

## 인 메모리 데이터 그리드 21세기 컴퓨팅의 새로운 기술

발표자 : Brad Buff

직함 : WW Sales – WebSphere Application Infrastructure



# Application Infrastructure Trends

**Rapid application development and delivery** driving simplified, integrated and automated development and operations lifecycles

**Explosion of mobile, social and cloud applications** driving new demands on middleware infrastructures

The combination of huge transaction volumes against massive amounts of data with little tolerance for delays is driving the need for **elastic caching** technologies

Use of **cloud delivery models** to provide elasticity, scale, multi-tenancy and context across different form factors and access methods

# Driving Application Innovation

---

## ● Application Server

- Rapidly create and deliver secure, reliable and high performing modular applications
- OLTP, Batch, Mobile Web, Web 2.0 and Communications enabled applications

## ● Elastic Caching

- Consistent application and transaction response times, even as data grows with capability for linear scaling and fault tolerance

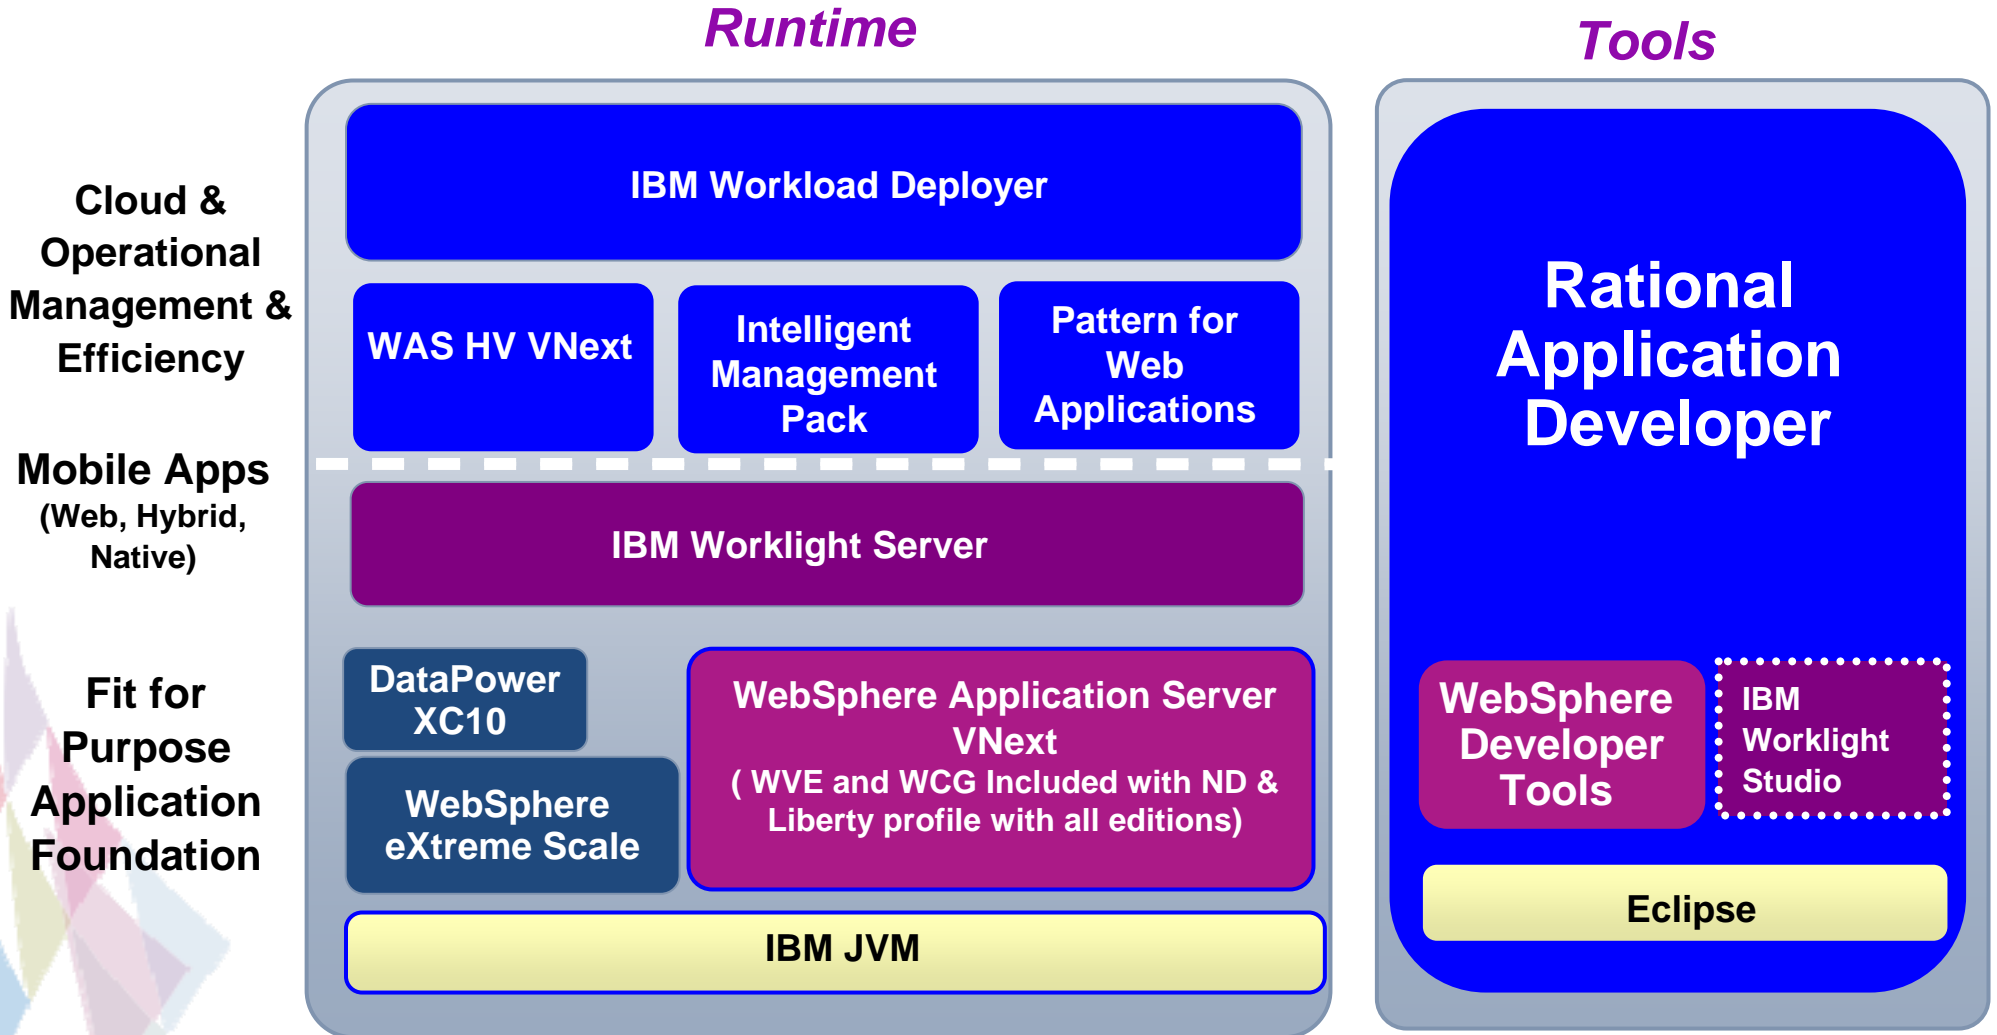
## ● Mobile

- Extend existing business capabilities to mobile devices
- Manage mobile devices and Secure mobile business
- Transform the business by creating new opportunities

## ● Cloud

- Improved IT & business efficiency
- Lower operational & energy costs by maximizing utilization of resources
- Rollout new services faster
- Superior customer service by dynamically delivering needed resources to highest priority apps

# WebSphere Application Infrastructure



# Elastic Caching solves the problem of Transaction Overload

*Web Server Tier*



*App Server Tier*



*In memory data grid*

1



DataPower XC10 for simple data oriented scenarios.

