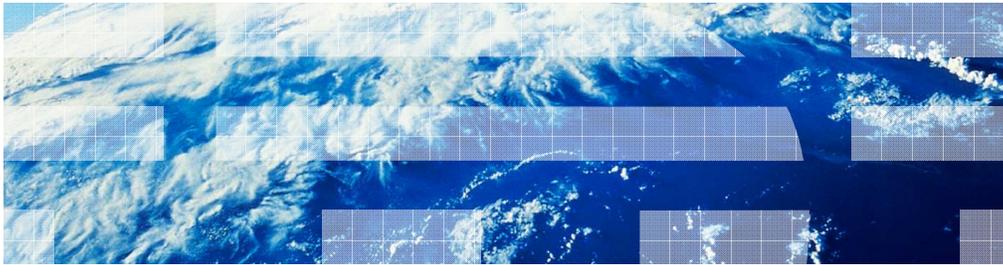


IBM WebSphere eXtreme Scale 7.1

Overview and new features



This presentation covers version 7.1 of WebSphere® eXtreme Scale, an elastic, scalable, in-memory data grid.

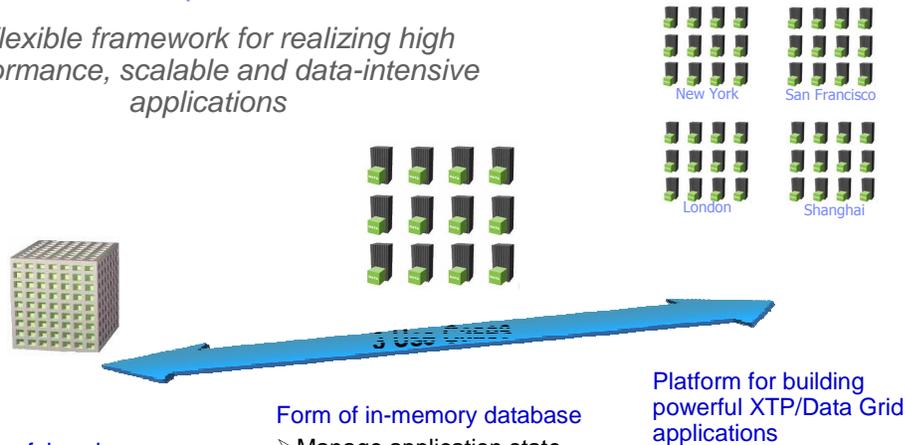
Overview

WebSphere eXtreme Scale dynamically caches, partitions, replicates, and manages application data and business logic across multiple servers. WebSphere eXtreme Scale performs massive volumes of transaction processing with high efficiency and linear scalability, and provides qualities of service such as transactional integrity, high availability, and predictable response times.

WebSphere eXtreme Scale can be used in different ways. It can be used as a very powerful cache or as a form of an in-memory database processing space to manage application state or as a platform for building powerful extreme transaction processing (XTP) applications.

What is IBM WebSphere eXtreme Scale?

A flexible framework for realizing high performance, scalable and data-intensive applications



Very powerful cache

- Scales from simple in-process topologies to powerful distributed topologies

Form of in-memory database

- Manage application state
- Scales to 1000's of servers
- Sometimes referred to as Distributed Application State Management

Platform for building powerful XTP/Data Grid applications

What is IBM WebSphere eXtreme Scale? IBM WebSphere eXtreme Scale is a flexible framework for realizing high performance, scalable and data-intensive applications. It's a single map addressable space of virtualized JVM heap spaces.

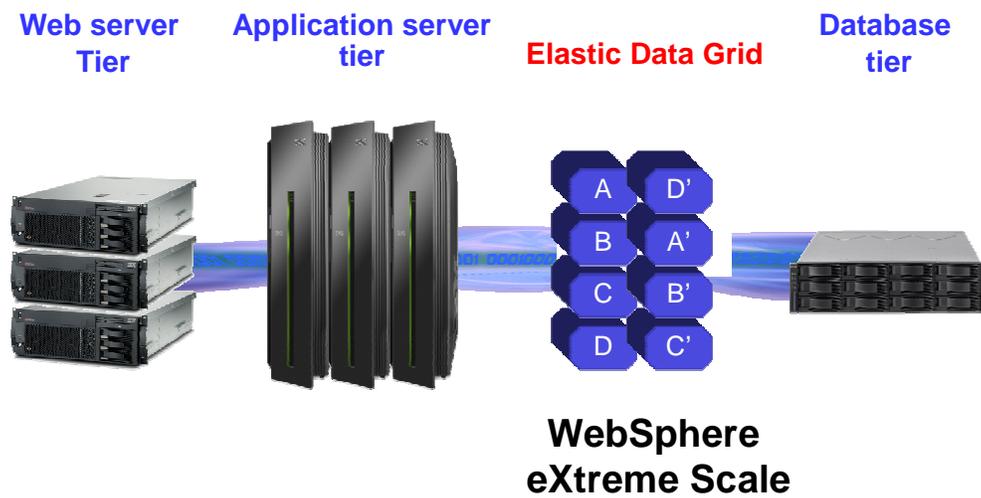
First use of WebSphere eXtreme Scale is a power cache for your applications.

