



IBM Software Group

IBM WebSphere eXtreme Scale V7.0

Architecture



@business on demand.

© 2009 IBM Corporation
Updated July 6, 2009

This presentation will cover the WebSphere® eXtreme Scale V7.0 architecture.

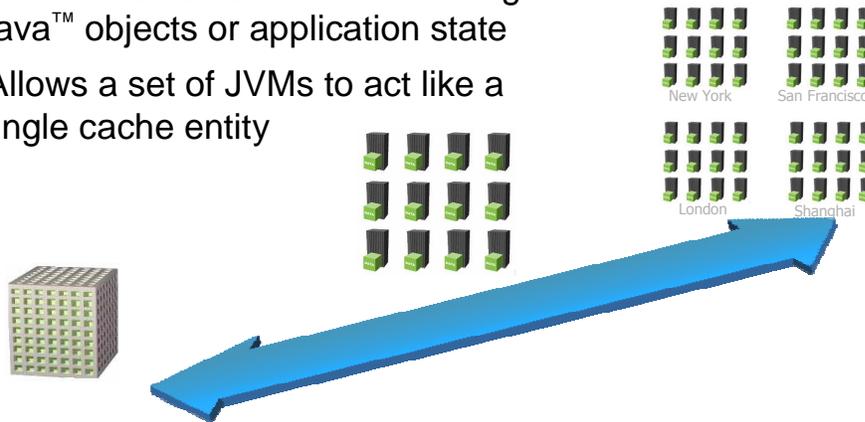
Agenda

- eXtreme Scale grid
- Grid containers
- BackingMaps
- Partitions
- Shards
- eXtreme Scale clients and servers
- Dynamic clustering
 - ▶ Catalog server
 - Quorum
 - ▶ Zones

The agenda is to cover the eXtreme Scale terminology in order to understand how an eXtreme Scale grid is structured and how applications use the grid.

eXtreme Scale grid

- An eXtreme Scale grid is a logical grouping of data
 - ▶ Enables a JVM to host one or more grid containers that can be used as a cache for storing Java™ objects or application state
 - ▶ Allows a set of JVMs to act like a single cache entity



WebSphere eXtreme Scale enables a Java Virtual Machine to host one or more grid containers that can be used as a cache for storing Java objects or application state. An eXtreme Scale grid is a self managed group of connected grid containers that act like a single cache entity.

Because grid containers can be added to the eXtreme Scale grid as more and more data needs to be stored, WebSphere eXtreme Scale provides a very scalable caching solution.

Grid containers

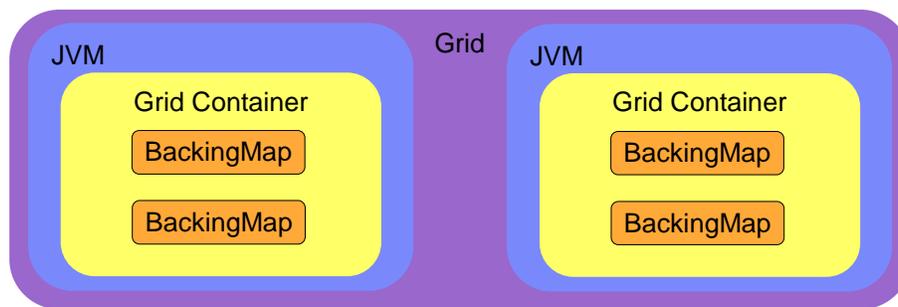
- A grid container holds data in a collection of maps
 - ▶ In the eXtreme Scale programming model, these are called BackingMaps
 - ▶ Provides services such as security, transaction support, JNDI lookup, and remote connectivity



Much like J2EE containers, like Web or EJB containers, a grid container provides the grid application services such as security, transaction support, JNDI lookup service, remote connectivity and so forth. It houses data in a collection of maps called BackingMaps.

BackingMaps

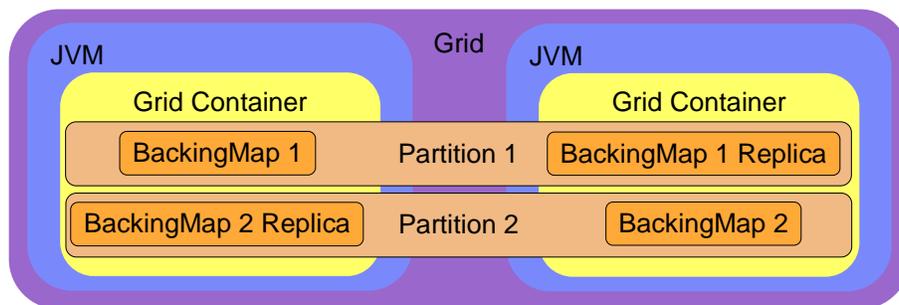
- A map is an interface that stores data as key/value pairs
- A BackingMap contains cached objects that have been stored in the grid



A BackingMap is a type of map that contains cached objects that have been stored in the grid. Each BackingMap has its own loader instance. It will request any needed data that it does not contain from the loader, which in turn, will retrieve it from the back end data store.

