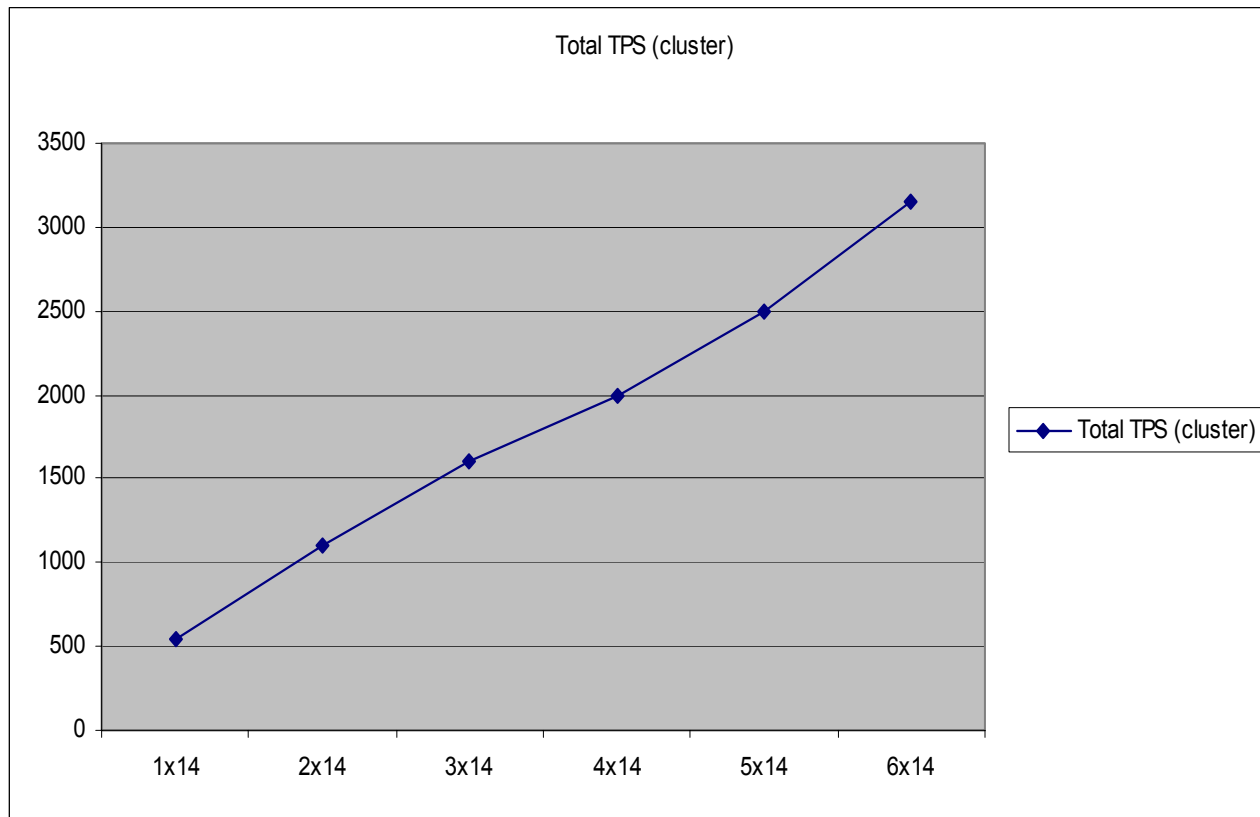
WebSphere Extended Deployment Provides Highly Reliable and Scalable Transaction Processing



WebSphere Partitioning Facility – Use Cases

- **Linear Scaleable J2EE applications with excellent high availability**
- **Distributed Cluster Wide Cache**
 - ▶ Using DynaCache as the caching system
 - ▶ Enables a logical cache to be split over multiple physical servers
 - ▶ Enables large caches (> 2 GB)
- **Creation of High Availability Singletons**
 - ▶ Create a singleton service that runs on a single JVM but it still highly available
- **Creation of Application-based Affinities**
 - ▶ Bind to a particular server based on application defined parameters
 - Not just HTTP session ID or Transaction ID

