GlassFish High Availability Overview

Shreedhar Ganapathy
Engg Manager, GlassFish HA Team
Co-Author Project Shoal Clustering

Email: shreedhar_ganapathy@dev.java.net
http://blogs.sun.com/shreedhar

What we will cover today

- What is Project GlassFish ?
- Explain GlassFish v2 Clustering
  - Cluster Creation & Configuration
  - Cluster Lifecycle components
  - Runtime Dynamic Clustering - Project Shoal
- GlassFish High Availability options
  - HADB based
  - In-Memory Replication based
- In-memory replication – how it works
- Configuration
- Q&A

What Is Project GlassFish ?

- Open Source Project developing Java™, Enterprise Edition (Java EE platform) Application Server
  - https://glassfish.dev.java.net
- Free for development, deployment, and redistribution
- Fully Java EE 5 compliant
- Open Source
  - Dual OSI Licenses – CDDL, GPL v2 w/CPE
- Community at java.net
  - Source Code, Bug Database, Discussions/Forums at java.net
  - Roadmaps, Architectural documents
  - Code contributions: Sun Microsystems,Oracle & others

GlassFish Clustering/HA Introduction

- **What is High Availability?**
  - Ensuring a certain absolute degree of operational continuity of business services even when part of the system has failed.
  - Continued availability of business services to users during the measurement period

- **How do you get High Availability ?**
  - Cluster server instances for redundant availability of service
  - Use Session State Persistence for redundant availability of business data/sessions
  - Size your middleware infrastructure for Fault Tolerance, Scalability, and Reliability
  - Supported by a Load Balancing tier to distribute load