Partitions

- Process of splitting data into smaller sections
- Data related by schema are all in the same partition
- Each partition has a primary shard and an optional set of replica shards



6

Partitioning is the process of splitting data into smaller sections. It allows the grid to store more data than can be accommodated in a single Java Virtual Memory heap. The data partitioning is accomplished by using an application-defined schema. Each partition is made up of a primary shard and one or more replica shards that can be replicated to multiple servers. (Note that some applications may allow only one partition.)

Shards

- A shard is an in-memory database
 - ▶ Represents an instance of a partition that is placed on a container
 - ▶ 1 partition=>multiple shards (1 primary + replicas)
- Shards can be replicated from one JVM to another
 - ▶ For high availability, a primary shard and replica shards exist for each partition in separate grid containers
 - ▶ Multiple shards that represent different partitions can exist on a single grid container
 - ▶ The replica shards are either synchronous (sync) or asynchronous (async)
- Shard placement is determined by catalog servers
 - ▶ The types and placement of replica shards are directed by using a deployment policy, which specifies the minimum and maximum number of synchronous and asynchronous shards

7

The term shard is used to define a single instance of a partition. Each partition is made up of a primary shard and zero or more replica shards. The shard distribution algorithm ensures that the primary and replica shards are never in the same container to ensure fault tolerance and high availability. Shards that represent different partitions can exist in a single grid container however.

Shard placement is the responsibility of the catalog servers. The type and placement of replica shards are directed by a deployment policy. So, as the grid members change and JVMs are added to or removed from the grid, the catalog server will redistribute shards accordingly to best fit the available JVMs. The ability of the grid to scale out or scale in provides tremendous flexibility to the changing nature of the eXtreme Scale infrastructure.

Shard types

■ Primary

- ▶ Receives all insert, update and remove operations.
 - Replicas are read-only
- ▶ Adds and removes replicas
- ▶ Replicates data to the replicas
- ▶ Manages commits and rollbacks of transactions
- ▶ Replicas turn into primaries to handle failure events or rebalancing

■ Synchronous replica

- ▶ Maintains the same state as the primary
- ▶ Receives updates as part of the primary's transaction to guarantee consistency

■ Asynchronous replica

- ▶ Might or might not be at the same state as the primary
- ▶ Receives updates after the transaction commits on the primary
 - Primary does not wait for the asynchronous replica to commit

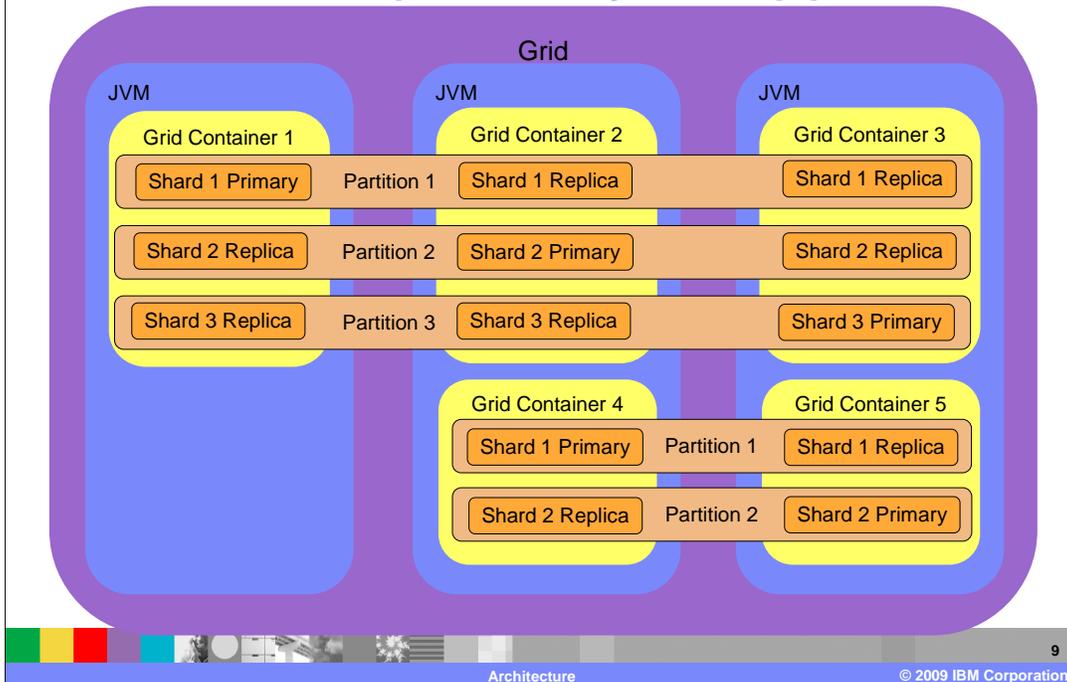


Primary and replica shards are balanced across containers. Replicas turn into primaries to handle failure events. This slide discusses some of the details of each shard type.

