

Smarter Workload Optimization – IBM Software And POWER Unleashed




WebSphere On POWER – Optimized Web Facing Engine

IBMDiscoveryDays2011

Copies of Today's Presentations:

<http://www.ibm.com/developerworks/offers/techbriefings/details/power.html>



Workload Optimized Systems

Meet The Insatiable Demand For Processing

IBM Software

Optimized to exploit
IBM hardware
performance features



IBM Power Systems

Performance redefined.
Delivering new services faster
with more cores per server.
Massive parallel architecture
1024 threads and 256 cores



**Lowest cost per
workload yields
huge savings!**



Service Oriented Finance Needs A Web Facing Engine To Expand Their Business

Our banking requirements have increased. I need a stronger infrastructure that has room for growth. Cost is also an important.



Service Oriented Finance CIO

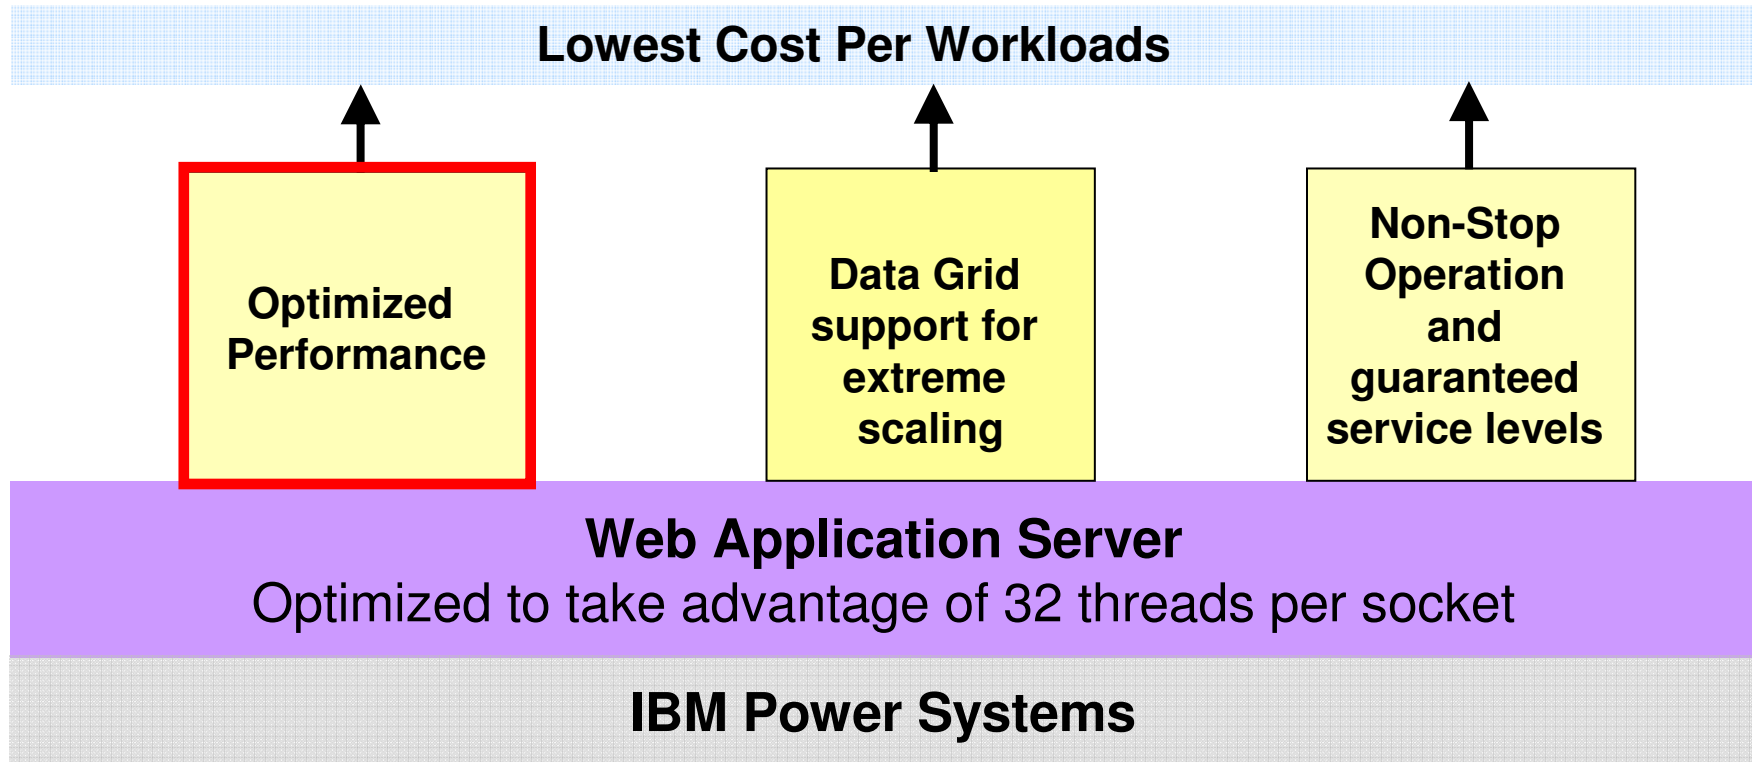
Let us take a look at some of the key factors you should consider for a strong web facing engine that is cost effective.



IBM


IT Demands More Work From Web Application Servers

Do you have a Web Facing Engine optimized to leverage all of the threads available?

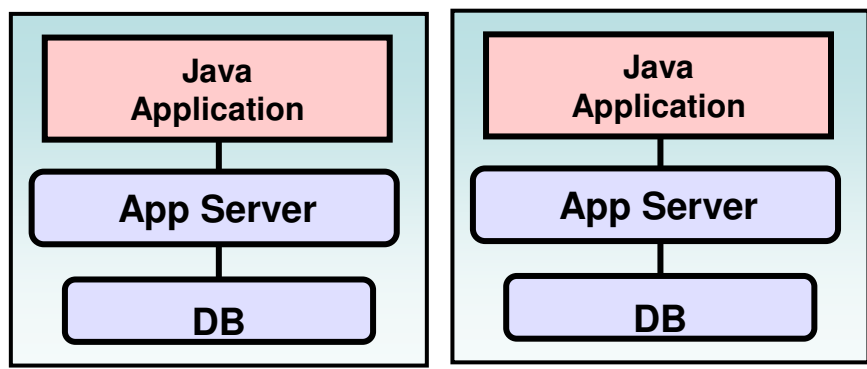


An Optimized Web Application Server Must Address Every Application Server Pattern

1. Single JVM



2. Multi JVM




3. Cluster



WebSphere Application Server Optimized To Exploit 32 Hardware Threads In POWER7 Socket

Results:

Day Trader 2.0 – Single JVM



Near linear scaling on Power Systems, up to 32 threads

Optimizations improve performance by 85% over non-optimized single instance

The WebSphere Application Server Is Optimized For Power Systems

WebSphere Application Server

Reduced code path length when running on POWER
Optimized to reduce the amount of lock contention on POWER
Access to AIX environment parameters

Java Virtual Machine

Uses 64K pages by default on AIX
Elimination of unneeded exit tests in the Just-In-Time (JIT) compiler

Operating System

Caching optimizations in AIX to improve performance when lock contention is occurring in WebSphere

POWER Hardware

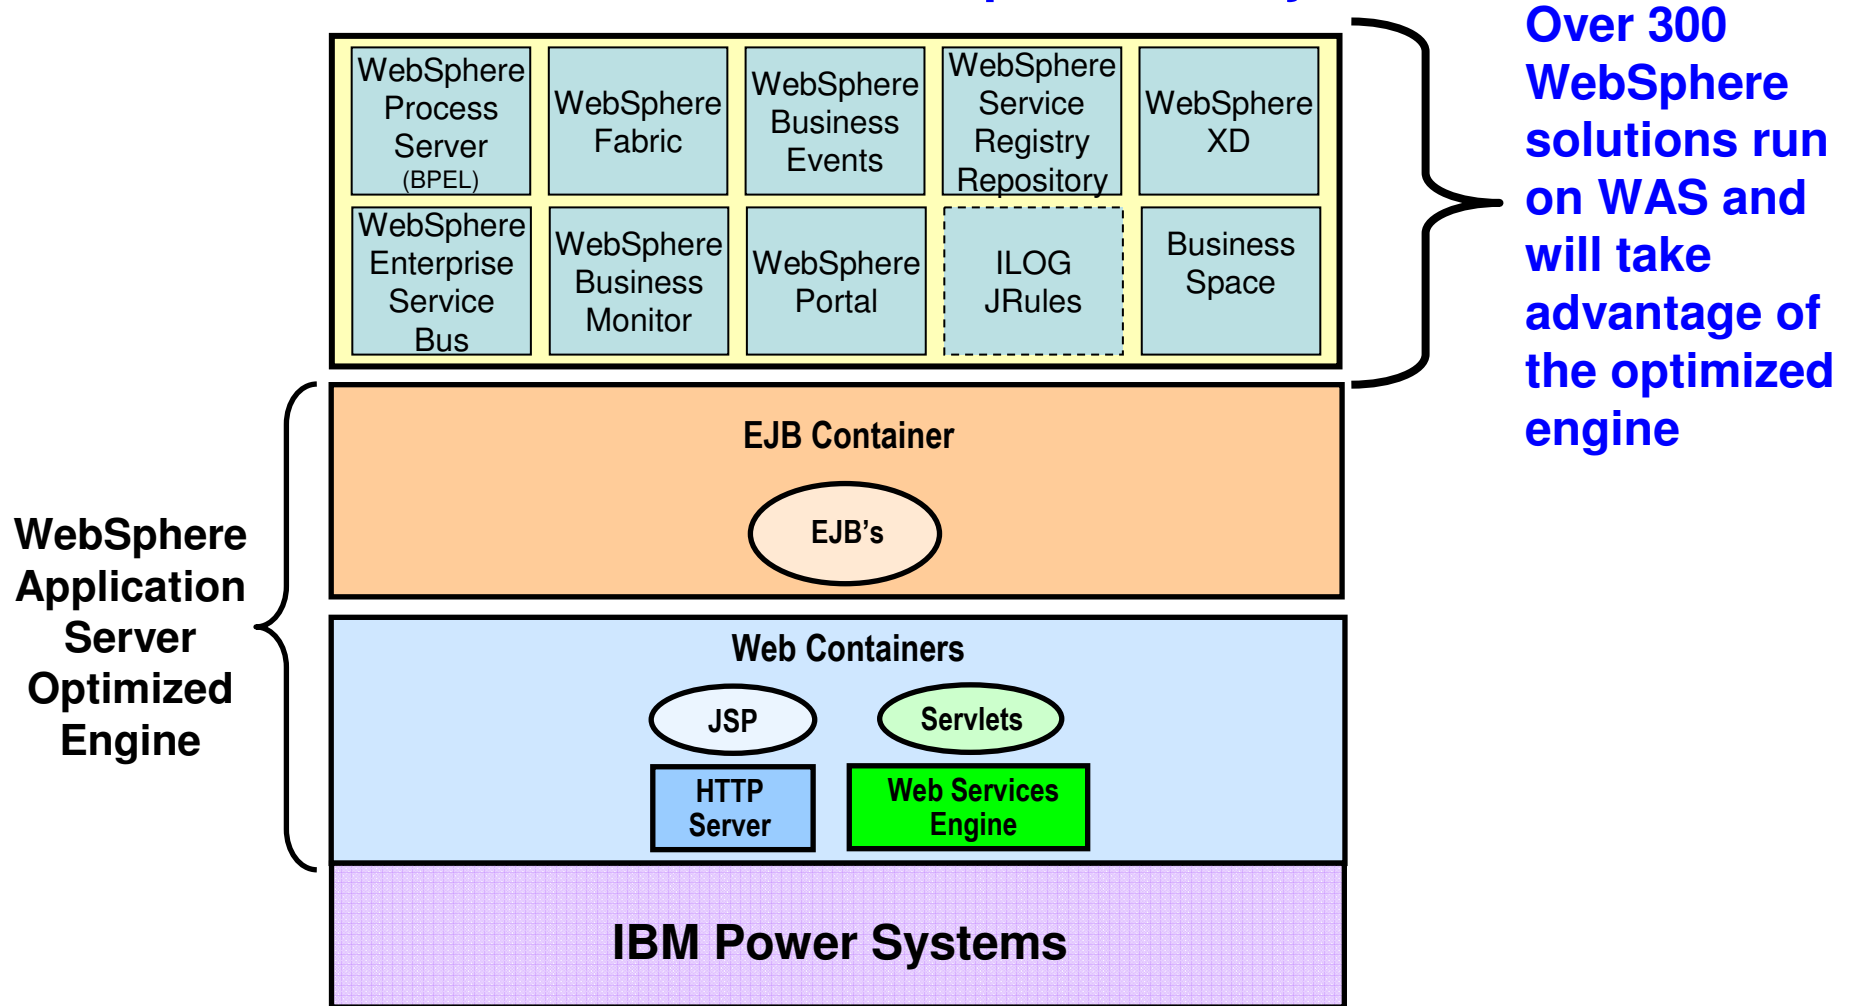
WebSphere takes advantage of the increased number of hardware threads available in Power Systems