WebSphere eXtreme Scale can move that collection of application caches into a grid that will be highly available, elastic, and self-healing.

Second use of WebSphere eXtreme Scale is a form of in-memory database.

Third use of WebSphere eXtreme Scale is a form of redundancy across datacenters.

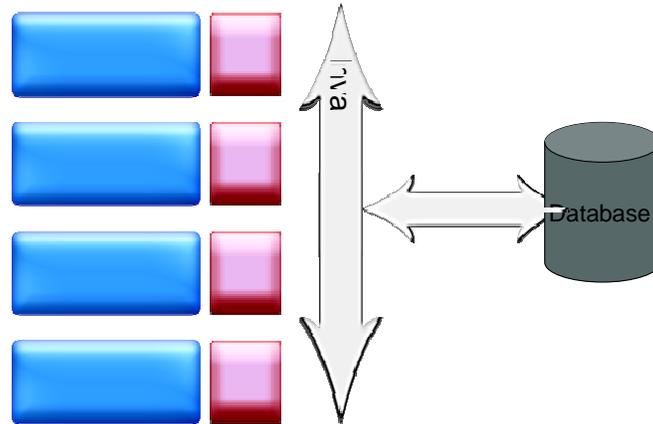
Scale with simplicity



Here is a very simple depiction of WebSphere eXtreme Scale core competency. Positioning WebSphere eXtreme Scale between the database tier and application tier. Database tier is usually more expensive to scale. An elastic data grid like WebSphere eXtreme Scale has predictable scaling at predictable cost, preserving your database tier and increasing its life span.

Traditional cache operation (1 of 2)

- Cache capacity determined by individual JVM Size.
 - Size of each cache = M
 - # JVMs = N
 - **Total cache = M**
- Invalidation load per server increases as cluster grows.
- Cold start servers hit the database.

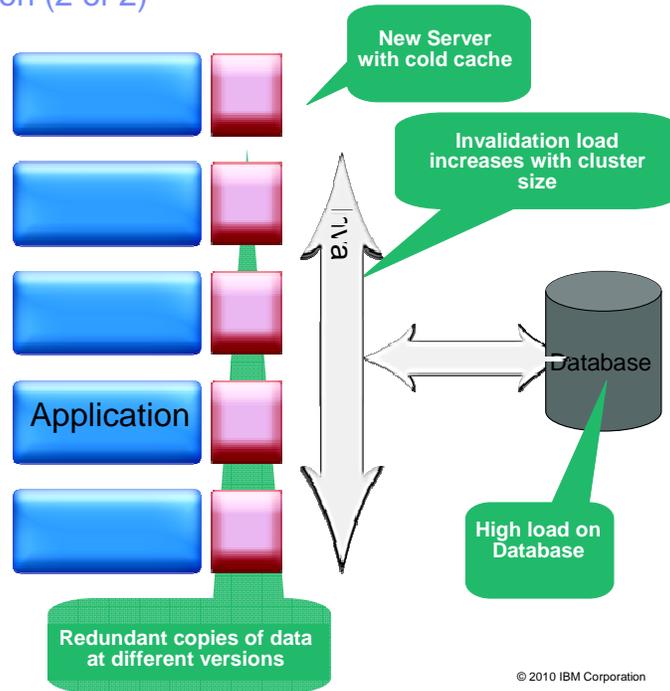


What are the differences in traditional cache versus in-memory grid?

A cache is only as big as what can be held in a single JVM. In a traditional cache in this example the four application instances with separate caches each has to go to the database four different times.