GlassFish v2 Clustering Lifecycle components

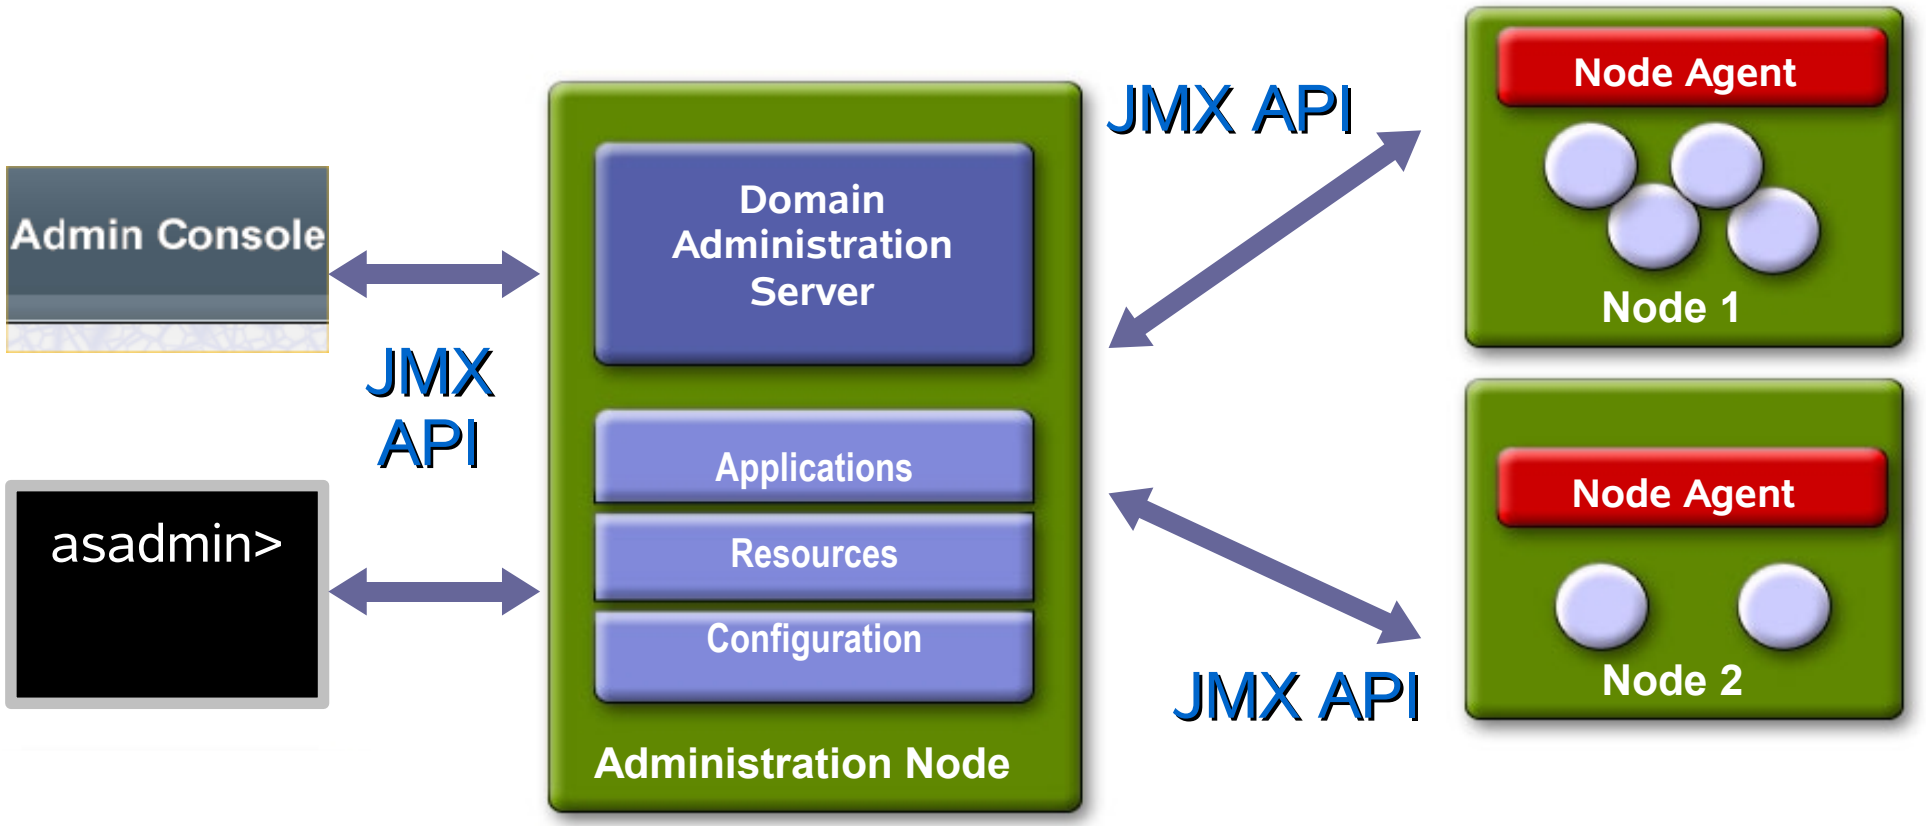

JMX = Java Management Extensions

GlassFish v2 Web Admin Console

- Creating a Cluster

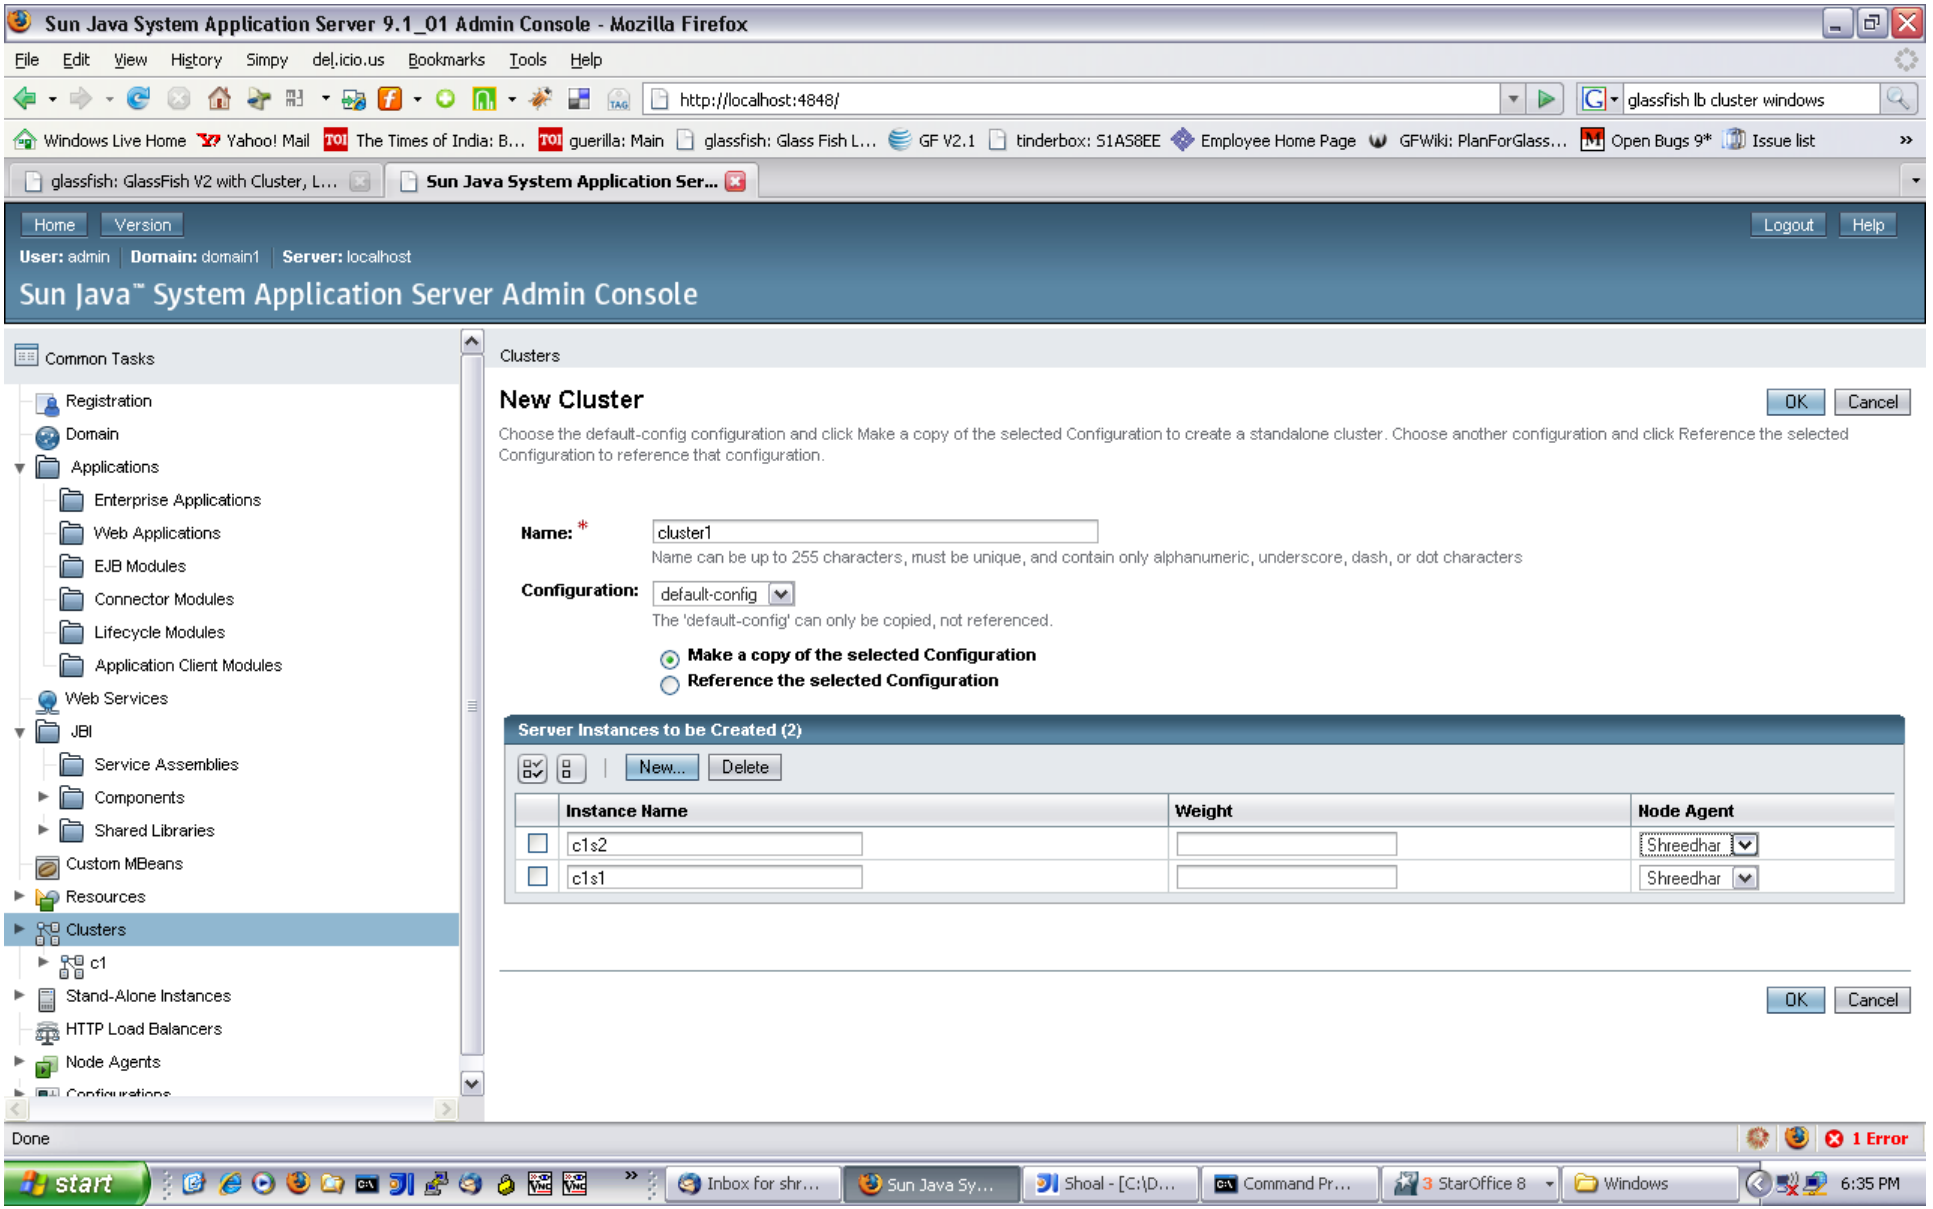

Clustering Architecture