Notes:

- **WebSphere Application Server optimizations in Fix Pack 9 (7.0.0.9)**
- **JVM optimizations to be available with WAS Fix Pack 9 as a separate download**
- **AIX V6.1 optimizations (Technology Level 5)**

WebSphere Application Server's Optimized Engine Is The Foundation For The WebSphere Framework

The Fastest Web Engine In The Industry Is The Foundation For The WebSphere Family!



Why Is Processing Efficiency Important?

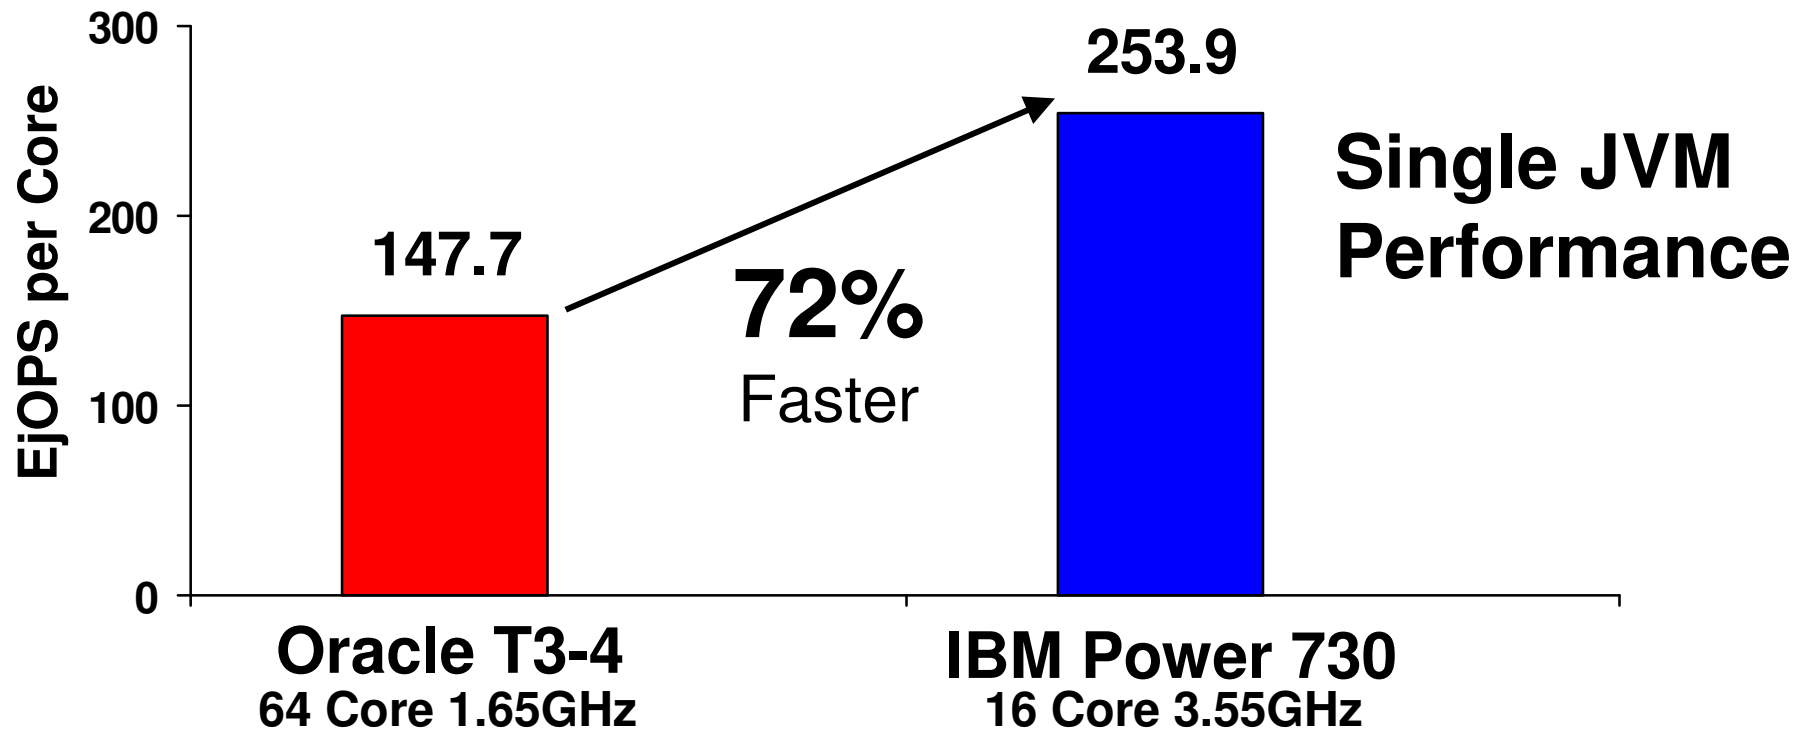
- More performance per core means fewer cores are required
- For a given workload
 - ▶ The hardware will cost less
 - ▶ The software will cost less
 - Commercial licenses are priced on a per core basis

WebSphere optimizations for IBM servers reduce hardware and software costs!

Power 730 vs. Oracle T3-4 SPECjEnterprise 2010 Benchmark Results

World's Fastest Single Server SPECjEnterprise 2010 Performance

SPECjEnterprise2010 performance **results per core**



Source: 2SPECjEnterprise2010 published results: <http://www.spec.org/jEnterprise2010/results/jEnterprise2010.html>

How Does WebSphere On POWER Compare To Oracle's New Exalogic Appliance?

- **Exalogic comes in two different architectures with **slower performance than Power****
 - ▶ EL X2-2
 - Based on Sun Fire x4170 M2 (2.93GHz XEON Class 12 Cores)
 - Currently available
 - ▶ EL T3-1B
 - Based on Oracle T3 Blade (1.6GHz T3-1 SPARC T3 64 Cores)
 - Not available ??


- **Neither Exalogic architecture is **cost competitive****
 - ▶ WebSphere on POWER **performance per core** is superior
 - ▶ WebSphere on POWER **cost per workload** is significantly less



IBM WebSphere Is Optimized For POWER7 Multi-JVM Performance

IBM WebSphere Application Server 7

Red Hat Linux
64-bit
Multi-JVM



IBM Power 730
12 cores / 3.7GHz

13,017 Transactions/second

\$19 per Transaction/second

Competitor Enterprise Application Server

Competitor Linux
64-bit
Multi-JVM



Sun Fire x4170 M2
12 cores / 2.93 GHz

8,354 Transactions/second

\$32 per Transaction/second

Competitor Enterprise Application Server

Competitor UNIX
with Containers
64-bit
Multi-JVM



T3-1B
16 cores / 1.65GHz

10,145 Transactions/second

\$22 per Transaction/second

IBM Online Banking benchmark. JVM selected and number of JVM's configured for optimum performance. No virtualization, except T3 configuration which required containers.

Exalogic Locks Customers In To A Hardware and Software Solution From Oracle That Is **Not Optimized**


- **Forces you to choose between 1/4, 1/2 and full rack – Very Inflexible**
 - ▶ 96, 192 or 360 cores
- **Built on already two slow processor technologies**
 - ▶ Sun Fire x4170 M2
 - ▶ T3-1B Blade
- **“Must” purchase full WebLogic Suite per core**
 - ▶ \$45,000 per core (Includes WLS EE, Web Caching and Coherence)
- **No virtualization support in this first version**
 - ▶ Thus, no way to achieve consolidation savings
- **Architected to have each http server outside Infiniband connections**
 - ▶ Resulting in the http server being a big performance bottleneck
- **No Cloud Support at all**
 - ▶ No virtualization which is fundamental for cloud computing
 - ▶ No Cloud computing management software of any kind

WebSphere Clusters - Proven To Support Demanding Workloads



Wimbledon

- Peak 1 million hits/minute
- 30K simultaneous access to the scoreboard



The IBM employee portal handles 30 million requests a day, maintaining sub-second transaction response times for many applications.