The synchronous replica shard data remains at the same level as the primary shard data. Whenever a transaction commits on the primary, the synchronous replica shard will update its contents. The primary shard waits until the replica shard responds that it got the data.

The asynchronous replica data shard is not guaranteed to remain at the same level as the primary shard data. When the primary shard sends data to the asynchronous replica, it does not wait for a response.

When shards are replicated they make up partitions



When shards are replicated they make up partitions. This figure shows three partitions on the top with two replicas each and on the bottom two partitions each with a single replica, making 13 shards total. All of the shards are located in grid containers. Grid containers 1, 2, and 3 each hold three shards, while grid containers 4 and 5 hold two shards each.

ObjectGrid clients and servers

- An ObjectGrid server
 - ▶ Contains grid containers which host the cache (stored in BackingMaps)
 - ▶ Is a JVM running on a particular host and port that
 - Listens on a port for client requests
 - Can communicate with other ObjectGrid servers to create a self-configuring grid
- An ObjectGrid client
 - ▶ Can connect to an ObjectGrid server to perform various operations, such as create, retrieve, update and delete the objects in the grid
 - ▶ Contains an ObjectMap and may contain a “near-cache” copy of a BackingMap



There are two components that can contain BackingMaps within the eXtreme Scale framework, the ObjectGrid server and the ObjectGrid client. The ObjectGrid server, or simply grid server, contains grid containers which host the cache, stored in BackingMaps. The grid server is a JVM running on a particular host and port that listens for client requests and can communicate with other ObjectGrid servers and catalog servers to maintain a self-configuring, self-healing eXtreme Scale grid.

The ObjectGrid client, or grid client, will connect to a grid server to perform operations such as create, retrieve, update and delete to the objects in the eXtreme Scale grid within the scope of a single transaction. The client contains an ObjectMap and may contain a copy of a BackingMap.

The grid server BackingMap is always shared between clients, while the client-side BackingMap, or near-cache, is shared only between threads of the grid client.

The terms “grid server” and “ObjectGrid server” can be used interchangeably.

ObjectGrid servers and catalog servers

- Two kinds of servers used in a dynamic eXtreme Scale environment
 - ▶ The ObjectGrid server which hosts shards in its grid containers
 - ▶ The catalog server which controls how
 - Partition shards are distributed
 - ObjectGrid servers are organized
 - ObjectGrid servers health is monitored
- When an ObjectGrid server starts up it registers with a catalog server
- When an ObjectGrid client wants to access the eXtreme Scale grid, it bootstraps to a catalog server to get ObjectGrid information
 - ▶ Catalog servers do not participate in the normal flow of requests from ObjectGrid clients to servers



The ObjectGrid servers that hold data in their grid containers are not the only type of eXtreme Scale servers. Catalog servers also play a role in the management of the eXtreme Scale grid. When an ObjectGrid server starts up it registers with a catalog server. The catalog server manages how the partition shards are distributed, monitors the health and status of the ObjectGrid servers, and enables the ObjectGrid clients to locate the primary shard for a requested object.

Catalog server responsibilities

- Registration of grid containers
- Core group assignment of containers
- Partition placement
- Replication group dynamic configuration
- Shard provisioning
- Partition balancing
- Interaction with WebSphere Virtual Enterprise placement
- Bootstrap for ObjectGrid client



The catalog server is the engine that drives the grid operations. It provides services essential to maintaining a scalable self-healing grid.

Clustered catalog servers

- Catalog servers can be clustered
- One of clustered catalog servers is elected as master catalog server
- When the master catalog server fails, one of standby catalog servers in the catalog server cluster is elected as the new master catalog server
- Any ObjectGrid clients and already started ObjectGrid servers will be automatically rerouted to the new master catalog server



Catalog servers can be clustered to ensure high availability. Initially, one of the catalog servers is elected as the master catalog server. If the master fails, a standby catalog server is elected as the new master. ObjectGrid clients and servers are automatically rerouted to new master catalog server.