2

**WebSphere. software**

eXtreme Scale for maximum flexibility and sophisticated application oriented scenarios.

*Back-end Systems Database Tier*



# In-Memory Computing Is Key for Cloud-Native Applications

## What Is In-Memory Computing?


A computing style by which it can be assumed that:

- The primary data store is in-memory ("RAM is the new disk")
- Data access latency is negligible
- Terabyte-size datasets can be stored in the computer RAM
- Spinning disk is "the new tape" (overflow, recovery)



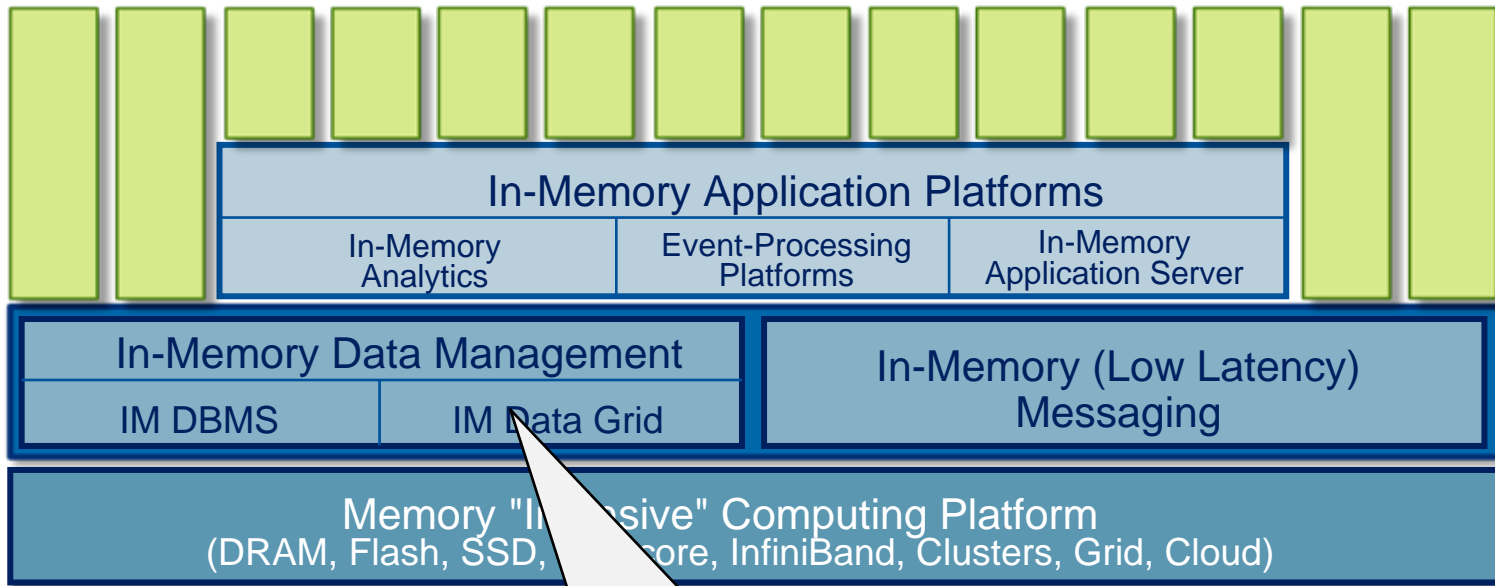
## Why Now?

- 64-bit processors can address up to 16 exabytes of data
- DRAM prices drop by 30% every 18 months
- 1Gb of NAND flash memory average price is 82 cents\*
- Commodity blades provide 1 terabyte of DRAM (or more)
- Multicore CPUs enable parallel processing of in-memory data
- In-memory-enabling software is amply available and proven!