eBay.com

- 1+ Billion page views/day



Schwab.com

- 16.5 million transactions per day



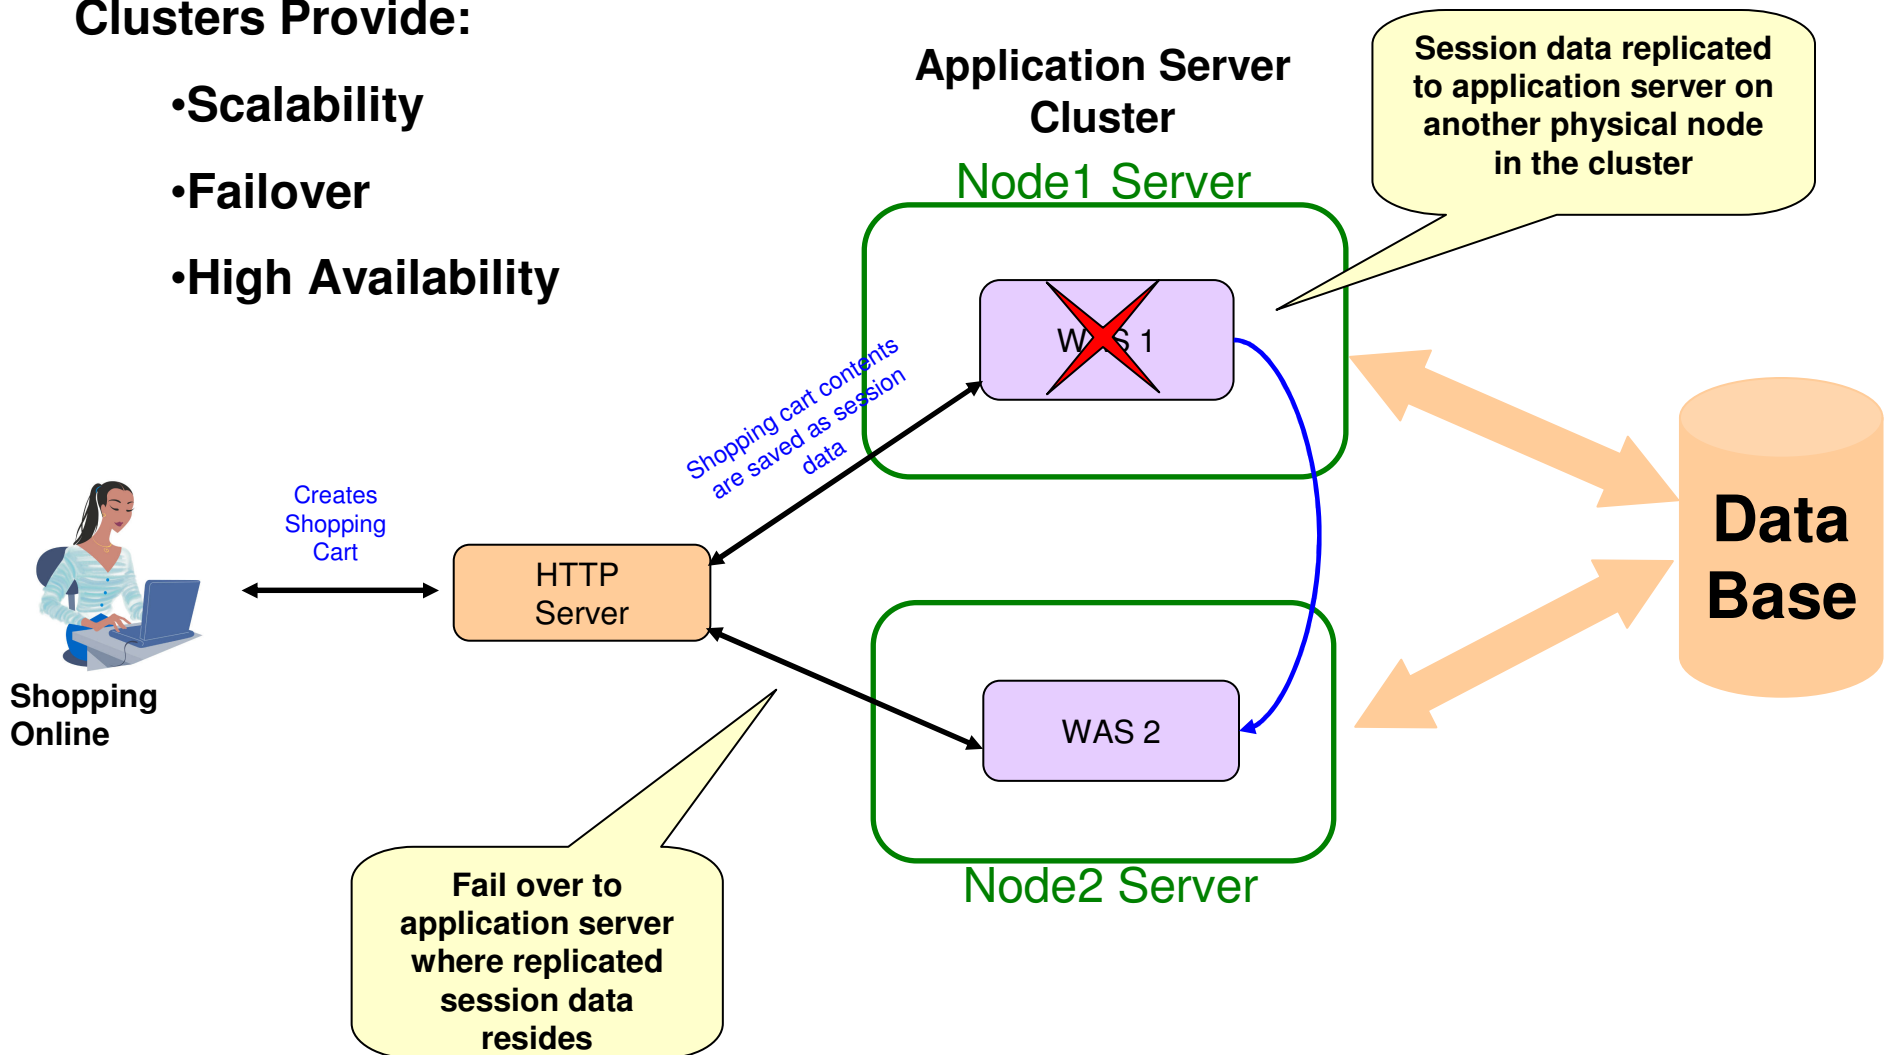
Redbox (\$1 DVD Rentals)

- 14000 Transactions per second
- 2% CPU utilization
- 25K Kiosks
- \$1.3 Million per day in Revenue

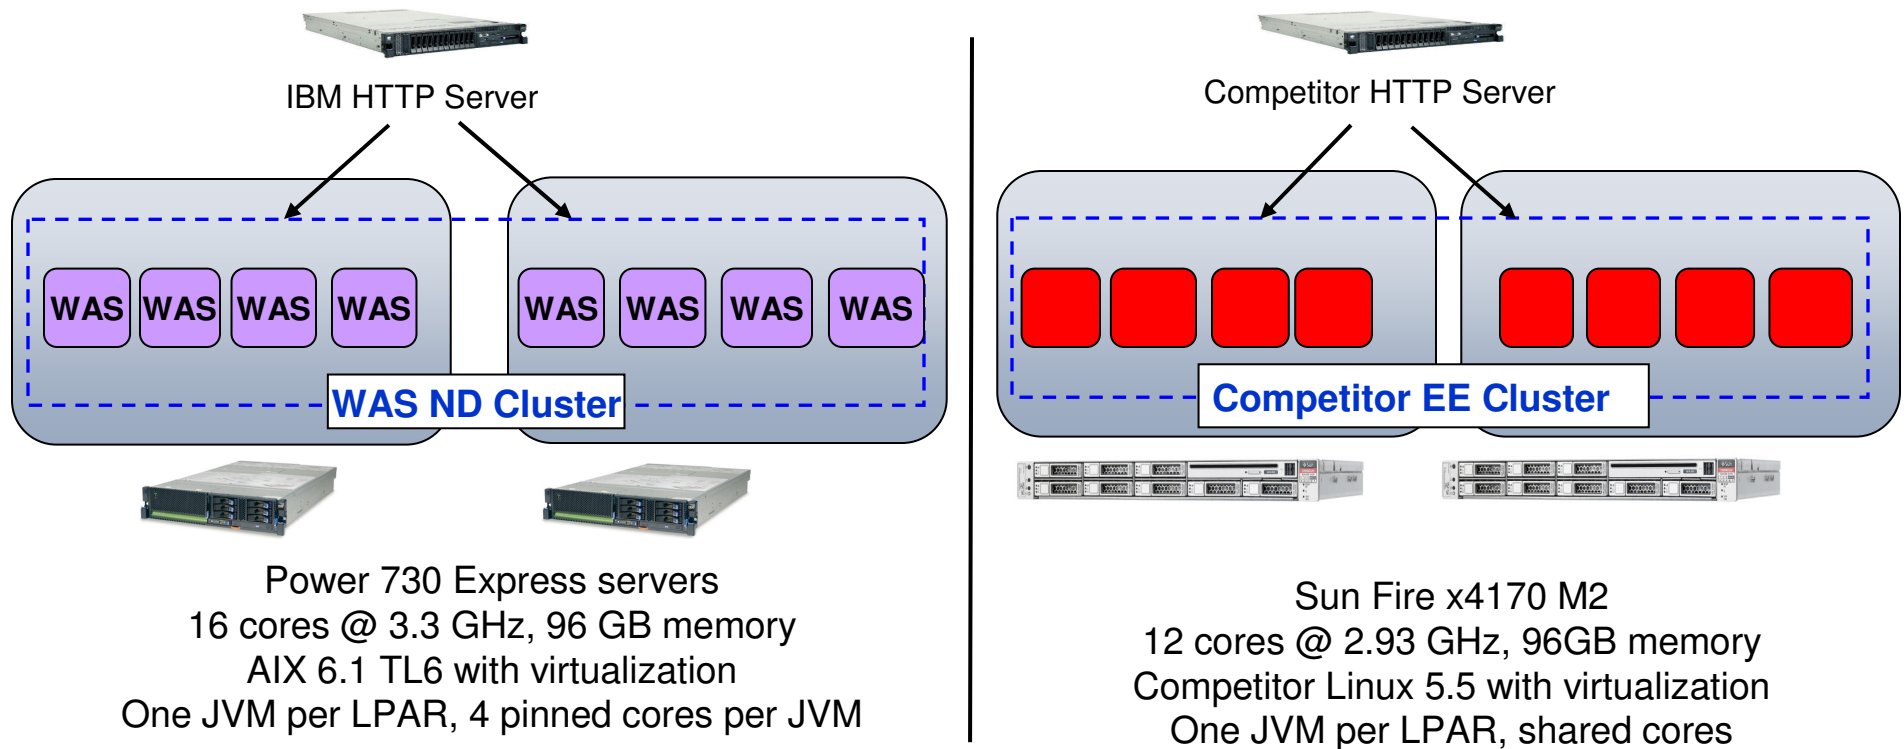
Clusters Offer Horizontal Scalability And Reliability For Web Facing Applications

Clusters Provide:

- Scalability
- Failover
- High Availability



IBM WebSphere Clusters On Power Has Better Throughput Than The Competitor

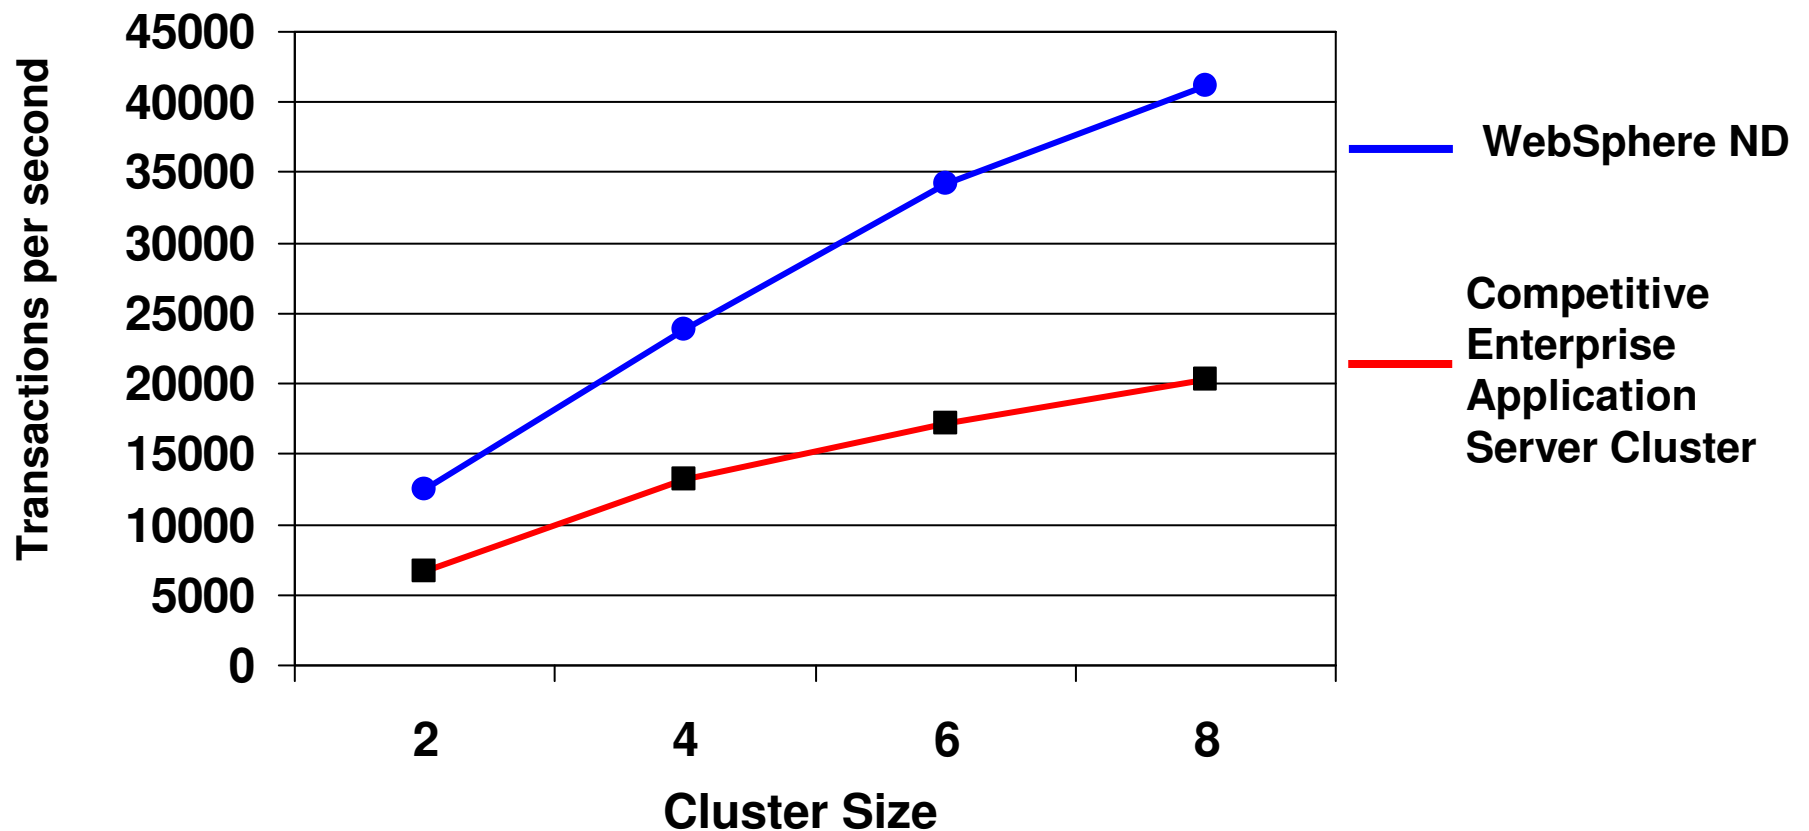


Cluster Size	WebSphere Throughput	Competitor Throughput
2	12,517	6,622
4	23,848	13,225
6	34,265	17,173
8 (pictured)	41,205	20,276

IBM Online shopping benchmark,
add 10 items per cart

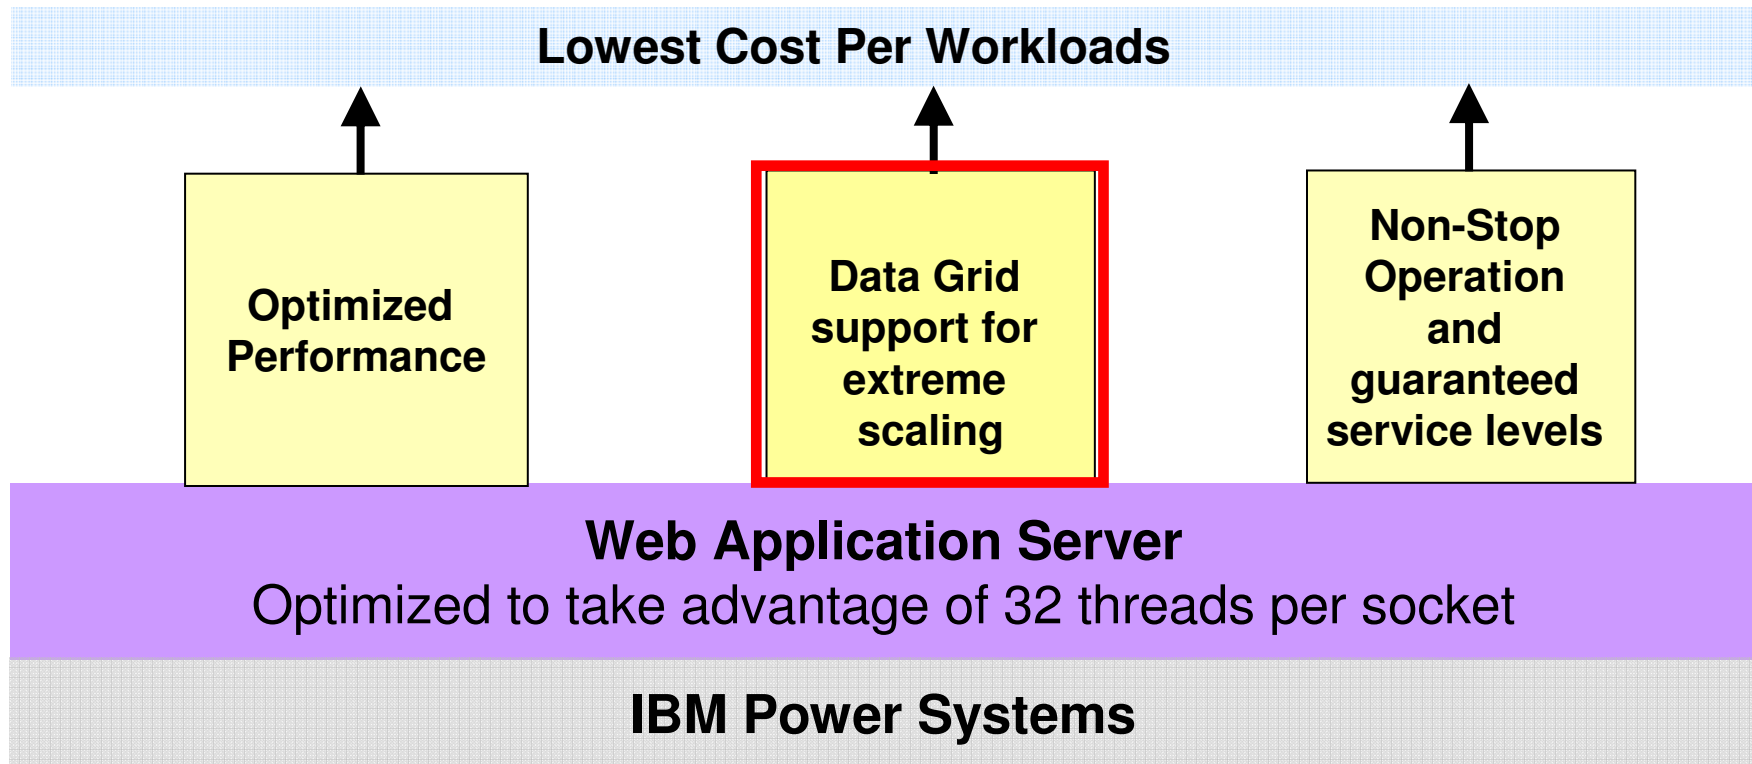
IBM WebSphere Clusters On Power Scales Better Than The Competition

WebSphere clusters provide the highest throughputs and also scale the best.....



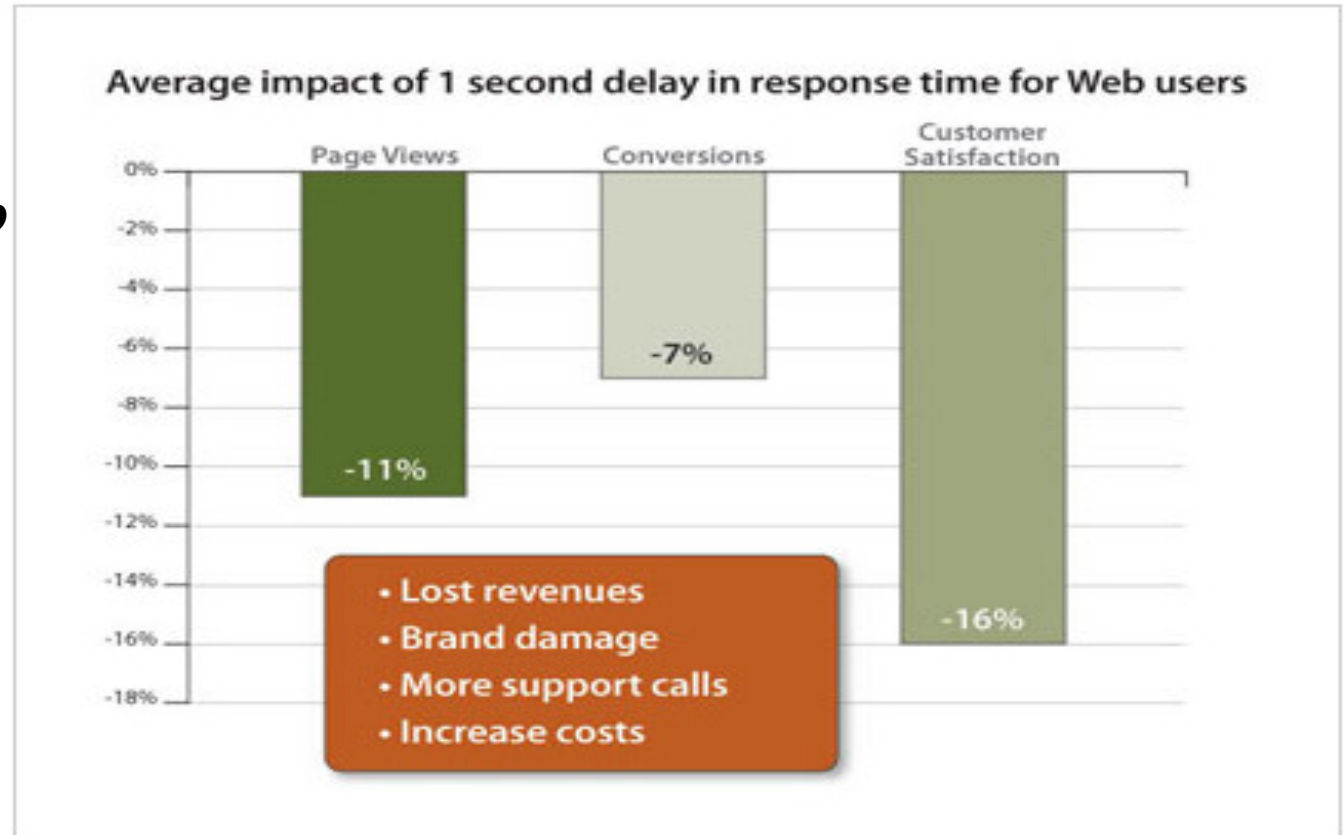
IT Demands More Work From Web Application Servers

Do you have a Web Facing Engine optimized to leverage all of the threads available?



Some Workloads Are Even More Demanding

Aberdeen Group found that an average of one second delay in Web page response time negatively affected page views, conversions and customer satisfaction.