Linear Scalable J2EE Applications with Excellent High Availability



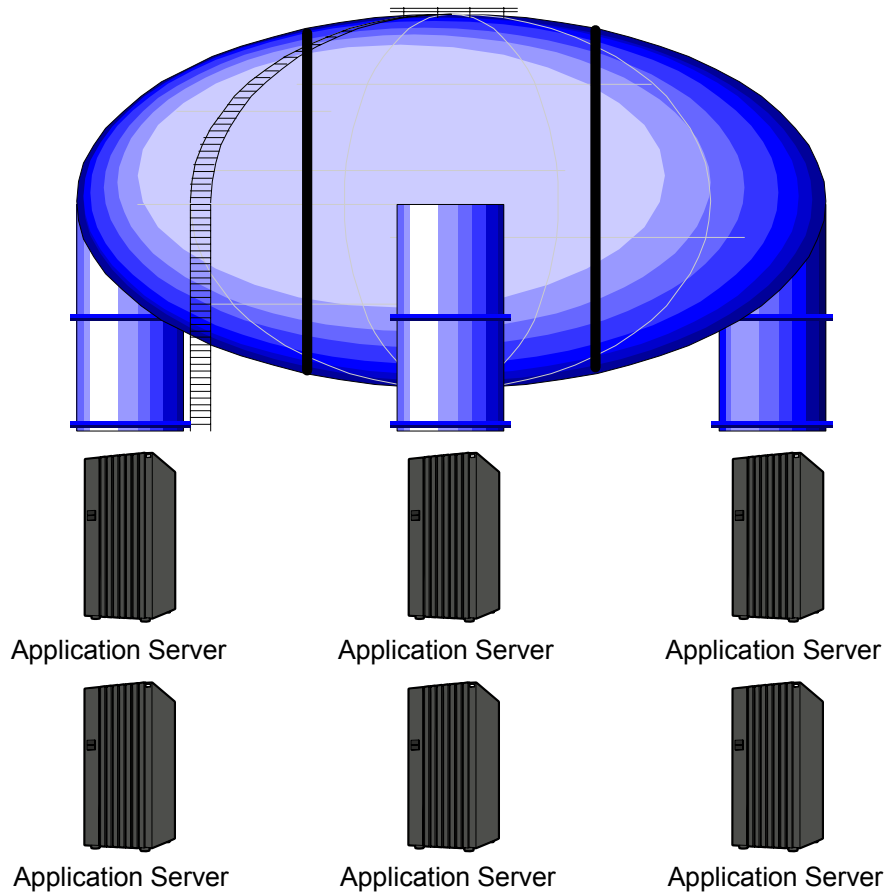
MxN means M boxes with N threads driving work.

Average response time is 31ms*

Note: TPS = transactions per second

*The performance data contained herein was determined in a controlled customer environment. Therefore, the results obtained in other customer operating environments may vary significantly.

Distributed Clustered Cache



WPF Features

- Accessible from multiple protocols
 - ▶ IIOP with client side partition identification handler
 - ▶ HTTP through the partition.xml configuration in the ODR
 - ▶ JMS through an application design pattern to “pull” the appropriate messages into the server
- Simple Programming Model
 - ▶ Add a single new Stateless Session EJB
 - ▶ EJB Implements the PartitionHandler Interface
 - `getPartitionNames()`, `loadPartition(x)`, `unloadPartition(x)`
- Flexible Partition Definition
 - ▶ Policy files allow control over partition assignment, grouping, and failure policy
 - ▶ Allows external coordination and fail-back policy
- Integration with Blade Center hardware for more predicable failure scenarios
- Virtual Data Source Support for Partition-aware JDBC Access

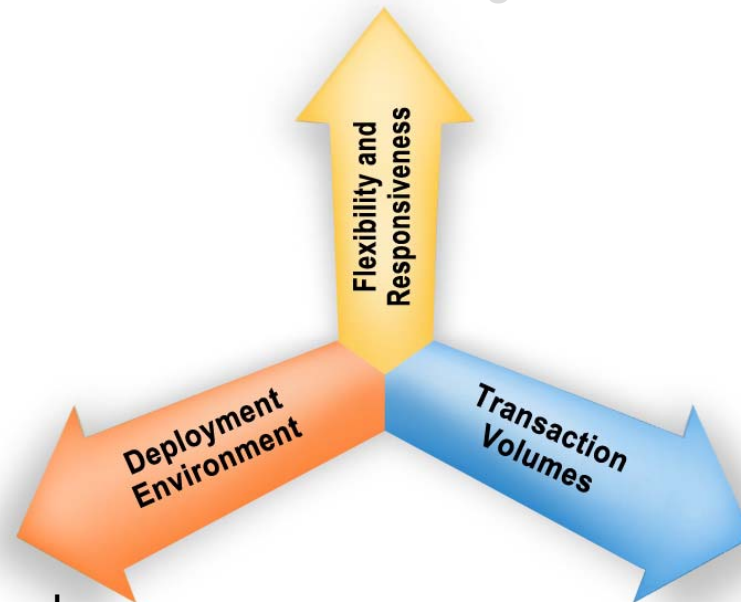
Extended Manageability with WebSphere Extended Deployment