Traditional cache operation (2 of 2)

- Cache capacity determined by individual JVM Size.
 - Size of each cache = M
 - # JVMs = N
 - Total cache = M**
- Invalidation load per server increases as cluster grows.
- Cold start servers hit the database.



6

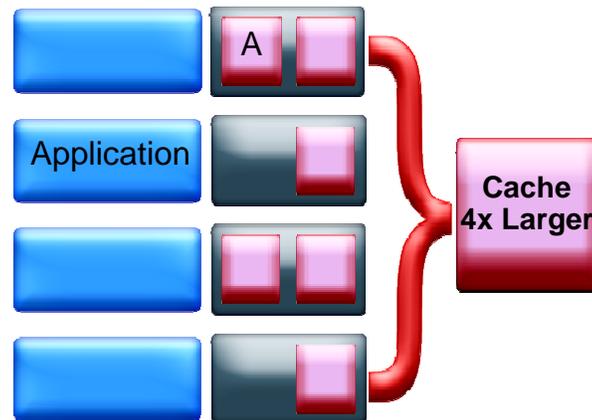
Overview and new features

© 2010 IBM Corporation

As you add more application servers the invalidation chatter actually grows. Its an inversely scalable problem, and increase the load on the database tier.

WebSphere eXtreme Scale (1 of 3)

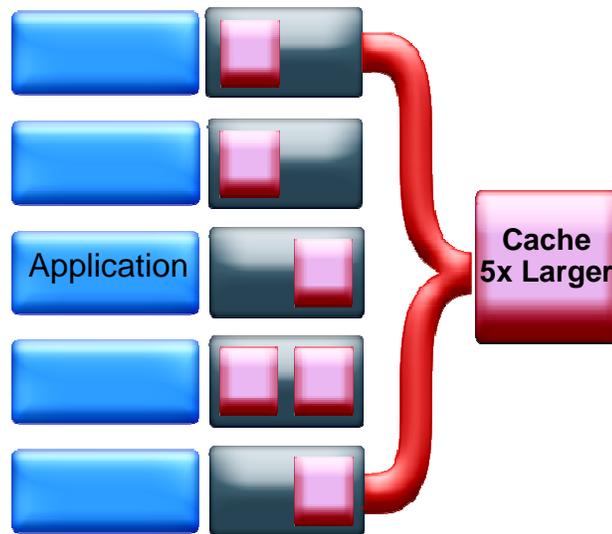
- Cache capacity determined by total cluster size
 - Size of each cache = M
 - # JVMs = N
 - Total Cache = M x N
- No invalidation chatter
- Linearly scalable
- Less load on database and no cold start spikes



In WebSphere eXtreme Scale there is only one copy of “A” in the grid. For example, if the application server on the bottom needs “A”, it will fetch it from the top application server. In this scenario, the scale is linear and provides less load on the database.

WebSphere eXtreme Scale (2 of 3)

- Cache capacity determined by total cluster size
 - Size of each cache = M
 - # JVMs = N
 - Total Cache = $M \times N$
- No invalidation chatter
- Linearly scalable
- Less load on database and no cold start spikes



Versus traditional cache, if a new application server is added, it benefits immediately from the cache shared by the other application servers. The total cache size grows further reducing the load on the database.

WebSphere eXtreme Scale (3 of 3)

- Proven mature product:
 - Fourth major release of product with V7.1
 - Public references
 - Private references
 - Used at some of the largest websites in the world
- Lightweight runtime footprint (15MB jar)
- Integrates with all versions of WebSphere Application Server and almost any Java™-based application container or Java Virtual Machine (1.4.2 and higher)
- Proven multi-data center capabilities
- Proven low-latency access to data



WebSphere eXtreme Scale 7.1 is a proven and mature product with a lightweight footprint. WebSphere eXtreme Scale 7.1 integrates with all versions of WebSphere Application Server and almost any Java-based application container or Java Virtual Machine.

Installation options

- Standalone installation
 - Requires J2SE, no other requirements
 - Main JARs: ogclient.jar, objectgrid.jar
 - Client JVMs should use the ogclient.jar
 - Server JVMs should use the objectgrid.jar
 - A **-Djava.endorsed.dirs** JVM parameter must be specified to enable the IBM ORB (ORB is supplied with eXtreme Scale from V6.1.0.5 onwards)
- Integrate eXtreme Scale with WebSphere Application Server
 - Integrates with WebSphere Application Server V6.0, V6.1 and V7.0
 - Main JARs: wsogclient.jar, wsobjectgrid.jar
 - Java EE applications can be clients, containers or both.
 - As an alternative to installing eXtreme Scale on top of WebSphere Application Server Network Deployment, an application can simply include the wsogclient.jar in the module enabling a client-only operation thus avoiding installing additional software on top of Network Deployment.