Required: fast response times, even under heavy load

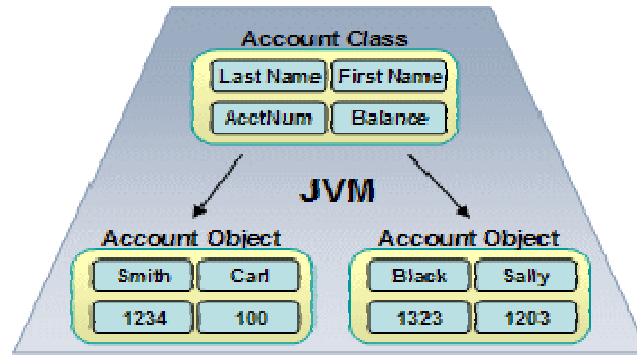
Source: "The Performance of Web Applications: Customers Are Won or Lost in One Second,"

Bojan Simic, Aberdeen Group, November 2008.

What Causes Response Time Delays?



Access to database for reads or updates can cause delays



Smith	Carl	1234	100
Black	Sally	1323	1203
Jones	Bob	9312	33203
....
....

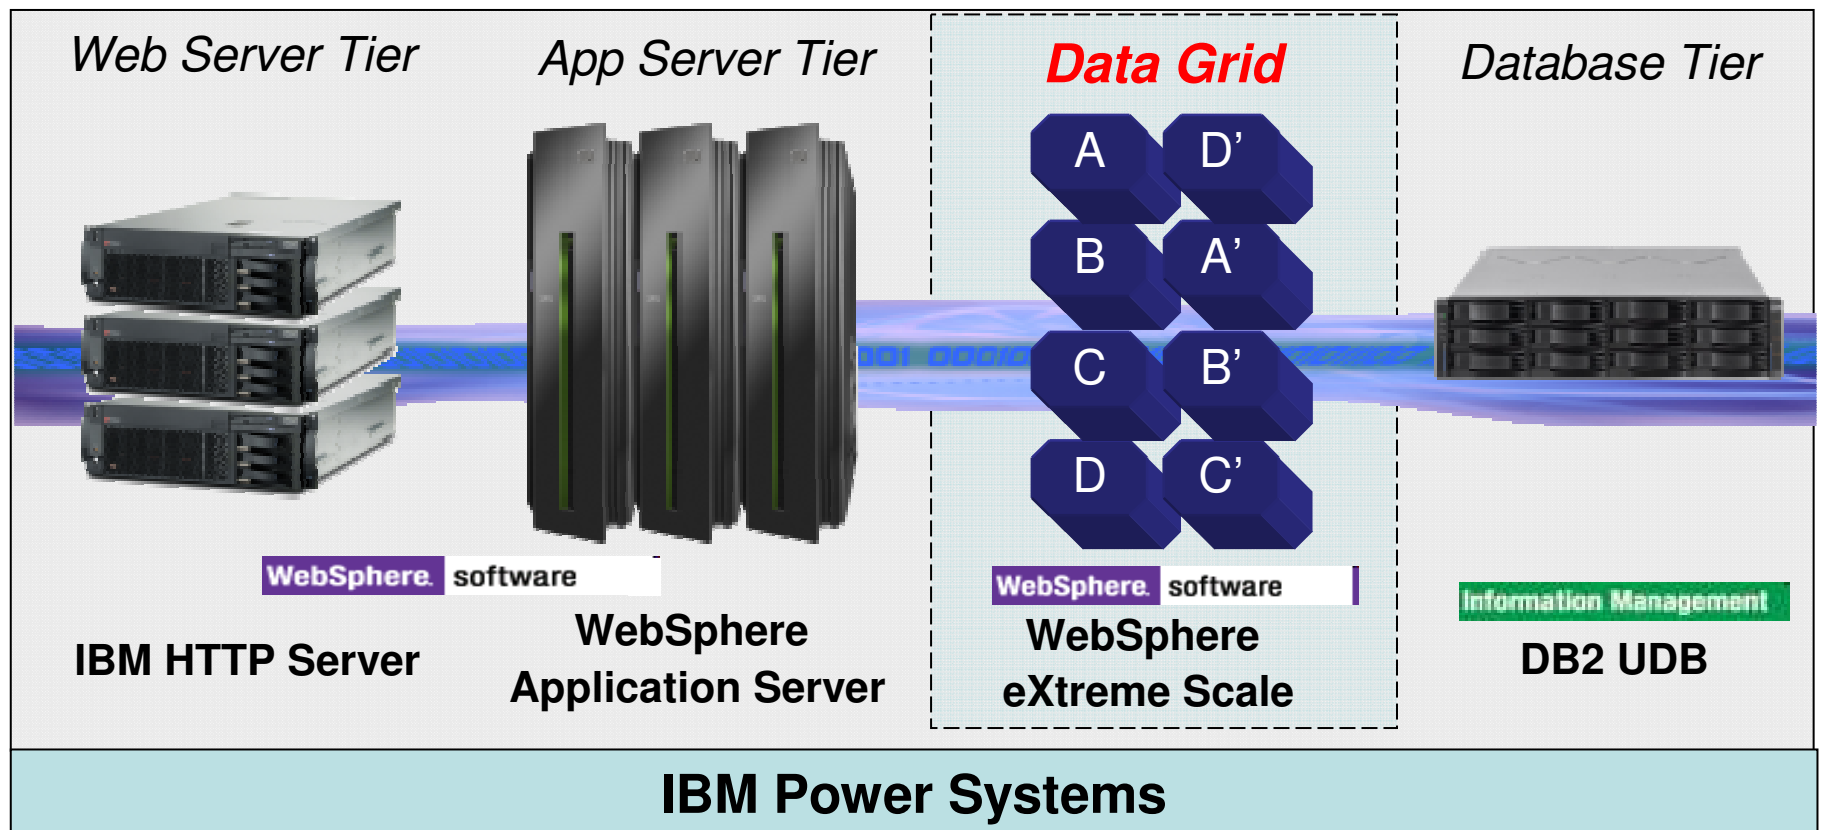
Convert Java object programming model to relational data base rows

Solution: Provide in memory cache to store data in Java object form

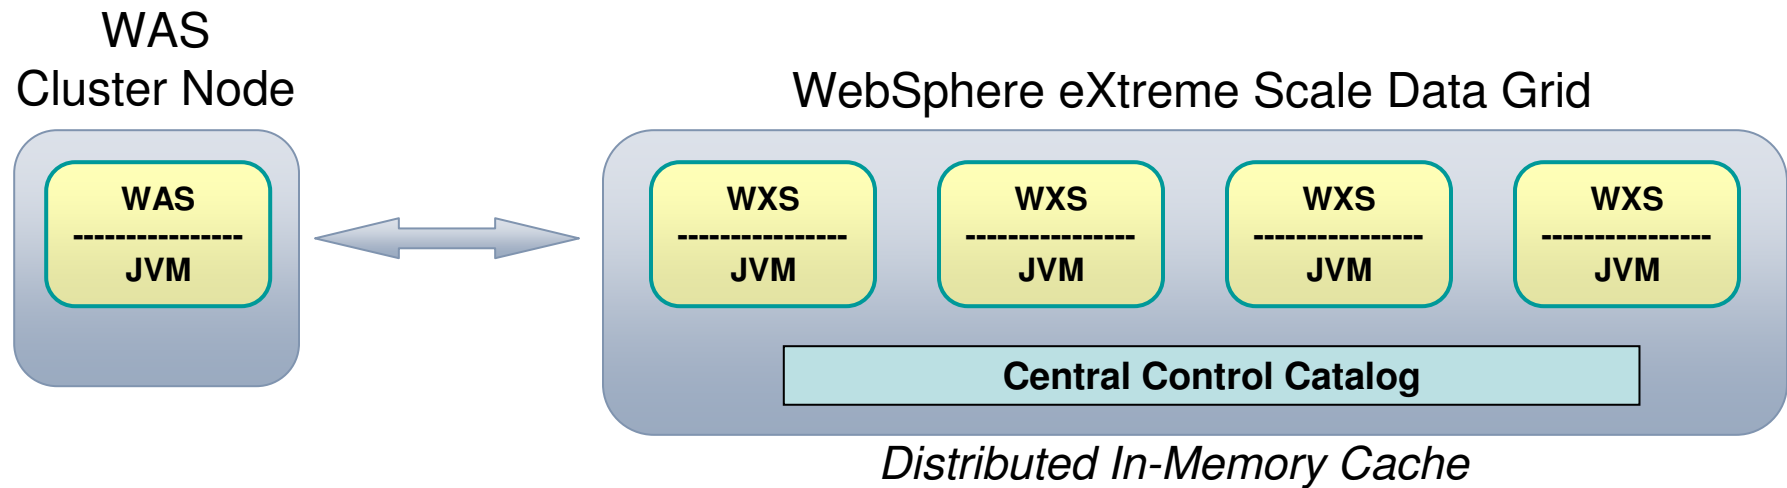
WebSphere eXtreme Scale – Advanced Data Grid For Extreme Scaling

A more efficient caching method to dramatically improve response times in single servers or cluster environments

1. Within single servers or JVM's
2. Across a cluster of servers
3. Across clusters of clusters that are geographically dispersed



WebSphere eXtreme Scale Data Grid Works in LANs And WANs




- **WXS nodes can be clustered to form a “data grid”**
 - ▶ Linear scaling with predictable response time
 - ▶ Stored objects are transparently accessed anywhere in the cluster
 - ▶ Everything is synchronized with a Central Control Catalog
- **Grids can be used in different “patterns” depending on need**
 - ▶ With or without database
 - ▶ For session data
- **Grids can be replicated and synchronized across distance (WAN) for local performance (Oracle Coherence can't do this)**

LAN = Local Area Network
WAN = Wide Area Network

DEMO: WebSphere eXtreme Scale – Increase Application Performance Dramatically


WAS ND V7.0.0.11

WebSphere eXtreme Scale Data Grid V7.1



Demo Condition:

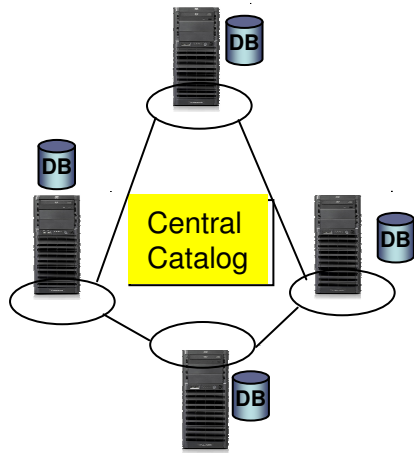
- CPU intensive workload
- Response times reduce and transactions per second increase



Dramatically Improve Performance Across Three Different Patterns

WebSphere eXtreme Scale Works In All Environments – small to extra large

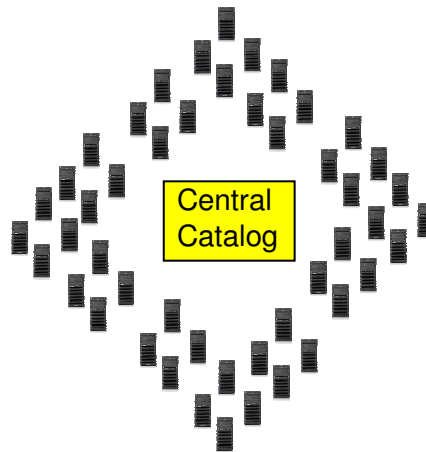
Small Cluster (LAN)