Dynamic Operations

- Virtualized WebSphere Environment
- Goals-Directed Infrastructure
- Autonomic Management

Extended Manageability

- Visual Operational Monitoring
- Health Monitoring



High Performance Computing

- Highly Scalable, Partitioned Applications

Extended Manageability – Key Concepts

■ Visual Operational Monitoring

- ▶ Operations console for managing and monitoring a production environment against defined goals
 - At a glance runtime map
 - Runtime Operational State (where applications are running and how are we doing)
 - Detailed goal-oriented performance charting
- ▶ Simplified Administration Support for Clustered Environments

■ Health Monitoring

- ▶ Monitoring of certain aspects of the WebSphere system to detect and take action on potential problems
- ▶ When a problem condition is detected, one of three action styles can be put into effect
 - Monitored
 - Supervisory
 - Automatic



Visual Operational Monitoring

The screenshot displays the WebSphere Administrative Console interface. The main content area shows the configuration for the 'StockNodeGroup' node group. Under 'Nodes', 'AppNode2' is at 99% and 'AppNode1' is at 88%. A 'Starting Provisioning Server Instance' dialog box is visible over AppNode2. The 'Dynamic Clusters' section lists 'StockTrade', 'Account/management', and 'FinancialAdvice'. A chart titled 'Avg. Response Times (ms) Per Dynamic Cluster in Node Group:StockNodeGroup' shows response times for various DC's (AccountManagement, StockTrade, Bronze_Goal, Gold_Goal, Platinum_Goal, Silver_Goal) from 08:31:05 to 08:32:26. The chart shows that most DC's maintain low response times, while 'AccountManagement' shows a significant spike at the end of the period.

Time	AccountManagement	StockTrade	Bronze_Goal	Gold_Goal	Platinum_Goal	Silver_Goal
08:31:05	~30	~30	~1200	~1800	~1500	~2100
08:31:25	~30	~30	~1200	~1800	~1500	~2100
08:31:45	~30	~30	~1200	~1800	~1500	~2100
08:32:05	~30	~30	~1200	~1800	~1500	~2100
08:32:26	~1000	~30	~1200	~1800	~1500	~2100

At the bottom, the 'WebSphere Status' section shows configuration problems: 0 total configuration problems, 6 total warnings, and 0 total errors. The system time is June 24, 2004 8:32:06 AM CDT.

Health Monitoring

→ Step 2 : Define Health Policy Health Condition Properties

Maximum Age	* 7	Days	▼	<p>i The age to allow the Application Server Instance to "live" prior to corrective action. This value can go from 1 hour to 365 days.</p>
-------------	-----	------	---	---

Previous

→ Step 2 : Define Health Policy Health Condition Properties

Total Requests	* 20000000			<p>i The total number of requests that the application server instance serves before corrective action. It must be between 1000 or greater.</p>
----------------	------------	--	--	--

→ Step 2 : Define Health Policy Health Condition Properties

Total Memory Used	* 85	%		<p>i The threshold value for the percentage of memory over the max heap size used for the JVM process. This value can go from 1 to 99.</p>
Time Over Memory Threshold	* 5	Minutes	▼	<p>i The time that the total memory must be over the threshold value prior to corrective action. This value can go from 1 second to 60 minutes.</p>

Previous

→ Step 2 : Define Health Policy Health Condition Properties

Response Time	* 10	Seconds	▼	<p>i The average time that requests can take prior to corrective action. This value can go from 1 millisecond to 1 hour.</p>
---------------	------	---------	---	---

Previous Next Cancel

E-Mail Notification

- Health monitoring events can be sent automatically to administrator's e-mail
- The notification configuration looks something like this:

General Properties		
SMTP HostName	* localhost	<i>i</i> Specifies the SMTP server to connect to when sending mail.
SMTP Port Number	* 25	<i>i</i> Specifies the SMTP port number to connect to when sending mail.
SMTP UserId		<i>i</i> Specifies the user ID to use when the SMTP mail server host requires authentication.
SMTP Password		<i>i</i> SMTP Specifies the password to use when the SMTP mail server host requires authentication.
Enable Notifications	<input type="checkbox"/>	<i>i</i> Specifies whether or not email notifications are enabled.
Email Addresses	<i>i</i> Specifies the list of email addresses to send email to when notifications are enabled.	
	<p>Email Address</p> <input type="text"/>	<p>Current Email Addresses</p> <div style="border: 1px solid black; height: 150px; width: 100%;"></div>
	<input type="button" value="Add >>"/>	<input type="button" value="<< Remove"/>
<input type="button" value="Apply"/> <input type="button" value="OK"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>		

WebSphere Extended Deployment Design Principles

- An “add-on” or “extension” to the WAS ND environment
 - ▶ Prereqs WAS ND instead of bundling it
- Installation is a simple delta to an existing environment
 - ▶ Does not require “migration” or restructuring of current installation
- Totally integrated into the WAS ND environment
 - ▶ Extends the WAS Admin Console
 - ▶ Extends the wsadmin scripting environment
- Meaningful without implementing full Autonomic concepts
 - ▶ Manual and Supervised modes allow autonomies to be adopted gradually
 - ▶ Goals-directed WLM can be implemented without autonomic features
 - ▶ Extended Manageability features such as visualization and health monitoring have broad appeal

WebSphere XD V5.1 Beta Participants

Objectives for use of WebSphere XD:

- Want to see the cost reductions that can be realized when existing capacity is used to maximum efficiency
- Want to dynamically distribute the workload across the available capacity to satisfy business goals and reduce costs
- Want to provide differentiated qualities of services to authenticated users
- Interested in implementing a Grid strategy

“We think WebSphere XD has a potential capability for our next Grid infrastructure. We are expecting XD innovation to meet our goal in the near future”.





WebSphere software

WebSphere Extended Deployment Version 6.0

Messaging and Content Preview



IBM Software Group


ON DEMAND BUSINESS™

WebSphere XD 6.0 Focus Areas

- WebSphere Business Grid
 - ▶ Long Running Workloads
 - J2EE Batch
 - J2EE Parallel Computationally Intensive
 - ▶ Combined with OLTP
- ObjectGrid
 - ▶ Transactional Replicated Distributed Caching for transparent object data access
 - ▶ WebSphere Partitioning Facility (WPF)
- Application Versioning and Continuous Availability
 - ▶ Managed Rollout of New Applications
- Scale Out
 - ▶ Support for non-WebSphere server environments
- Enhancing the Dynamic Operations Environment
- Enhancing the Extended Manageability Capabilities