WebSphere eXtreme Scale has two installation options.

A stand-alone installation that requires only Java 2 Platform, Standard Edition .

You can also Integrate WebSphere eXtreme Scale with WebSphere Application Server.

You can install on top of WebSphere Application Server which augments the profile. Now any application defined in your cell can become a container server.

Deployments comparison

Feature	Stand-alone (J2SE)	WebSphere Application Server
Installation	Installation wizard pointed to empty directory	Installation wizard pointed to WebSphere Application Server directory, need to augment profiles
Java runtime	Use latest and greatest IBM or Sun Java runtimes and ORB.	Limited to WebSphere Application Server supported Java runtimes and ORB.
Administration	Server start scripts, xsadmin script, MBeans	WebSphere Application Server administrative console, wsadmin and xsadmin scripts, MBeans
Scalability	Simply add more JVMs	Limited by size of cell (A multi-cell configuration is possible.)
High availability	Fully configurable	Fully configurable
Memory footprint	Only eXtreme Scale and application data in memory	Additional memory for the WebSphere Application Server Runtime
Configuration	XML, Spring and Annotations	XML, Spring and Annotations
WebSphere Application Server integration	N/A	Grid servers automatically integrated with application life cycle. Integrated with logging and PMI, optionally WebSphere Application Server transaction demarcation aware
Licensing	Only eXtreme Scale license	WebSphere Application Server Network Deployment + eXtreme Scale license
Programming model	All APIs are available	All APIs are available

Here is a list comparing the various differences of WebSphere eXtreme Scale Deployment attributes.

New features

Next you will see some of the new features in WebSphere eXtreme Scale 7.1.

Spotlight on WebSphere eXtreme Scale V7.1

- Broadened market applicability
 - More robust support for cross data-center replication usage scenarios, using multi-master replication to propagate the data modifications made by each member to the rest of the group.
 - Windows®/.NET support through ADO.Net/ WCF Data Services V1.0 specification support
- Improved reporting and monitoring
 - New monitoring console which can provide real time and historical views into the WebSphere eXtreme Scale data grid
 - New WebSphere Application Server Console extensions that can provide improved integration with WebSphere eXtreme Scale
 - Administrator options to include grid queries in tracking the source of SQL requests to IBM DB2® through support for DB2 Performance Monitor
- Enhanced foundational capabilities
 - Infrastructure improvements to accelerate performance and optimize resource usage

In WebSphere eXtreme Scale 7.1, three main usage scenarios was targeted.

Broadened market applicability with .NET support through ADO.Net and more robust support for cross data-center replication usage scenarios, using multi-master replication to propagate the data modifications made by each member to the rest of the group.

New monitoring console which can provide real time and historical views into the WebSphere eXtreme Scale data grid.

Provide infrastructure improvements to accelerate performance and optimize resource usage

CAP theorem

- In a distributed environment, especially spanning data centers across LANs and WANs there are three core requirements for a service:
 - Consistency
 - Either the service works or it fails
 - Traditional ACID of databases provides consistency and isolation
 - Availability
 - Extremely important in web business model
 - In a large distributed system, one may have to compromise with consistency for the sake of availability
 - Partition tolerance
 - Network partition will happen when not all machines are connected
 - “No set of failures less than the total network failure is allowed to cause the system to respond incorrectly” – Seth and Lynch
 - Quorum is used to guard against split brain syndrome
- Brewer’s CAP conjecture states that
 - One can achieve only two not all three of the above mentioned requirements

There are three aspect of a distributed environment.

Consistency: Either the service works or it fails

Availability: Extremely important in web business model

Partition tolerance: Network partition will happen when not all machines are connected.

Partition in CAP refers to network partition, like some communication failure between data centers or even machines. Do not confuse it with data partitioning.