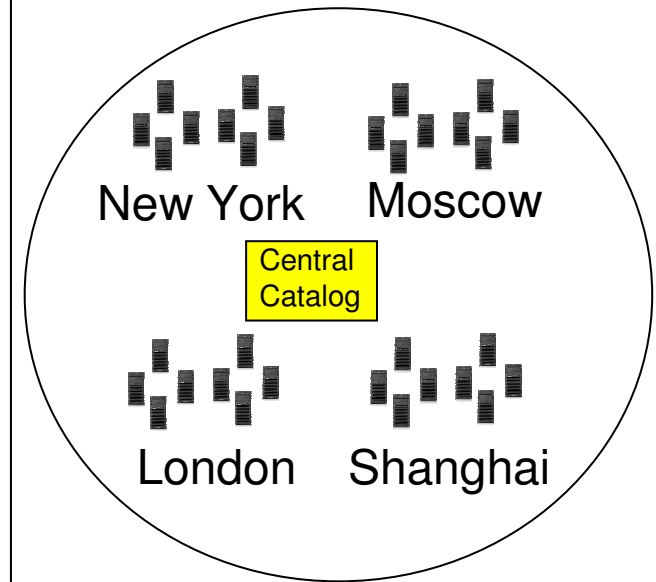
Memory shared across Servers to reduce DB access. Managed by Central Catalog

LAN = Local Area Network

Large Cluster (LAN)



Geographically Dispersed Clusters (LAN and WAN)



WAN = Wide Area Network

A Large Investment Bank Increases Transactions By 40 Times With WebSphere eXtreme Scale

Investment Banking

12 Million
orders per day

4x
increase in
revenue

40x
number of
transactions
supported

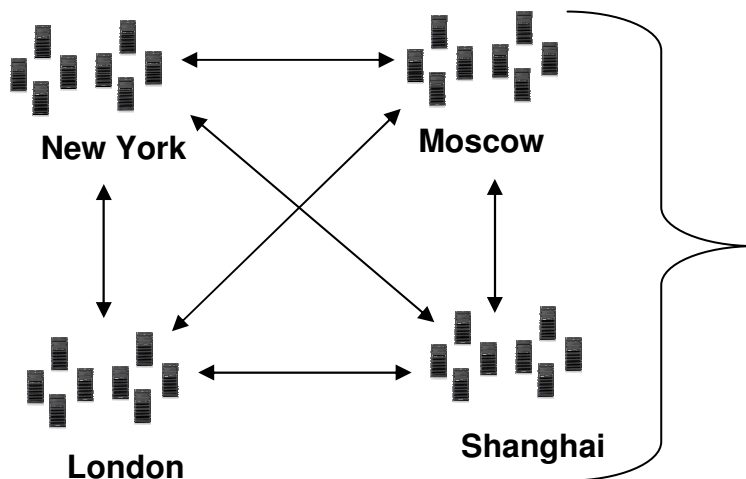


Next-generation Order Management System with WXS

- **Before:** Oracle RAC based architecture unable to scale to necessary demands
- **After:** 300K transactions / day → 12M / day
- Revenue up 4X and growing ... “all because of WebSphere eXtreme Scale”
- Response time drops to 2.5ms
- Moving to “22 x 7” operations (more than 9AM - 4PM)

Oracle Exalogic Cannot Support Geographically Dispersed Data Grids

Oracle Coherence does NOT have a central catalog
- Coherence replication model is flawed



Point to point multi-casting will not provide consistent replication across geographically dispersed clusters

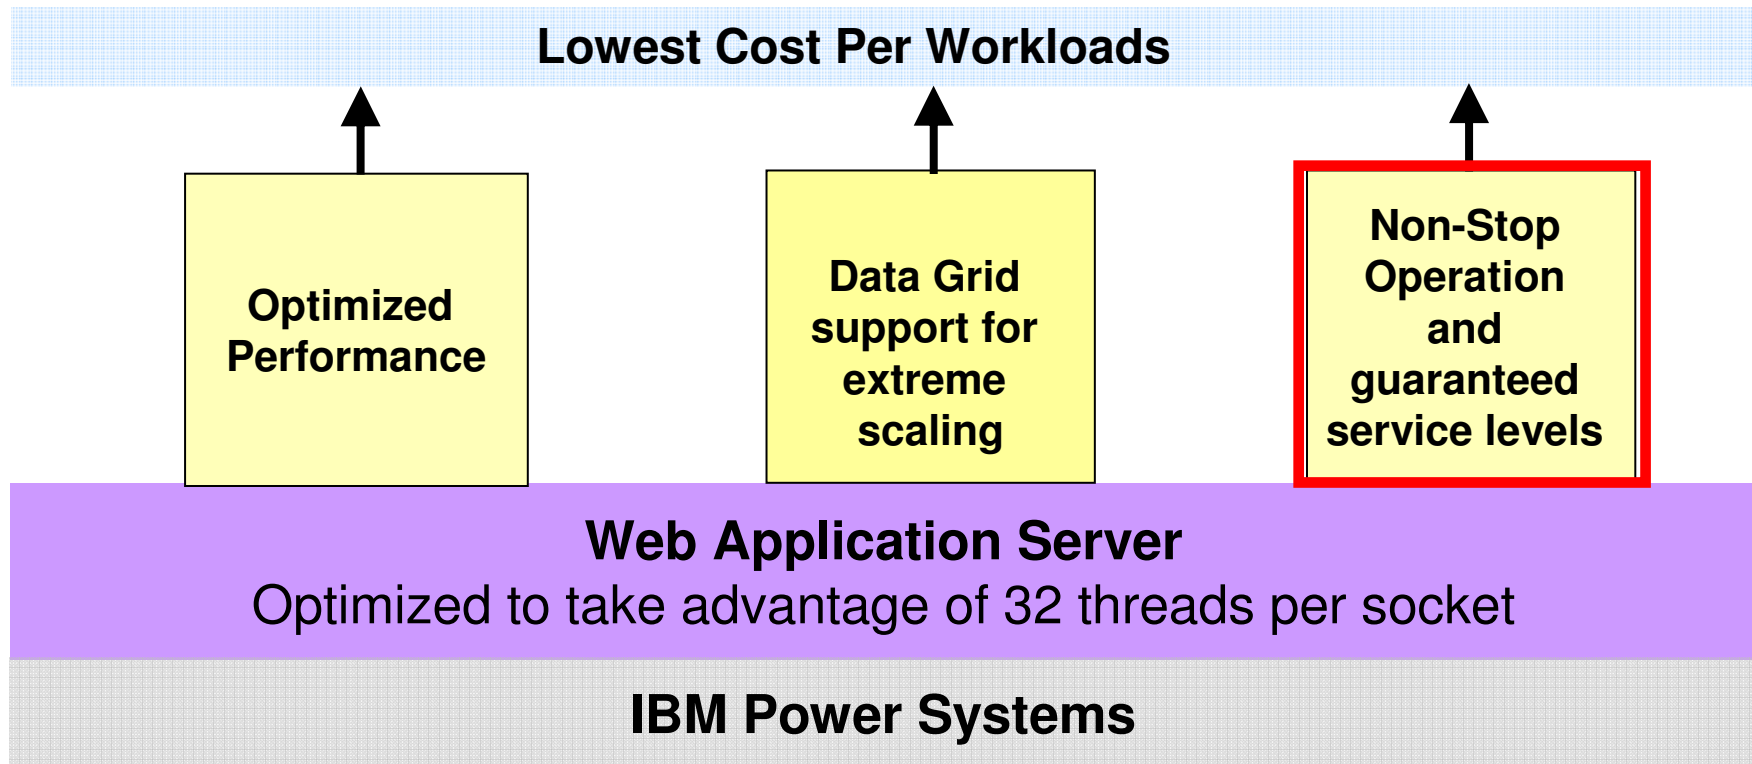
- No central catalog backup
- No verification of replication between clusters

WXS fully supports Geographically Dispersed Clusters

- Supported via request – receive protocol support like TCP/IP

IT Demands More Work From Web Application Servers

Do you have a Web Facing Engine optimized to leverage all of the threads available?



WebSphere Virtual Enterprise - Assures High Availability And Enforces Service Policies

1. Non-stop operation

- ▶ Continuous availability during application maintenance



2. Performance Management

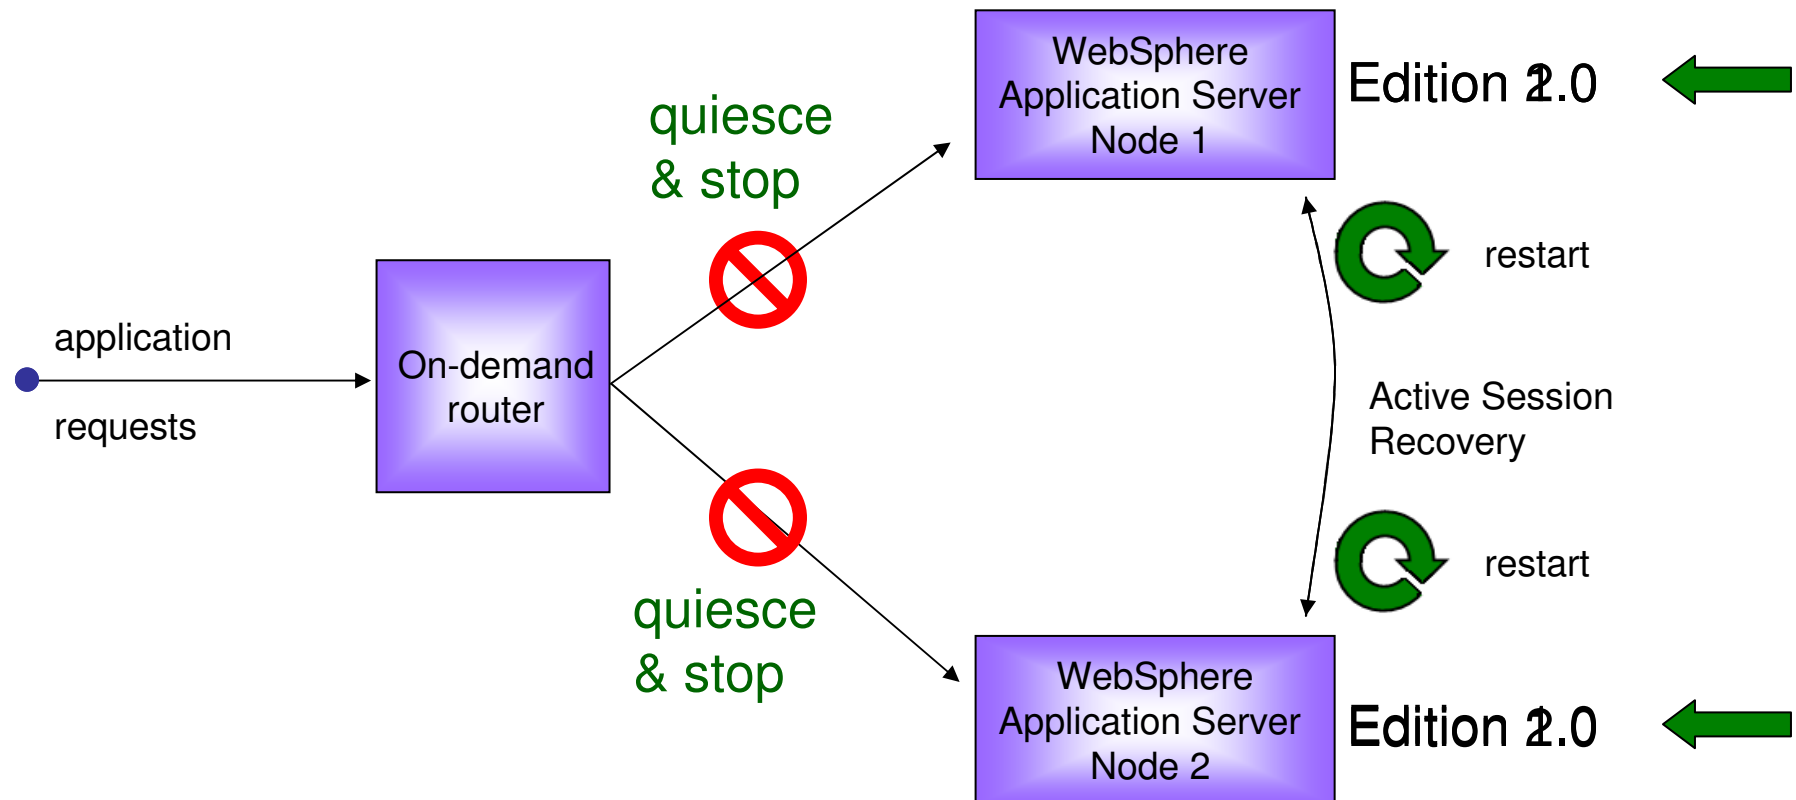
- ▶ Define **service policies** based on **response time goals**
- ▶ **Dynamic cluster** capability maintains response time objectives **despite variable workload** demands



DEMO: WebSphere Virtual Enterprise

Non-stop Operation

- Deploy new applications with lower risk of losing service.
- Deploy application versions without interruption.