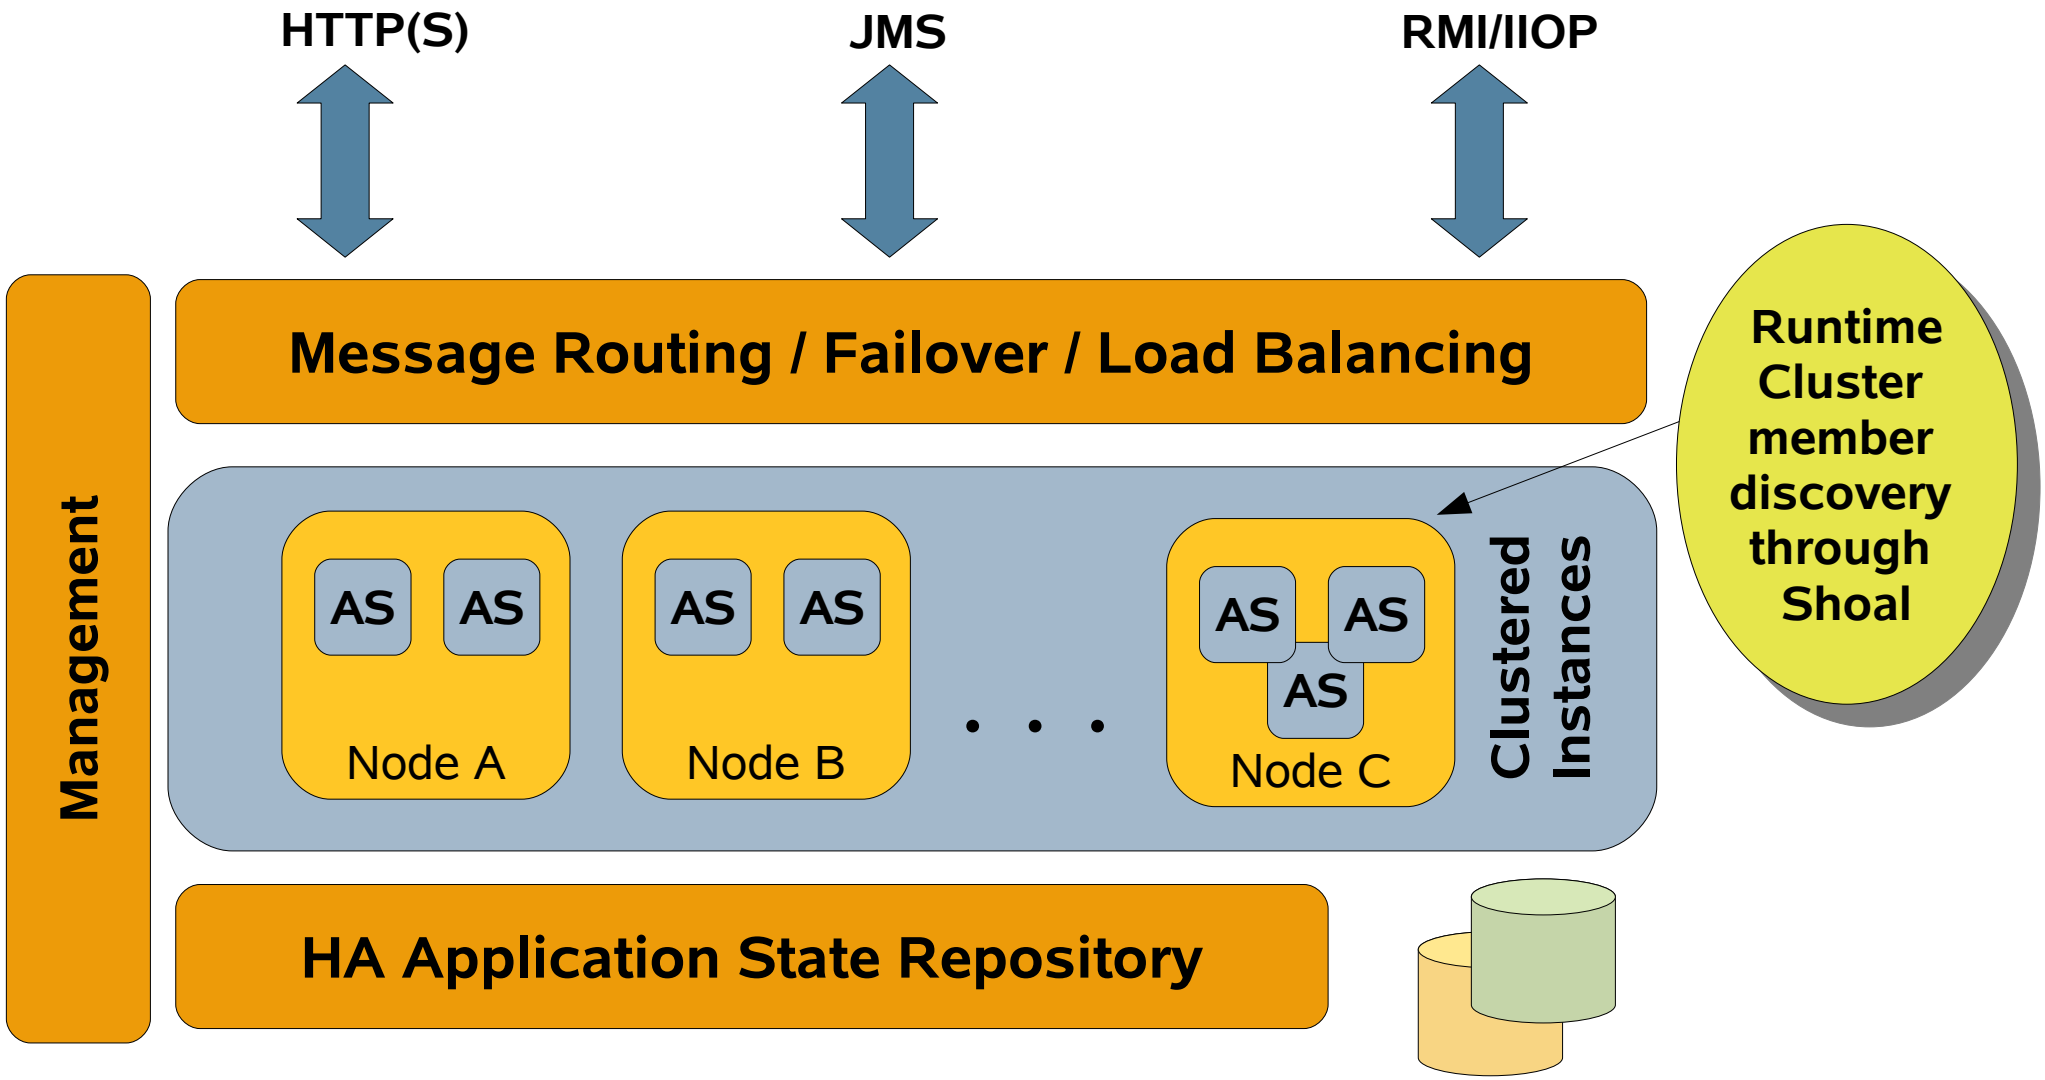

Runtime Dynamic Clustering – Project Shoal (https://shoal.dev.java.net)

- A Java language based clustering framework
- Pluggable into any product for clustering
- Provides an event model for Cluster events
  - Joining of instances
  - Failure suspicion
  - Failure confirmation
  - Planned Shutdown
  - Recovery Selection

- Provides messaging to group, or individual members
- Basis for building fault tolerance solutions
- GlassFish components such as In-Memory Replication module use Shoal.