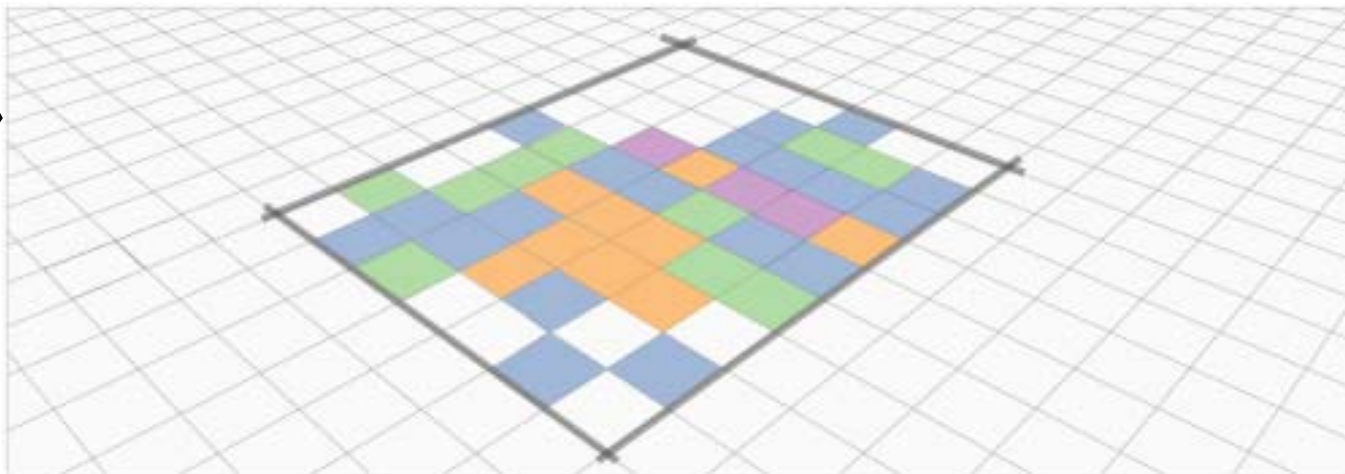


Goal: One Virtualized Infrastructure, One Management Environment



 **Computational Grid – LR Workloads**

 **Transactional Grid – OLTP workloads**

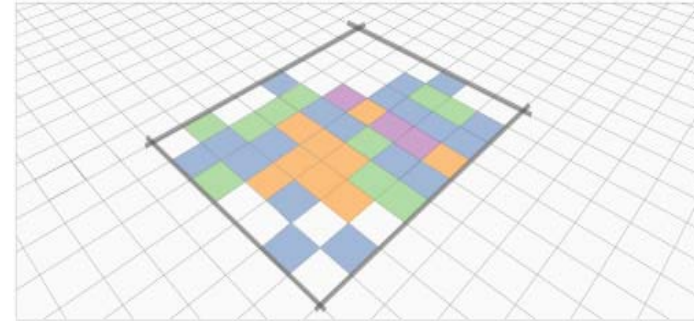


 **WebSphere Business Grid – mixed workloads (LR + OLTP)**





WebSphere Business Grid – Overview

- Delivers
 - ▶ Goal-oriented resource management of long running and transactional workloads sharing the same resource pool (whitespace harvesting)
- Expands
 - ▶ Asynchronous job execution
 - ▶ Parallel work execution
 - ▶ Container-managed Batch execution
 - ▶ Transparent Data Access via the Global Cache
 - ▶ Non-J2EE applications using SOA
- Integrates
 - ▶ Coexistence of J2EE OLTP, J2EE Batch, J2EE CI, and Non-J2EE applications in a Dynamic Cluster
 - ▶ Balancing of these diverse workloads
- Flexible Scheduling for LR Workloads
 - ▶ Specify when requests execute (time)
 - ▶ Event-based Scheduling
 - ▶ Policy/Goal Based Scheduling
 - ▶ Concurrent with OLTP




WebSphere Partitioning Facility Use Cases

- Linear Scaleable J2EE applications with excellent high availability
- Distributed Cluster Wide Cache
 - ▶ Using DynaCache as the caching system
 - ▶ Enables a logical cache to be split over multiple physical servers
 - ▶ Enables large caches (> 2 GB)
- Creation of HA Singletons
 - ▶ Create a singleton service that runs on a single JVM but it still highly available
- Creation of Application-based Affinities
 - ▶ Bind to a particular server based on application defined parameters
 - Not just HTTP session ID or Transaction ID



ObjectGrid

- Distributed Transactional Object Cache constructed from a portable, pluggable core with extensions
- Provides numerous application APIs to access cache from
 - ▶ POJOs, SDOs, JDBC, CMP EJBs, HTTP Session, etc
- Transactional
 - ▶ 1-Phase Tx Support (2-Phase Tx in Phase II)
 - ▶ Maintain consistency across transactions
 - ▶ Supports multiple isolation levels
 - ▶ ACID properties will be observed.
- Provides cache lifecycle features
 - ▶ Declaration, Configuration, Invalidation, Size Management, Cache Loading
- Provides Distribution features
 - ▶ Write-through, replication, triggers, partitions, WLM-integration, shared processes, file-based on NAS
- Provides hierarchical tag-based invalidation
- Componentized as a standalone feature
 - ▶ Run in WAS
 - ▶ Run in any J2SE environment (1.4.x)



Application Versioning/Continuous Availability

- Supports online application upgrade scenario
- Explicit control over application editions
 - ▶ creation/activation/deletion
- Continuous operations
 - ▶ Explicit orchestration between routing agent and application servers during upgrade
- Admin Console and scripting support
- Support multiple rollout policies
 - ▶ Rolling Upgrade
 - ▶ Divide and Switch
- Support Validation Mode
- Gradual workload increase
 - ▶ Users (generic or specific)
 - ▶ Requests
- Helps to reduce cost
 - ▶ Redundant Cells/Hardware

The screenshot displays the WebSphere Administrative Console in a Microsoft Internet Explorer browser window. The address bar shows the URL `http://localhost:9090/admin/secure/logon.do`. The page title is "WebSphere Administrative Console - Microsoft Internet Explorer". The main content area is titled "Application Edition Control Center > PlantsByWebSphere >".