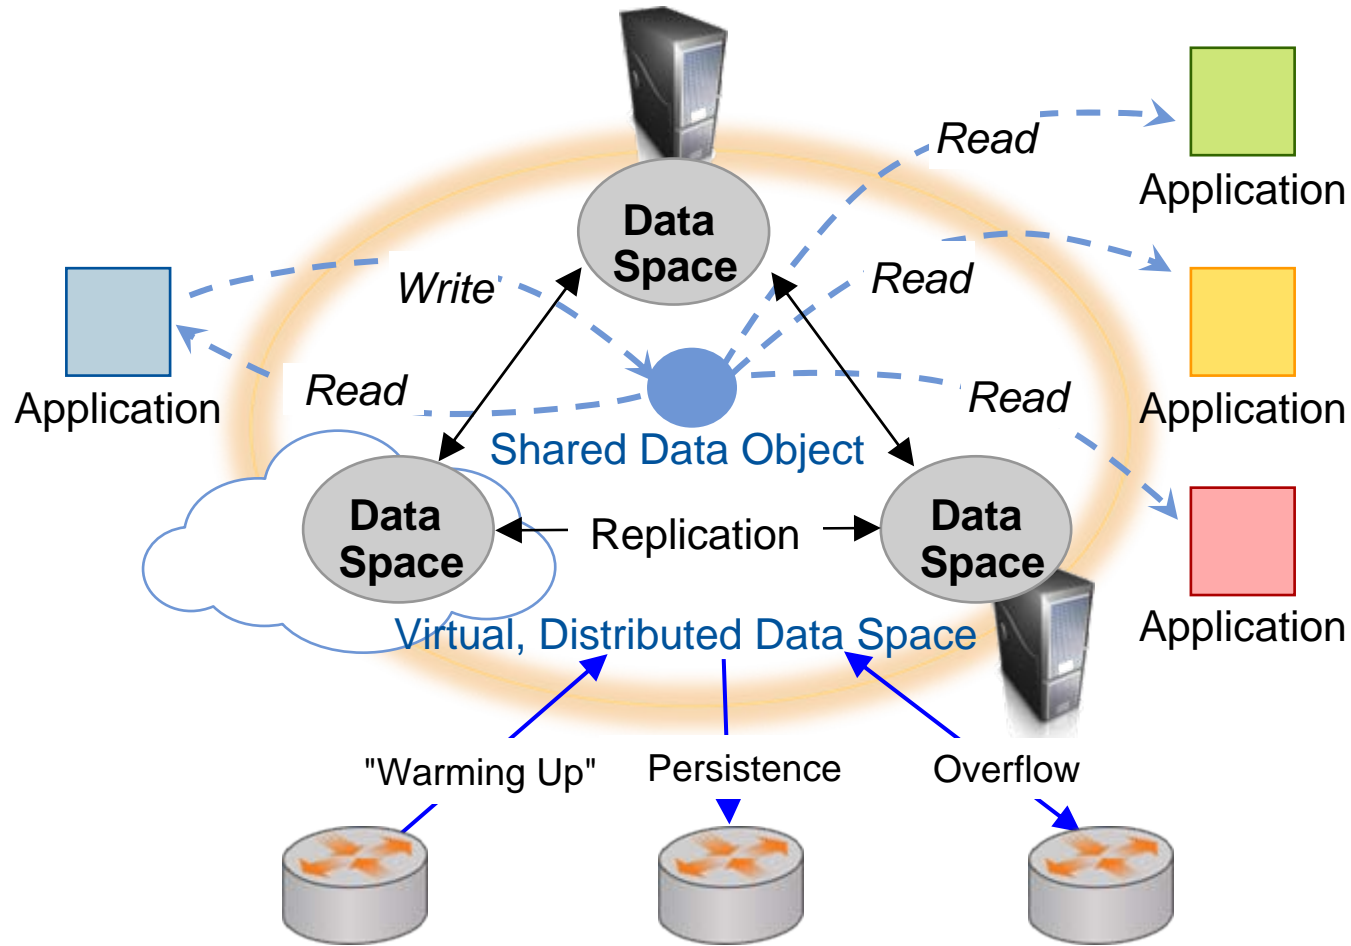
\* Per Gartner's "Weekly Memory Pricing Index, 7 October 2011," G00226221