GlassFish HA Options

- Enterprise Profile offers HADB based persistence solution
  - HADB
    - Enables proven “5—nines” availability
    - Highly reliable, scalable, but with lower performance
    - Right solution for mission critical, high reliability needs
    - More complex to administer, not (currently) open source
- Cluster Profile offers In-Memory Replication
  - In-memory replication
    - Ease to configure
    - A lighter-weight alternative to HADB
    - Provides high performance and scalability, lower reliability than HADB
    - Open Source – part of GlassFish project
    - Ideal for services for whom some risk of session loss is acceptable

GlassFish HA Options

- HADB and In-memory Replication solutions provide high availability for
  - Http session state
  - Stateful Enterprise JavaBeans™ (EJB™) technology session bean state
  - Single sign-on state
- GlassFish also provides JMS High Availability
  - GlassFish comes with a bundled Message Queue product
    - Sun Java System Message Queue
  - Message Queue provider integration is through resource adapters
  - Message Queue Broker clusters provide redundancy and availability of service and data
  - Data is persisted in HADB

MQ High Availability

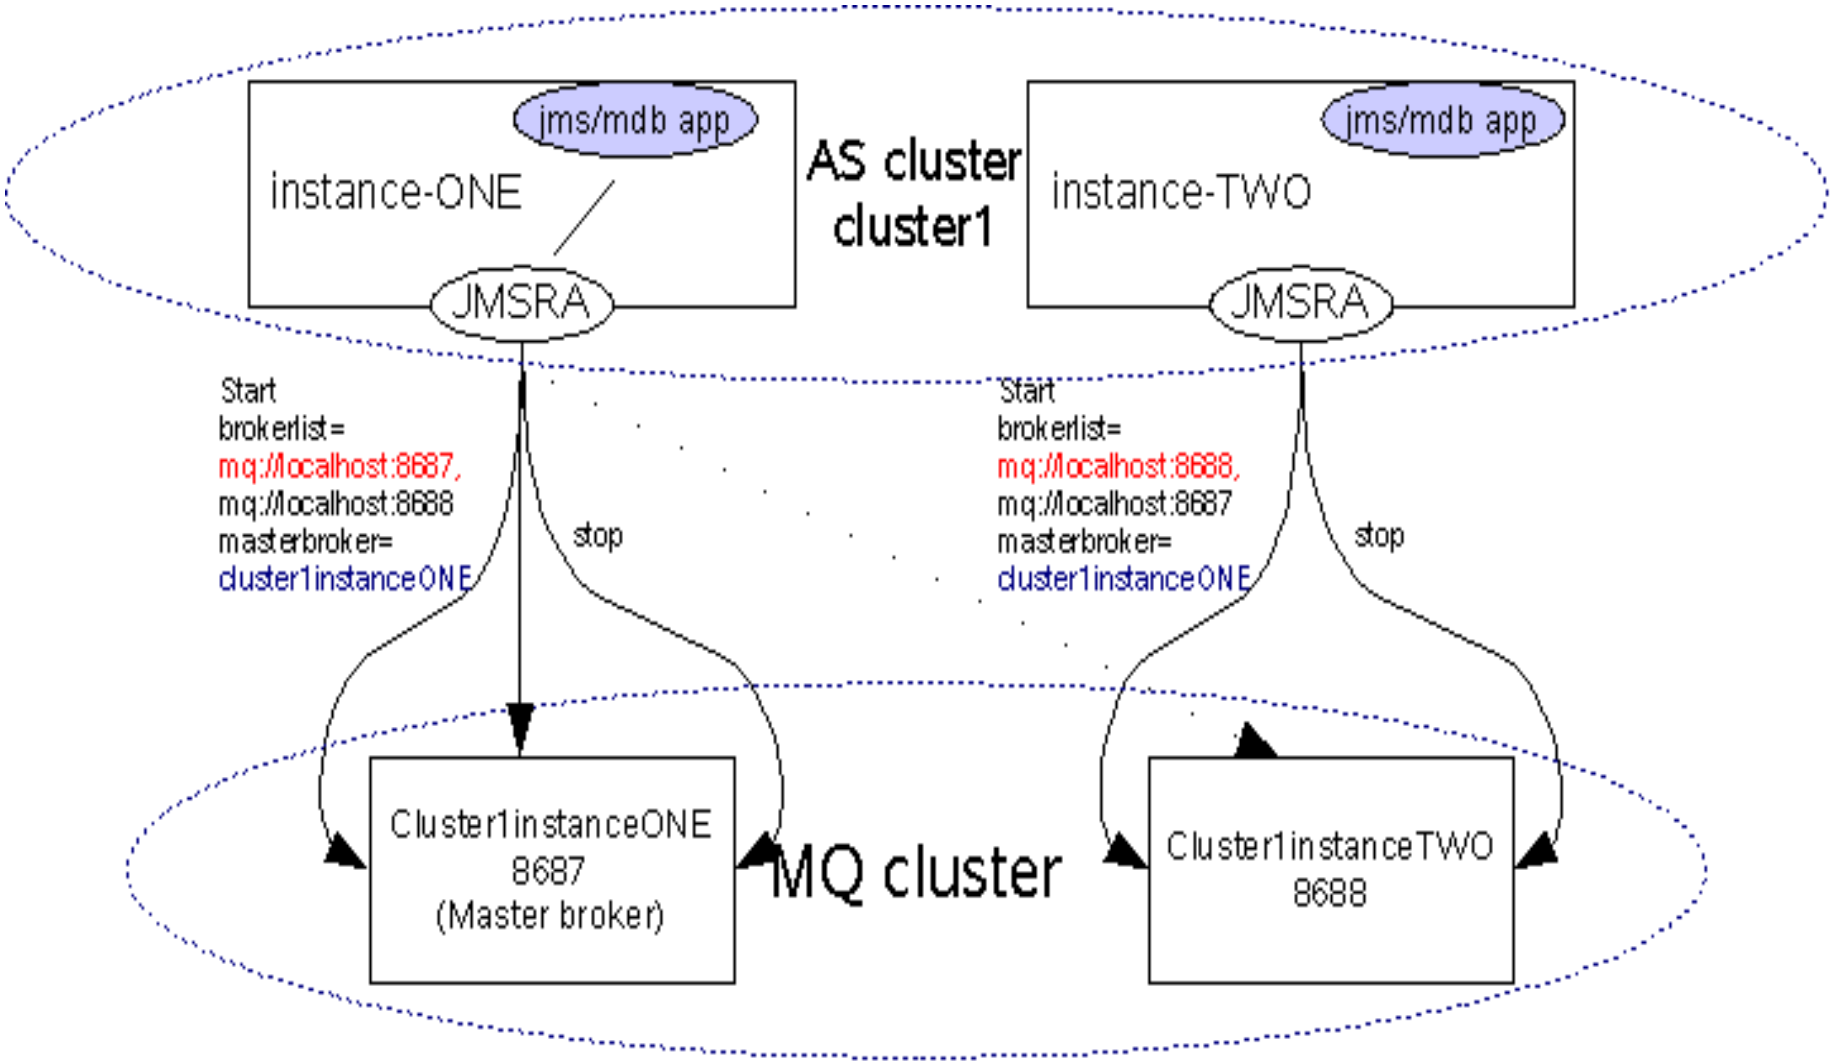

GlassFish In-Memory Replication Architecture - Internals

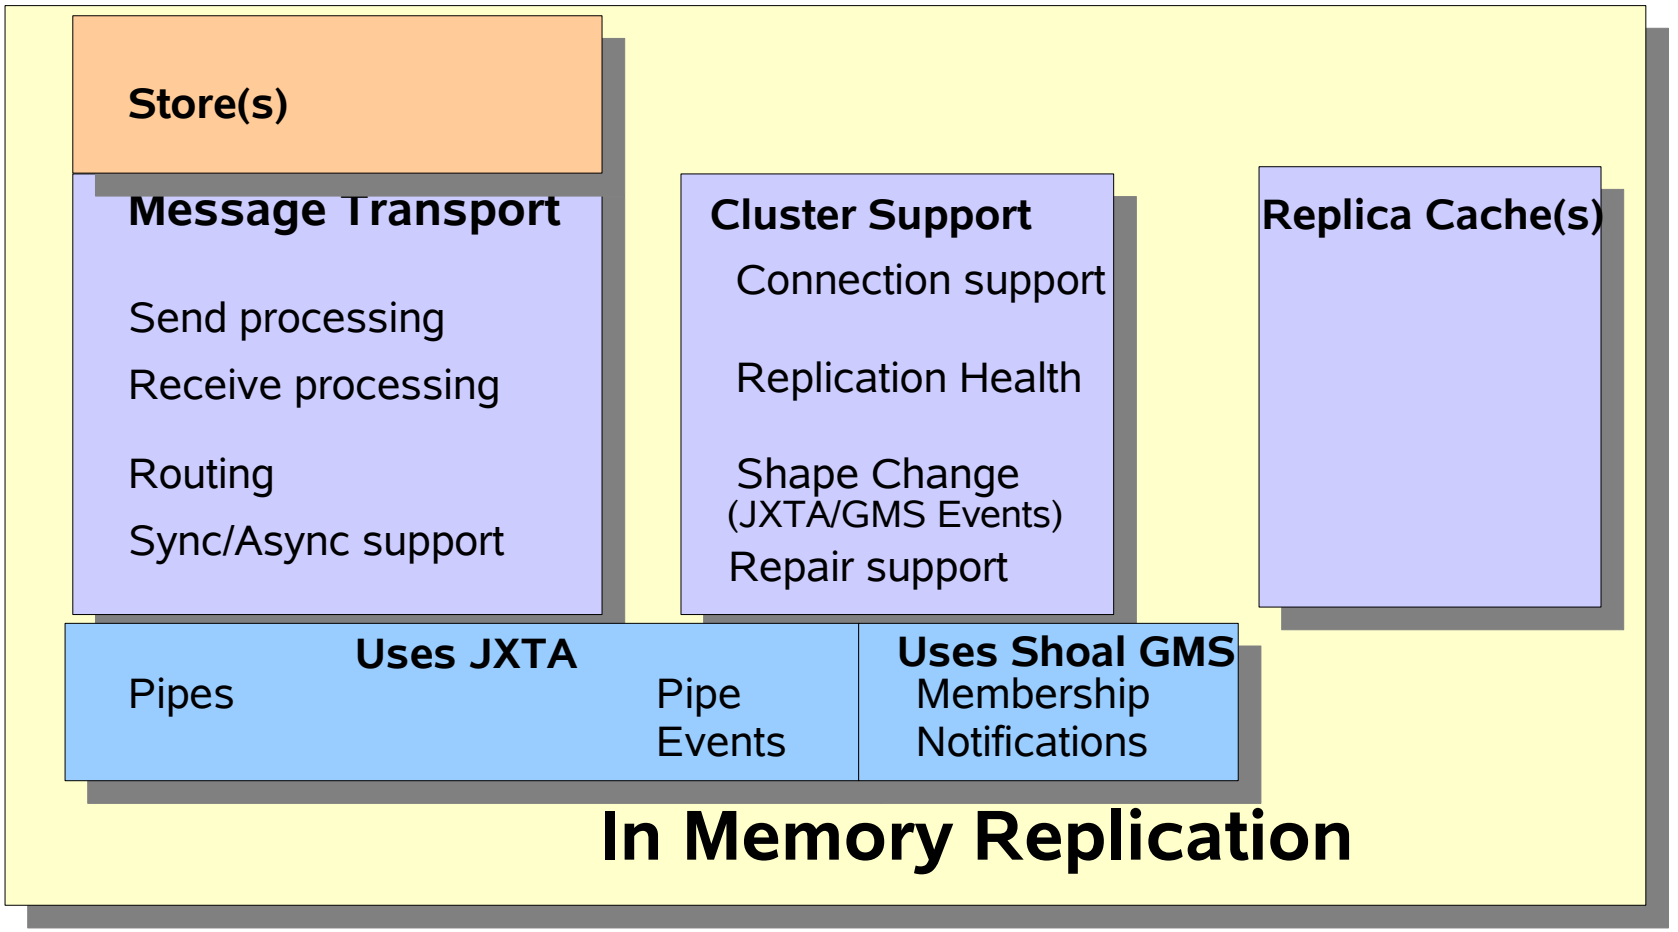