On the left side, there is a navigation menu with the following items:

- User ID: vignola
- cvignolaNetwork
 - ▣ Servers
 - ▣ Applications
 - [Enterprise Applications](#)
 - [Install New Application](#)
 - [Edition Control Center](#)
 - ▣ Resources
 - ▣ Security
 - ▣ Environment
 - ▣ System Administration
 - ▣ Troubleshooting

The main content area is titled "Manage Edition Deployments" and includes the following text:

View and change the application editions deployed and running on each server. [i]

Deployment: WebSphere:cell=cvignolaNetwork,cluster=productCluster1

Change the edition deployment to this target by selecting an edition number, one or more servers, then clicking the Deploy Edition button.

Editions:


- 1.0.0 Base edition.
- 1.5.0 Edition that supports automated orders submission to wholesaler.
- 2.0.0 Edition that supports customer personalization - under test.

Buttons:



Server	Node	Edition	Edition Status
<input type="checkbox"/> productionServer1	node1	1.0.0	➔
<input type="checkbox"/> productionServer2	node2	1.0.0	➔
<input type="checkbox"/> productionServer3	node2	1.5.0	➔

Scale Out

- Goals
 - ▶ Separate application server administration from application QoS (WLM, HA, session mgmt, etc.)
 - ▶ Apply QoS to non-WAS environments
 - ▶ Easily add and configure new application server instances (in a “blade-like” fashion)
 - ▶ Common admin and management across mixed application server environment
- Phase I
 - ▶ Componentized QoS Features
 - ODR, policy-based WLM for non-WebSphere and Standalone WebSphere environments
 - ObjectGrid for non-WAS J2SE/J2EE Environments
- Phase II
 - ▶ Scale-Out Admin
 - Template-based server configuration with hot-plugging support
 - Peer-to-Peer admin based on shared file system
 - No DMgr, Config from anywhere, peer-JMX routing



Scale Out Topology



Dynamic Operations Enhancements

- Support multiple protocols
 - ▶ IIOP, Messaging
- Support Landing Zones
 - ▶ Applications Configured but lazily started
- Support Routing Rules
 - ▶ Enables pilots and other advanced topologies
- Support flexible classification rules
 - ▶ Custom, By User/Role, By Virtual Port ID, By Business Process, By Header
- Integration with EWLM