# Gartner Taxonomy of "In-Memory Computing" Technologies



IBM WebSphere Extreme Scale;  
Oracle Coherence

# In-Memory Data Grids





# IMDGs: What are they? What are they used for?

- **In-memory, distributed and replicated object oriented data store**
  - High performance
  - Scalability
  - High-availability
  - Shareable across multiple applications (transaction management, locking etc.)
  - Persistence, warm-up and overflow management (integration with DBMS)
- **Simple (relatively) APIs**
  - Put, Get, Update, Delete
  - SQL-like query language (map/reduce, grid-like applications)
  - Eventing mechanism
- **Utilization**
  - High performance data store for low-latency oriented applications
  - Performance/scalability booster for established web-applications
  - Mainframe/DBMS offloading
  - Enabler for application infrastructure products (e.g., ESB, BPM, CEP, BI etc.)
  - Cloud-computing enabler

# IMDGs Capabilities

- Partitioning
- Replication (Synchronous/Asynchronous)
- APIs, Programming Languages, Query Languages
- Remote Agents
- Eventing Framework
- Transaction Management
- Persistency Management (Read Through, Write Through, Write Behind)
- DB/DG Synchronization
- Warm-up Management
- Overflow Management (Eviction, Expiry, Persistency)
- Security
- Development
- Monitoring/Management/Administration

# IMDGs User Benefits/Challenges

- **Benefits**

- High performance transactional data store (not a DBMS!)
- High scale (hundreds of nodes, multi-terabytes)
- Relatively simple programming model (CRUD)
- Distributed computing complexity is (mostly) hidden to developers
- Transactional and analytical applications support (e.g., query language)

- **Challenges**

- Lack of commonly agreed standards (=vendor lock-in)
- Skills availability
- Architecture tuning (e.g., Java garbage collection issue)
- Testing/Debugging
- Monitoring/Management

# In-Memory Data Grids Will Be a Pervasively Deployed Technology

By 2014, at least 40% of large organizations will have one or more advanced in-memory data grid products in use (an increase from fewer than 10% in 2011).

## Why advanced IMDGs will be pervasive:

- IMDG supports short-term requirements, such as Web application performance boosting, session data management, database and mainframe offloading.
- IMDG enables advanced scenarios, such as in-memory computing, SaaS, social networks and cloud computing.
- IMDG technology comes integrated in various software products, including OSs, application platforms, portals, CEP/BAM, BPMS, ESBs and packaged applications.

## Why it may not happen:

- Technical complexity of the technology determines a chronic scarcity of skills
- High costs narrow adoption to leading-edge enterprises, and only to the most-demanding and complex scenarios.
- Lack of common standards prevents mainstream adoption because of lock-in risks.

# Innovative Elastic Caching Solutions



## DataPower XC10 Appliance

- Drop-in cache solution optimized and hardened for data oriented scenarios
- High density, low footprint improves datacenter efficiency

*“Data Oriented”*

Session Management  
Extension for DynaCache  
Side cache  
Worldwide cache  
Database buffer  
Business Event Processing  
Petabyte analytics  
Extreme Transaction Processing

*“Application Oriented”*

*Elastic caching for linear scalability*

*High availability data replication*

*Simplified management, monitoring and administration*



## eXtreme Scale

- Ultimate flexibility across a broad range of caching scenarios
- In-memory capabilities for application oriented scenarios

# WebSphere eXtreme Scale

WXS provides an in-memory data grid which, dynamically processes, partitions, replicates, and manages application data and business logic across hundreds of servers

- Automatically handles replication which can be either synchronous or asynchronous
- Handles advanced placement so that replicas can be placed in different physical zones
- Fully elastic in that servers can be added and removed and it automatically redistributes data
- Allows clients using different object representations to share data stored in the data grid
- Provides automatic integration with databases
- Provides HTTP Session Management
- Flexible deployment model allowing significant customization
- Proven multi-data center capabilities
- Proven low-latency access to data