Create **Service Policies** That Provide Guaranteed Highest Quality Of Service

- Define service level goals with service policies
- Service policies specify the **response time goals** and the **relative importance** of the service policy relative to other service policies
- Application requests are mapped to service policies based on rules

The screenshot shows the 'Service Policies' management interface. It includes a navigation sidebar on the left and a main content area. The main content area has a title 'Service Policies' and a description: 'A Service Policy defines a business goal and an importance, and contains one or more Transaction Classes. The Service Policies define an Operational Policy which is used by a component in the Proxy Server to categorize and filter work in the queue.' Below this is a 'Preferences' section with 'New' and 'Delete' buttons. A table lists the policies:

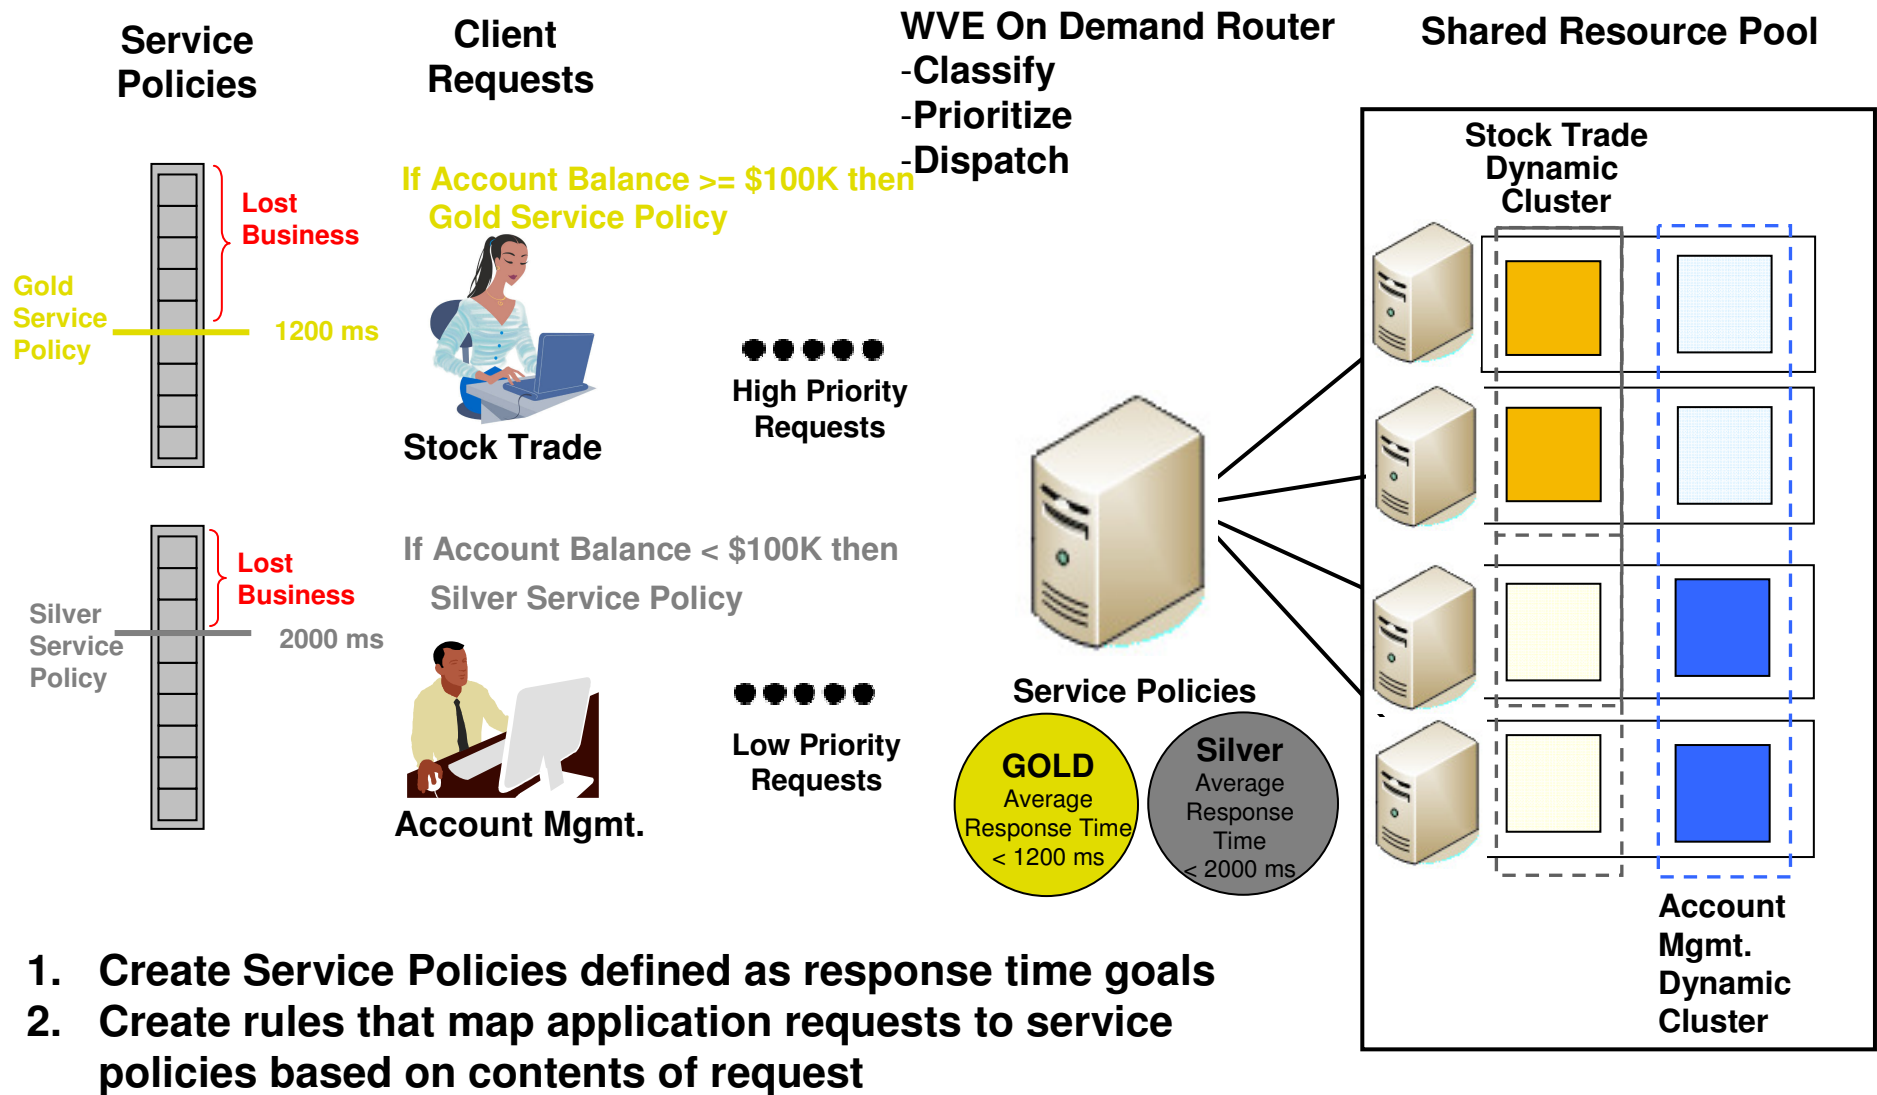
Select	Name	Importance	Goal	Description
<input type="checkbox"/>	Default SP		Discretionary	
<input type="checkbox"/>	Gold SP	High	Avg response 15 Seconds	Gold Service Policy
<input type="checkbox"/>	Platinum SP	Highest	Avg response 1500 Milliseconds	Highest SP

The 'Platinum SP' row is highlighted with a red circle around the 'Highest' importance and a blue circle around the '1500 Milliseconds' response time goal. The table footer shows 'Total 3'.