Brewer’s conjecture on CAP has been formally provided by Seth and Lynch of MIT

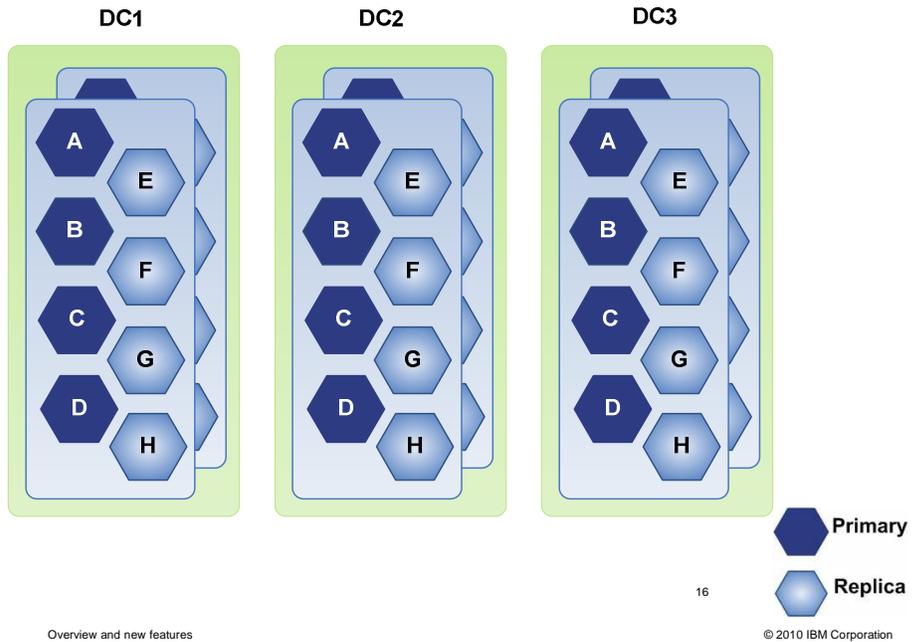
WebSphere eXtreme Scale and CAP theorem

- **CP** Only eXtreme Scale provides this option
 - Consistency and partition tolerance
 - eXtreme Scale never allows split brain syndrome
 - Across data centers eXtreme Scale provides a single grid
- **AP** (new in V7.1!) (Some competitors provide **AP**)
 - Availability and partition tolerance
 - Grid will never reject an update, but in the presence of a partition, the real value may not be known
 - Set precedence OR call out to user code
- **CA** does not scale (A single node cache)
 - Consistency and availability
 - Can be achieved by preventing network partitions to happen
 - Host everything in single machine
 - Not suitable in XTP environment

Only WebSphere eXtreme Scale provides this option of consistency and partition tolerance.

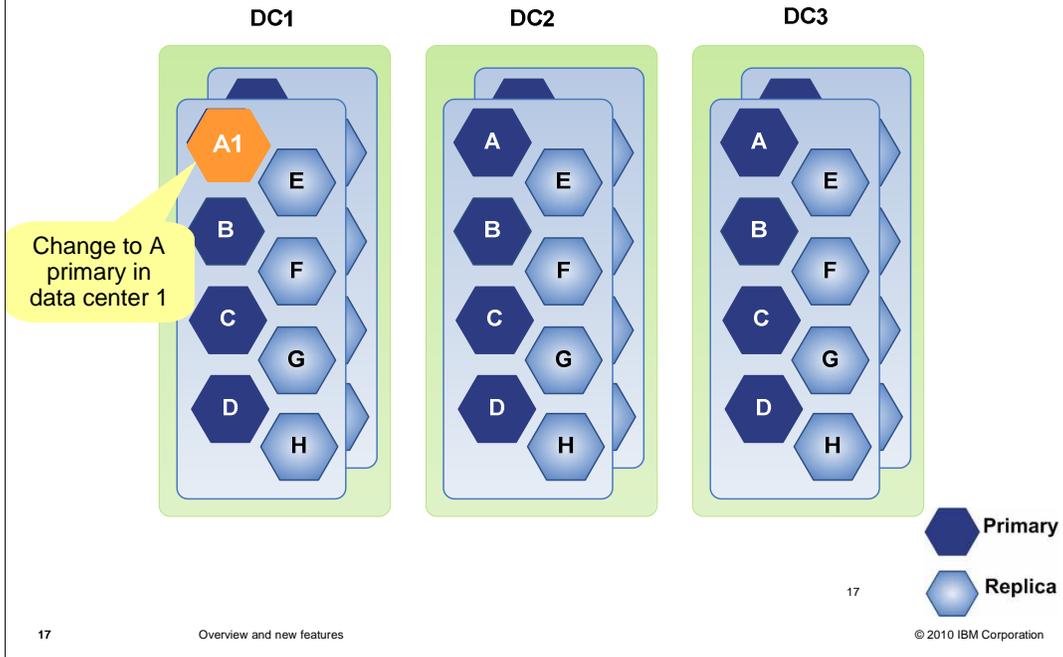
New in WebSphere eXtreme Scale 7.1 is availability and partition tolerance. The grid will never reject an update, but in the presence of a partition, the real value may not be known.