# IBM WebSphere DataPower XC10 Appliance

*The XC10 appliance provides a complete, purpose-built, easy-to-use solution for common distributed caching scenarios*

- Allows businesses to leverage the value of existing infrastructure investments
- Provides “drop-in” use for Http Session Management, Extended DynaCache service, and Simple Grid / ESB Caching scenarios, requiring little or no code changes to existing applications
- Offers the ability to quickly and easily increase cache capacity and throughput as needs grow
- Includes a flexible and simple user management interface for monitoring and administration
- Provides a large 240 GB cache with near-linear scalability
- Reduces risk of data loss via automatic replication, delivering high availability and fault tolerance for higher user satisfaction and faster task completion



# New: Announcing WXS 8.5 and XC10 V2.1

## *In Memory Data Grid for elastic scalability: WebSphere eXtreme Scale v8.5*

- **Support for elastic caching for WAS Liberty** brings scalability, fault tolerance and high availability
- **Dynamic cache replacement** for Portal content rendering improves throughput and performance.
- **Support for Install Manager** in WAS & Liberty environments reduces deployment time and costs
- **Grid Query Console** allows monitoring, inspection, and invalidation of specific cache content

## *Purpose-built elastic caching appliance:*

### *WebSphere DataPower XC10 Appliance v2.1*

- **Multi-data center support** allows customers to host data on XC10s in multiple locations with data kept in synch through multi-master replication
- **Support for elastic caching for WAS Liberty** provides scalability, fault tolerance and high availability
- **Dynamic cache replacement** for WebSphere Portal
- **Monitoring enhancements** ease administration and improve serviceability



*Improved TCO, increased operational efficiency and productivity, and better response time*





# Forrester TEI of WebSphere eXtreme Scale



A Forrester Total Economic Impact™ Study Prepared For IBM


## Total Economic Impact™ Of IBM WebSphere eXtreme Scale

February 2012

The financial analysis found that the organization experienced:

- ROI of 123%
- Payback Period of 9 months
- Net Present Value of \$5.9 million

Three-Year Risk Adjusted Analysis



Source: Forrester Research, Inc.

- **Benefits.** The organization Forrester interviewed experienced the following benefits:
  - **Reduction in hardware and software costs.** This benefit represents the hardware and software savings associated with eliminating the need to expand to additional databases.
  - **Annual ongoing staffing costs.** This represents the savings from the ongoing maintenance of additional databases.
  - **Incremental gross revenue (Not quantified)** This benefit represents the incremental revenue associated with the mitigation in user drop-off when users experience slow response resulting from a large surge in traffic during live events and product launches

*The following is taken from a commissioned study conducted by Forrester Consulting on behalf of IBM.”*

# Elastic Caching Integration

WebSphere eXtreme Scale and XC10 appliance can be easily integrated with other IBM application infrastructure products to provide a powerful, high-performance solution for your business needs

## Key value integration scenarios:

- WebSphere Portal Http Session Management to increase service capacity and provide high availability
- WebSphere Commerce & WebSphere Portal shared DynaCache replacement to reduce local memory requirements as well as providing cluster wide consistency
- Enterprise Service Bus side cache to offload redundant back-end processing and increase performance
  - ESB products supported include DataPower XI50, WebSphere Enterprise Service Bus, WebSphere Message Broker, and WebSphere Process Server



# Client Usage: Online Banking

## Retail Banking

**22 Million**  
online banking users

**35x**  
reduced  
response  
times

**\$500k**  
reduced  
costs per month

**20x**  
reduction in  
“FCIs”



## Next-generation Online Banking

- **Before:** 700ms to login with 2 backend calls
- **After:** 20ms to login with profile cache access
- \$6M/yr cost savings in MIPs reduction
- 3 datacenters: If one datacenter goes down it fails over seamlessly to the others
- 8Gb of data transfer per hour between DC's