-quorum true

- Writes to catalog service state are committed only when the majority of the catalog servers participate in the transaction
- Containers that are changing states cannot receive any commands until the catalog service transaction commits
- Catalog service hosted in a minority partition (no quorum)
 - ▶ Accepts heartbeat messages from core group leaders
 - ▶ Cannot accept server registrations or membership changes
 - ▶ Catalog state is essentially frozen until:
 - Quorum is reestablished
 - Quorum override is set on any catalog server that is online



By default, quorum mode is enabled. It allows writes to the catalog service state to be committed only when the majority of the catalog servers participate in the transaction. Containers that are changing states cannot receive any commands, unless the catalog service transaction commits first.

If the catalog service is hosted in a minority partition, meaning no quorum has been established, it will accept liveness messages. It cannot, however, accept server registrations or membership changes, because the state is essentially frozen until one of following two things occurs: quorum is reestablished or quorum override is set on any catalog server that is on-line.

This approach allows the catalog service to manage shard distribution and to maintain a consistent grid server membership list.

Replication using zones

- Zones are exclusive collections of JVMs
- Each JVM can be tagged with a zone name and will belong to exactly one zone
- Zone rules allow placement rules for shards within zones
- Rules can specify that different shard types can not be placed on machines within the same zone.
 - ▶ A replica of a partition can not be placed in the same zone as its primary for example
- Zones enable advanced placement of shards across racks, floors, buildings or cities



Zone capability is useful to ensure that primary and replica shards are placed in different physical locations for better fault tolerance. A zone is essentially an exclusive collection of JVMs. Rules are used to specify the possible set of zones that a shard can be placed in. The catalog server will follow the zone rules

to place the primary and replica shards across these groups of JVMs, or zones. As a general practice, you should place only synchronous replicas in the same zone and asynchronous replicas in a different zone to ensure high availability for a JVM failure in a local zone (synchronous replica) and in the case of a complete data center failure (asynchronous replica). Zones can exist across racks, floors, buildings, or even cities.

Using zones with HTTP Session Manager

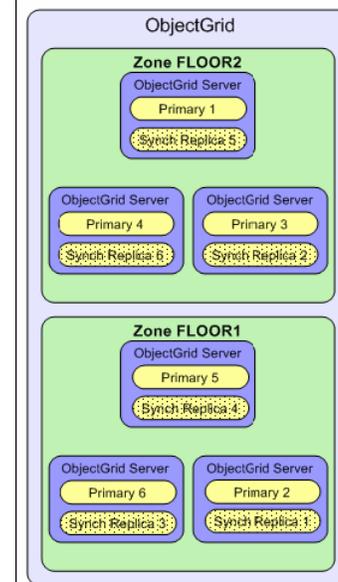
- The HTTP Session Manager for eXtreme Scale can also use zones.
 - ▶ Deploy a single Web application across multiple data centers
 - ▶ HTTP sessions are replicated across data centers
 - ▶ Sessions can be recovered even if an entire data center fails



Zone support is also provided for HTTP Session Manager, so HTTP sessions can be dispersed across different locations or zones for better fault tolerance.

Zone example

- A customer might name zones after floors in a building
- Use zones to ensure that primaries and replicas for the same data are on different floors
- A complete set of data will always exist on both floors



Here is an example of how to use zones. A customer using zones across two physical locations, such as floors in a building for example, would have the primaries distributed across both zones. The replica for each primary would be placed on a different floor than the primary. In this case, the floor name is used as the zone name. Zones ensure the data is adequately replicated between building floors and the data is available to the application regardless of its zoned location.

Summary

- eXtreme Scale grid enables a JVM to host grid containers
- BackingMap contains cached objects that have been stored in the grid
- Each partition has a primary shard and an optional set of replica shards
- Catalog server maintains the operation of the grid servers and shard distribution
- Zones enable shards to be replicated across physical locations
- eXtreme Scale clients and servers can be used in a distributed cache topology

This presentation covered the eXtreme Scale terminology to show how an eXtreme Scale grid is structured. WebSphere eXtreme Scale enables a JVM to host one or more grid containers that can be used as a cache for storing Java objects or application state. Cached objects are stored in a BackingMap. Partitions are used to split data into smaller groups and consist of a primary shard and zero or more replica shards. A catalog server plays an important role in the environment as the manager of shard distribution and the monitor of the availability and health of the grid servers. Zones can be used to ensure primary and replica shards are placed in separate physical locations. Finally, in a distributed cache topology, eXtreme Scale clients can be used to separate the application logic from the grid servers to reduce the load on your back-end data store.

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 e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WXS70_Architecture_Overview.ppt

This module is also available in PDF format at: [../WXS70_Architecture_Overview.pdf](..WXS70_Architecture_Overview.pdf)



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

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, [ibm.com](http://www.ibm.com), and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both: WebSphere

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

EJB, J2EE, Java, JVM, and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