Multi-master replication (1 of 5)



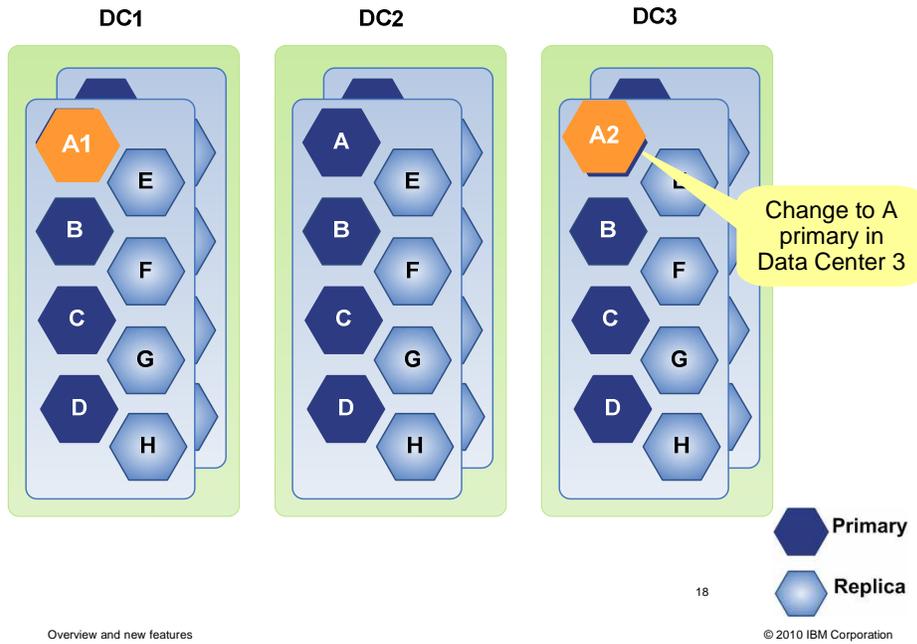
Here is an example of Multi-Master Replication with WebSphere eXtreme Scale 7.1. It has three data centers with a non-partitioned grid.

Multi-master replication (2 of 5)