Provide seamless cache infrastructure across applications

Deliver high performance & consistent response times

Ensure high availability of critical online applications

# Client Usage: Worldwide Fantasy Sports Web Site

Entertainment

7 Billion  
requests per day

10x  
reduced  
response  
times

## Fantasy Sports Web Infrastructure

- **Before:** 60ms response time against database
- **After:** WXS improved to 6ms response time
- 450k concurrent users
- 80k requests per second up to 1M in 2011
- 6 weeks from concept to production



Support transaction-intensive services

Deliver consistent & predictable response times

Take action on growing volumes of business events

Scale with simplicity and lower cost

# Client Usage: Investment Banking

## Investment Banking

**12 Million**  
orders per day

**4x**  
increase in  
revenue

**40x**  
number of  
transactions  
supported



## Next-generation Order Management System

- **Before:** Oracle RAC based architecture unable to scale to necessary demands
- **After:** 300K transactions / day → 45M / day
- Revenue up 4X and growing ... “all because of WebSphere eXtreme Scale”
- Moving to “22 x 7” operations (more than 9AM - 4PM)

Take action on growing volumes of business events

Deliver high performance & consistent response times

Ensure high availability of critical online applications

Scale with simplicity and lower TCO

# DataPower XI50/52 + XC10: Travel & Transportation

## Online Reservations

**100x**  
performance  
improvement

### Reservations System

- **Before:** 3-5 sec response time
- **After:** .01 -.05 sec response time
- Caching service requests
- Improved the average response time of the Global Distribution System requests for Fare Availability and Category Availability
- 52% caching rate
- 10 minute cache resulted in 40% reduction in load on the back-end systems
- Maintained high data integrity. Faster responses were also accurate
- POC in 3.5 hrs



Hot Deals,  
get our latest deals!

Sign up now →

Improved reliability and scalability of reservation channels

Reduced traffic to backend systems

Deliver high performance & consistent response times

Scale with simplicity and lower TCO

# Resources

- Fully functional J2SE trial download  
<http://www.ibm.com/developerworks/downloads/ws/wsdg/learn.html>
- Wiki documentation  
<http://www.ibm.com/developerworks/wikis/display/objectgrid/Getting+started>
- User's Guide to WebSphere eXtreme Scale  
<http://www.redbooks.ibm.com/abstracts/sg247683.html>
- Data Sheet  
<ftp://public.dhe.ibm.com/common/ssi/pm/sp/n/wsd14088usen/WSD14088USEN.PDF>
- XC10 Web Site  
<http://www-01.ibm.com/software/webservers/appserv/xc10/>
- Getting Started Wiki for XC10  
<http://www.ibm.com/developerworks/wikis/display/extremescale/2010/06/25/IBM+WebSphere+DataPower+XC10+appliance+now+available>

## Additional resources

Weekly video podcasts covering customers questions and forum posts on the IBM WebSphere eXtreme Scale product.



<http://www.youtube.com/user/ibmextremescale#p/a>

---



WebSphere Extreme Transaction Processing for Developers Space will discuss various topics for developing and deploying XTP applications and will point out emerging trends, benefits, challenges, and features associated with it.

<http://www.ibm.com/developerworks/spaces/xtp>



# 감사합니다





# Related Gartner Research

- ➔ **Taxonomy, Definitions and Vendor Landscape for Application Platform Products**  
Massimo Pezzini et al. (G00211411)
- ➔ **Innovation Insight: Invest in In-Memory Computing for Breakthrough Competitive Advantage**  
Massimo Pezzini (G00226070)
- ➔ **Hype Cycle for Application Infrastructure, 2011**  
Jess Thompson et al. (G00213370)
- ➔ **Predicts 2012: Cloud and In-memory Drive Innovation in Application Platforms**  
Massimo Pezzini et al. (G00226073)

For more information, stop by Gartner Solution Central or email us at [solutioncentral@gartner.com](mailto:solutioncentral@gartner.com).