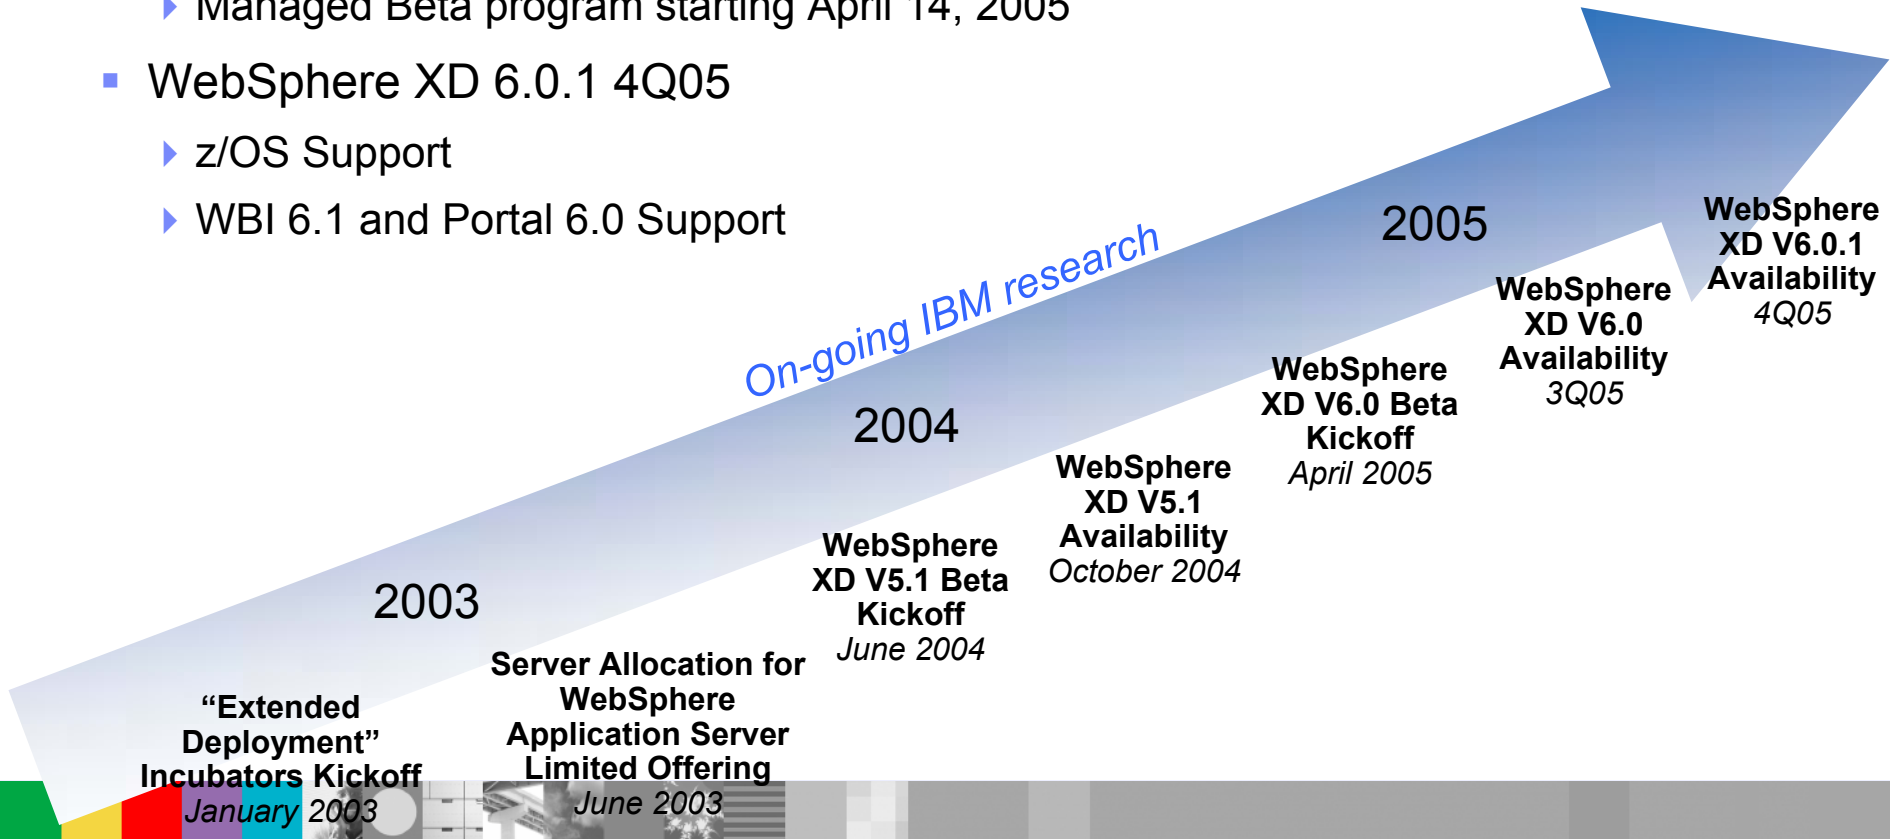
Extended Manageability Enhancements

- Phase I
 - ▶ Enhanced Visualization
 - ▶ Data logging and replay
 - ▶ Metrics for Charge back
 - ▶ HA DMgr (hot standby model)
 - ▶ Repository checkpoint/restore
 - ▶ Improved application monitoring/visualization
 - ▶ Health monitoring
 - ▶ Simplified node creation
 - ▶ Support for Visualization of partitions and grid applications
- Phase II
 - ▶ Topology-based viewer
 - ▶ Viz Server support
 - Standalone Visualization Server



Delivery Plan (normal caveats apply)

- WebSphere XD 6.0 for Distributed Platforms in 3Q05
 - ▶ Includes new distributed platforms like 64-bit, zLinux, and HP-UX
 - ▶ Based on WAS ND 6.0.2
 - ▶ Managed Beta program starting April 14, 2005
- WebSphere XD 6.0.1 4Q05
 - ▶ z/OS Support
 - ▶ WBI 6.1 and Portal 6.0 Support





IBM Software Group

Thank you



@business on demand software

© 2005 IBM Corporation
IBM Confidential