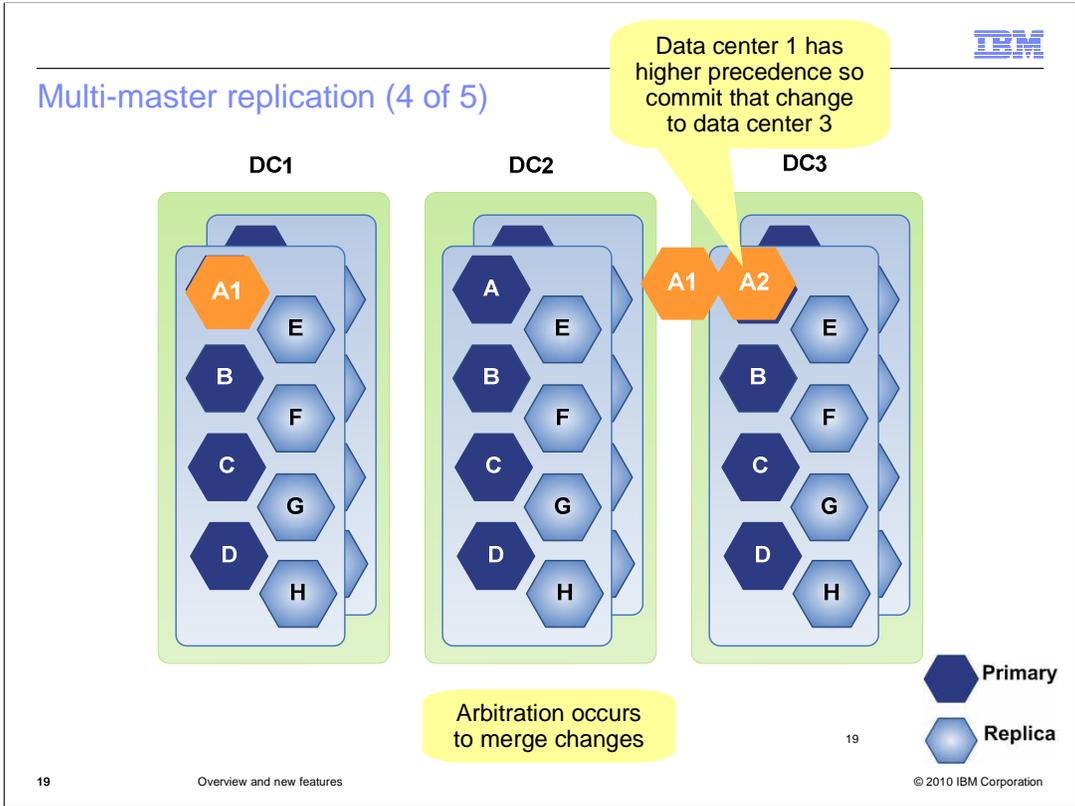


Change to A primary in data center 1.

Multi-master replication (3 of 5)

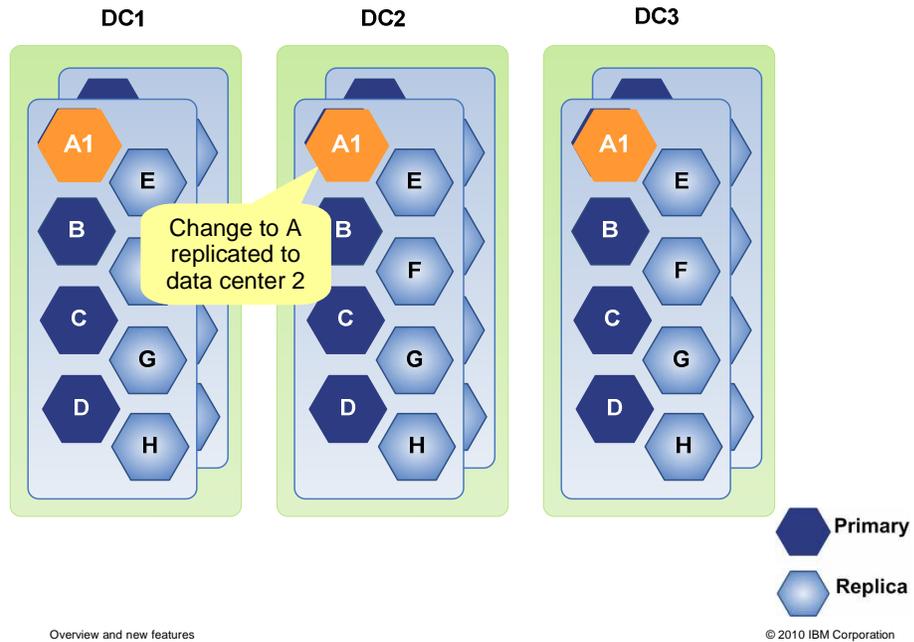


Change to A primary in data center 3. Now there are two different versions.



There is currently a conflict of values for 'A' availability and partition tolerance (AP) policy states data center 1 has higher precedence so commit that change to data center 3.

Multi-master replication (5 of 5)



First there is a change to A in data center 3 to A1, then change to A replicated to data center 2. Now the new value of A is consistent through all data centers.