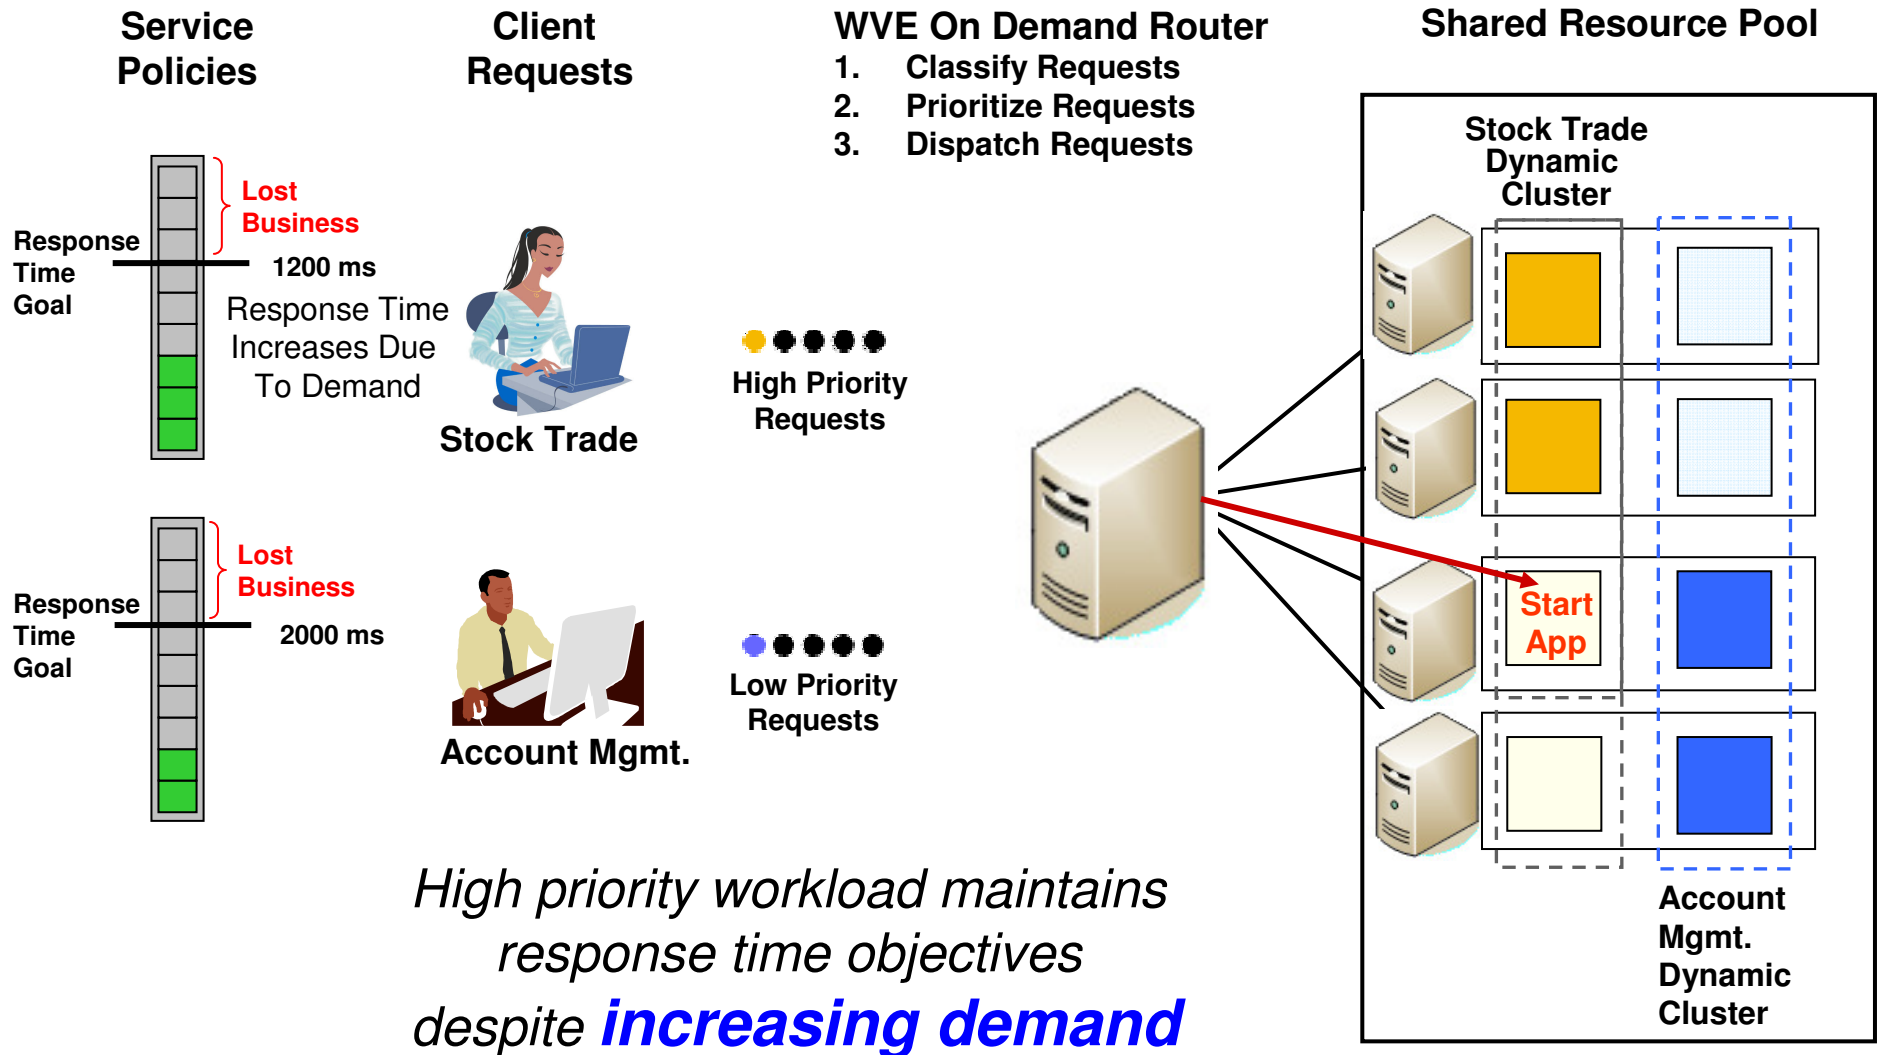
Service Policies define the **relative importance** and **response time goals** of application services

WebLogic cannot do this! Exalogic cannot do this!

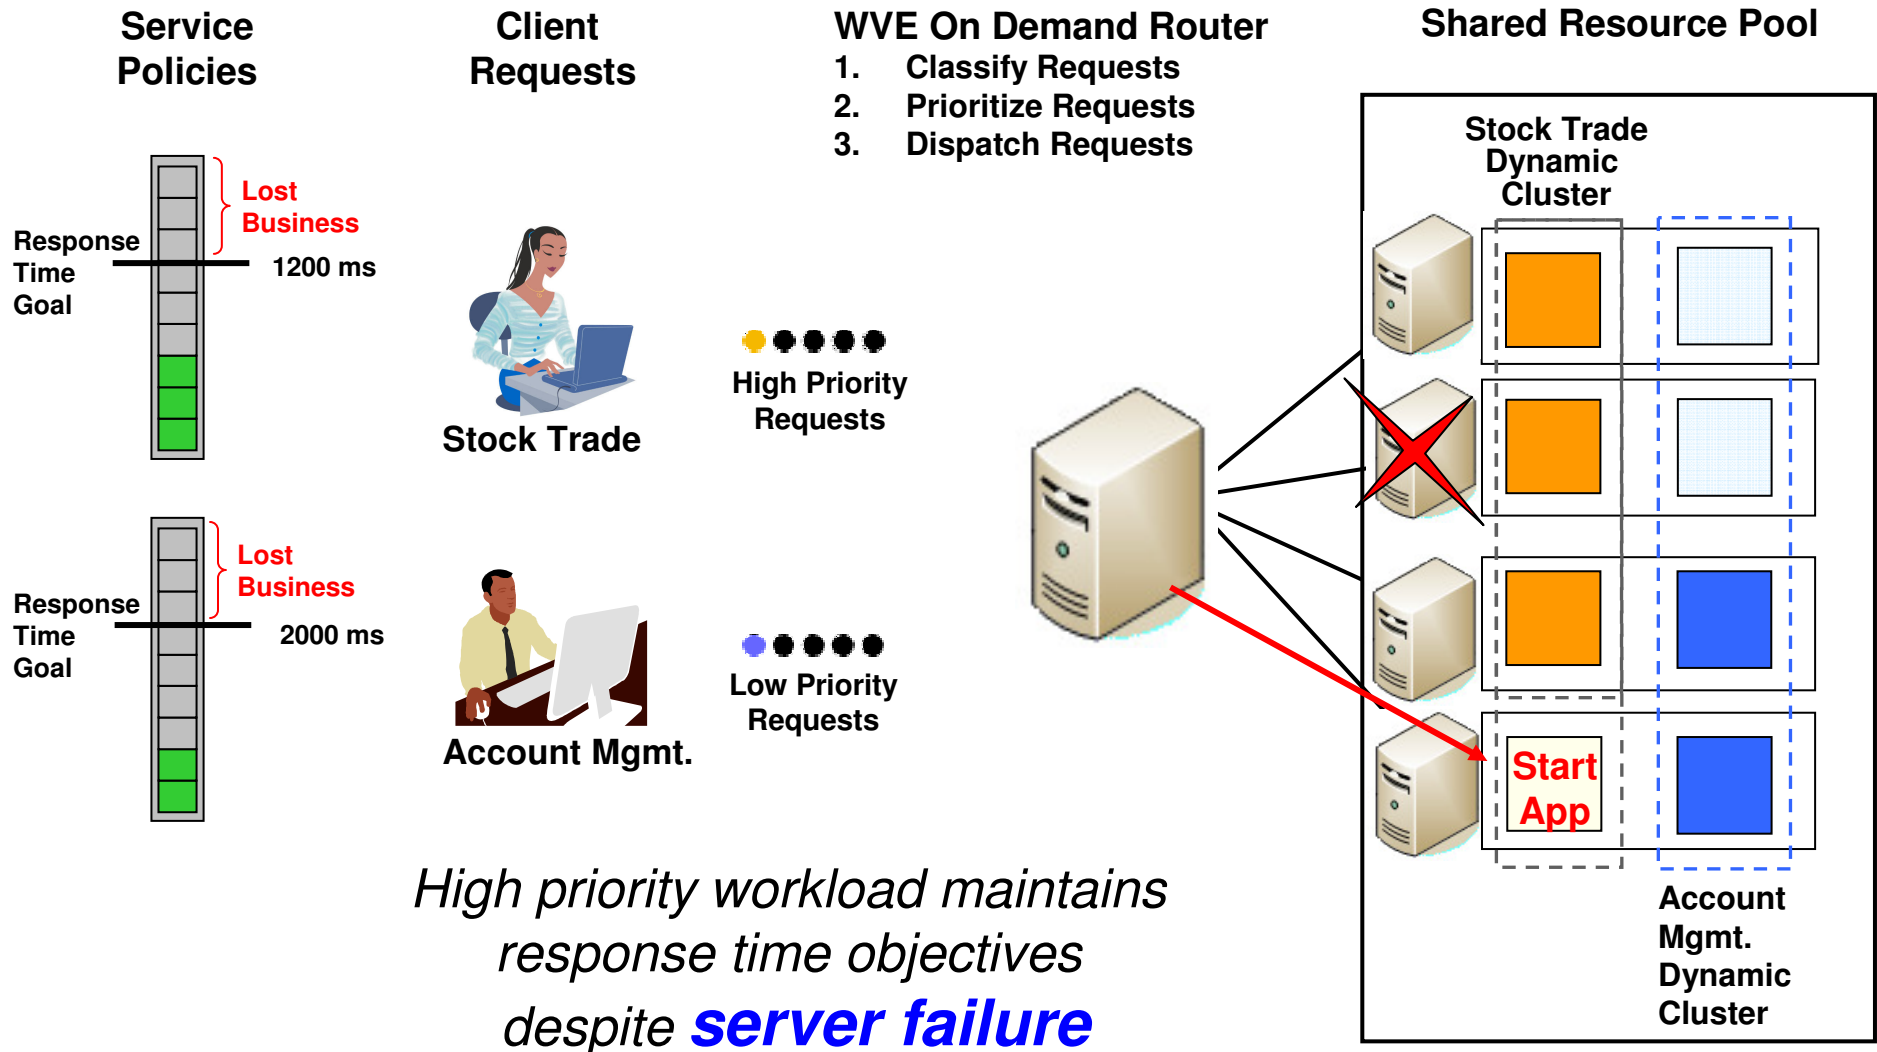
Create Service Policies That Provides Best Service For Most Valuable Customers



WebSphere Virtual Enterprise - Maintain Response Time Objectives



WebSphere Virtual Enterprise - Maintain Response Time Objectives



WebSphere –The Best Optimized Web Engine For Your Business At The Lowest Cost

	IBM	Oracle Exalogic
Best Single JVM Performance at the lowest cost	YES	NO
Best Multi-JVM Performance at the lowest cost	YES	NO
Best CLUSTER Performance at the lowest cost	YES	NO
Data Grid Scaling across Geographically Dispersed Clusters	YES	NO
Service Policy Management to assure best cluster performance	YES	NO