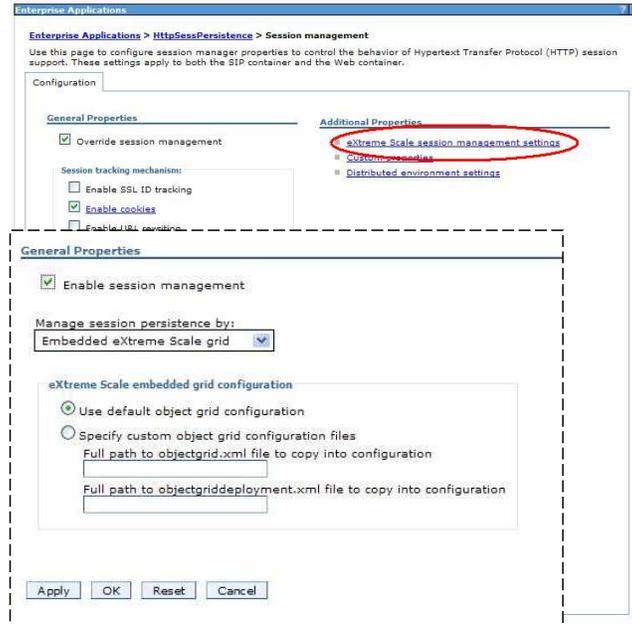
Combination CP and AP support

- For some applications, it may be desirable to have a mix of CP and AP support for different types of requests
- Example: Store front application
 - **Shopping Cart -> AP configuration**
 - Shopping cart must always be available though there may be consistency issues if the data centers are partitioned and the data is updated separately
 - **Purchase Transaction -> CP configuration**
 - Once the user decides to purchase what is in the shopping cart then the application will use a CP configuration to ensure that the transaction is committed exactly once

Since there are instances where both consistency and partition configurations and availability and partition configurations are needed. WebSphere eXtreme Scale 7.1 provides both to get the optimal replication strategy across datacenters.

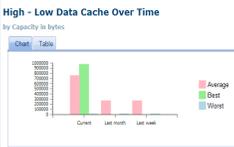
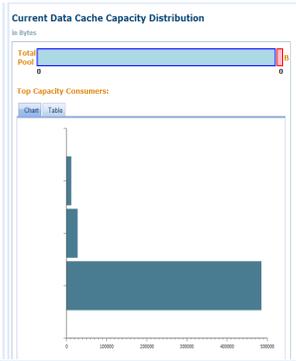
WebSphere Application Server console extensions

- Administrators and configurators can create or register catalog services
- Administrators and deployers can enable eXtreme Scale session persistence



When you augment the profile of WebSphere Application Server. You enable the ability to distribute HTTP Sessions in the grid. New options appear in the Session Management section of the WebSphere Application Server administration console.

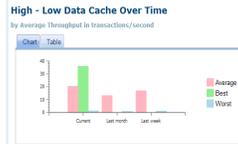
New monitoring console - All data caches overview



Top 5 Data Caches

By Average Transaction Time in milliseconds

Data Caches	Current	Data over last 59 min	[Min Max Avg]
TXS	.11		[.001 .62 .07]
TestGrid	.04		[.001 .12 .07]
SessionGrid	.06		[.002 .45 .08]



Data Caches

Current Data over last 59 min [Min | Max | Avg]

Data Caches	Current	Data over last 59 min	[Min Max Avg]
TXS	1.20		[.13 1.53 .89]
TestGrid	1.73		[.13 2.40 1.27]
SessionGrid	17.33		[1.3 22.73 12.66]



- The most broad of the overviews
- Various charts on capacity, throughput, and transaction time

New in WebSphere eXtreme Scale 7.1, is a monitoring console.

References

- WebSphere eXtreme Scale V7.0 Online Information Center
 - <http://publib.boulder.ibm.com/infocenter/wxinfo/v7r0/index.jsp>
- WebSphere eXtreme Scale Wiki
 - <https://www.ibm.com/developerworks/wikis/display/extremescale/Home>
- WebSphere eXtreme Scale Trial Download
 - <http://bit.ly/eXtremeScaleTrial>
- Run WebSphere eXtreme Scale in the clouds! -- on Amazon EC2
 - <http://bit.ly/extremescaleAM>
- WebSphere Extended Deployment Forum
 - <http://www.ibm.com/developerworks/forums/forum.jspa?forumID=778>
- IBM Redbook: User's Guide to WebSphere eXtreme Scale
 - <http://bit.ly/eXtremeScaleRedBook>
- Article: Understanding WebSphere eXtreme Scale and how it works
 - <http://bit.ly/eXtremeScaleIntro>
- Article: Leveraging WebSphere eXtreme Scale as an in-line database buffer
 - <http://bit.ly/Vuong0906>
- WebSphere eXtreme Scale Videos
 - <http://www.youtube.com/ibmextremescale>
- Read Billy Newport's blog on XTP
 - <http://www.devwebsphere.com/>

For further understanding of WebSphere eXtreme Scale 7.1 and its benefits, visit these links for more information.



Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WXS71_CEP.ppt

This module is also available in PDF format at: [../WXS71_CEP.pdf](..WXS71_CEP.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DB2, and WebSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. J2SE, Java, JVM, Sun Java, and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2010. All rights reserved.