GlassFish Memory Replication Architecture - Container interactions

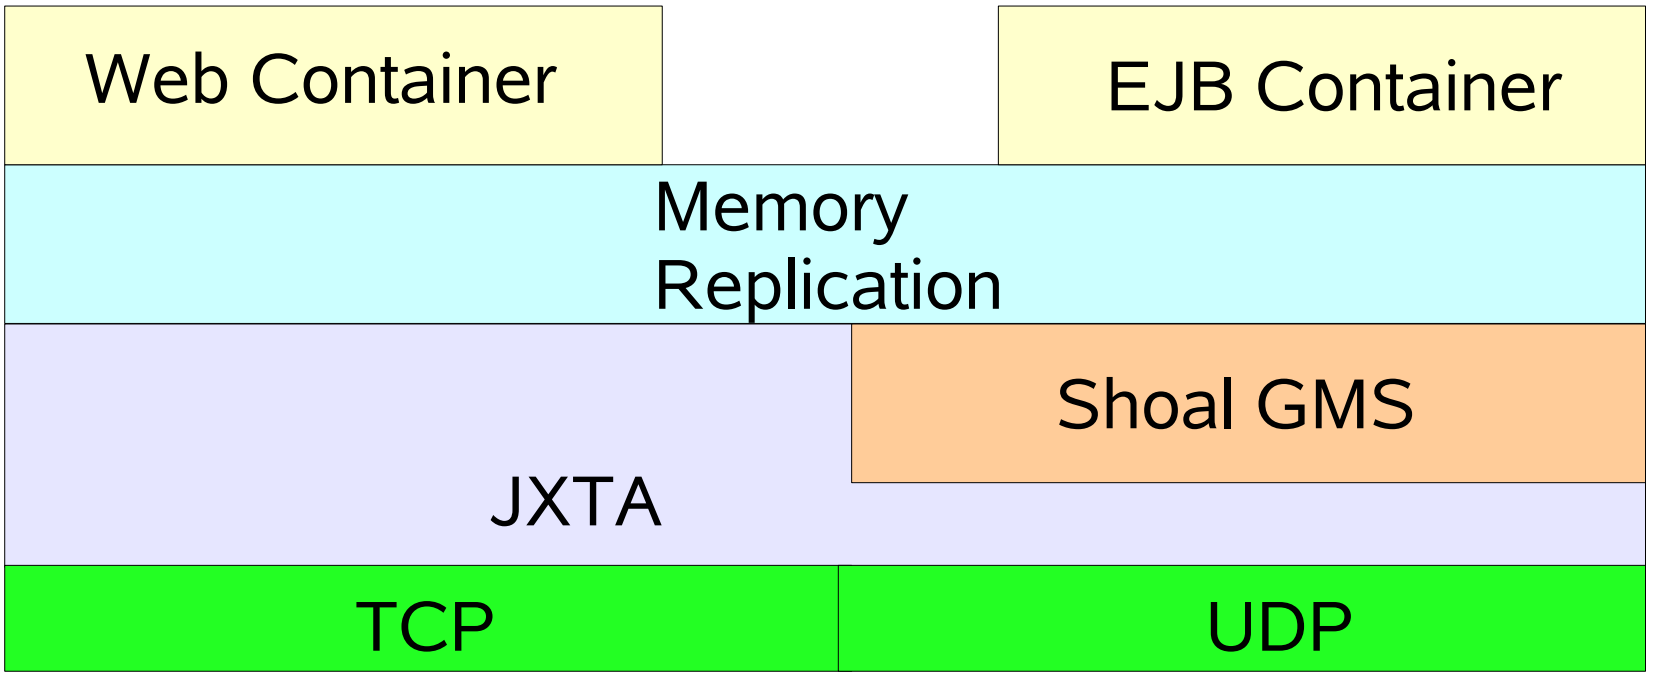

GlassFish Memory Replication Architecture

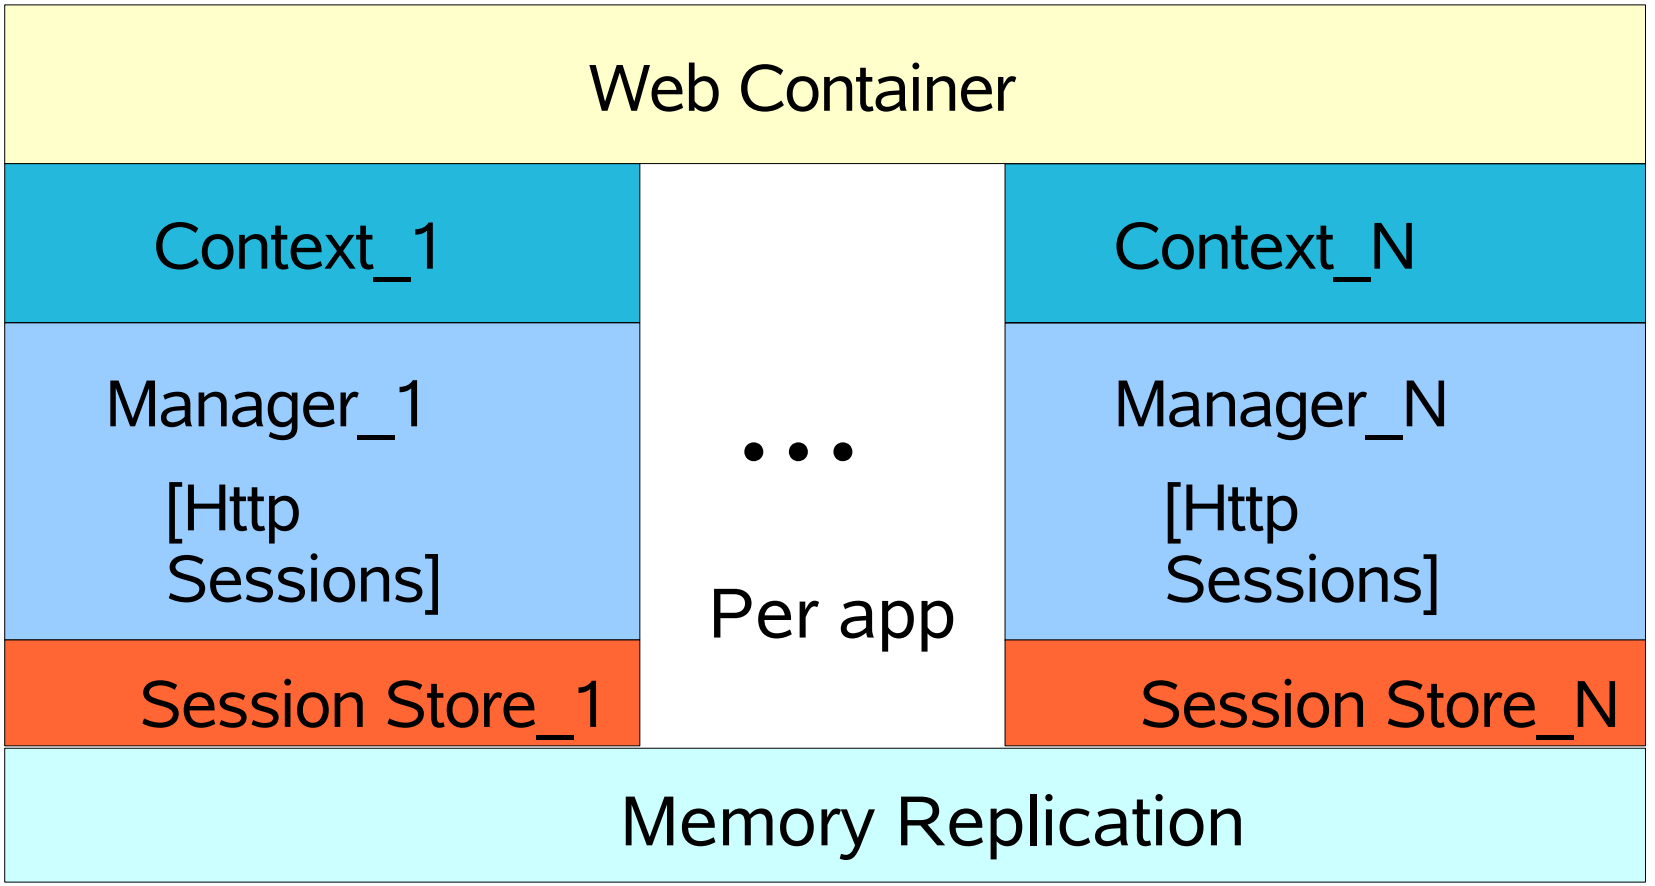

EJB containers are similar

Memory Replication

Typical cluster topology

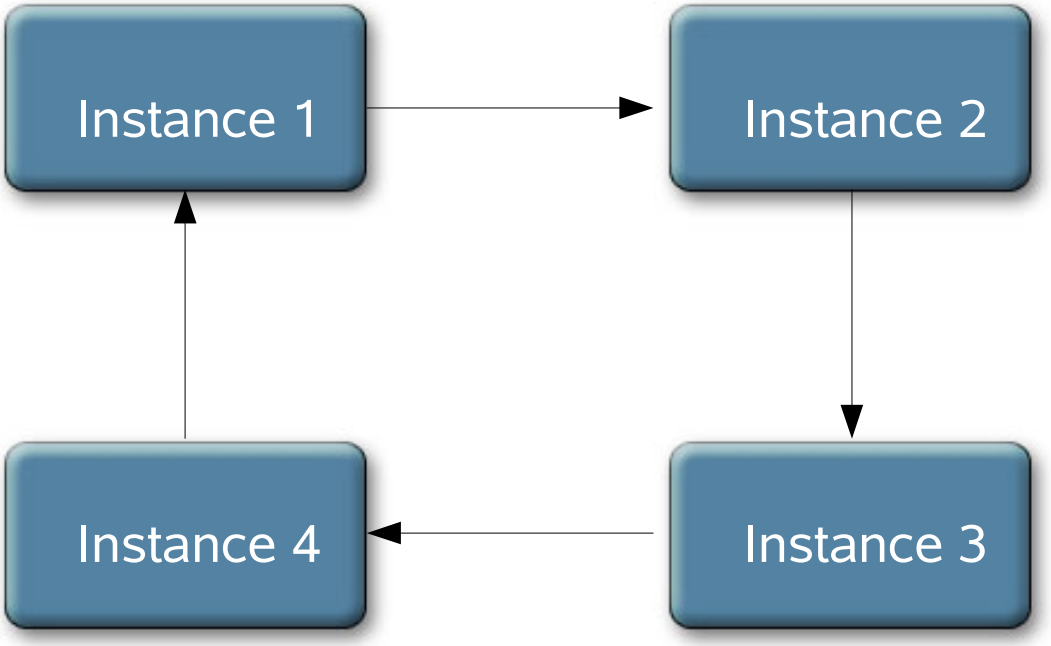

Memory Replication

Typical cluster topology

Example: Maximize Availability on 4 instance cluster on 2 nodes(machines)

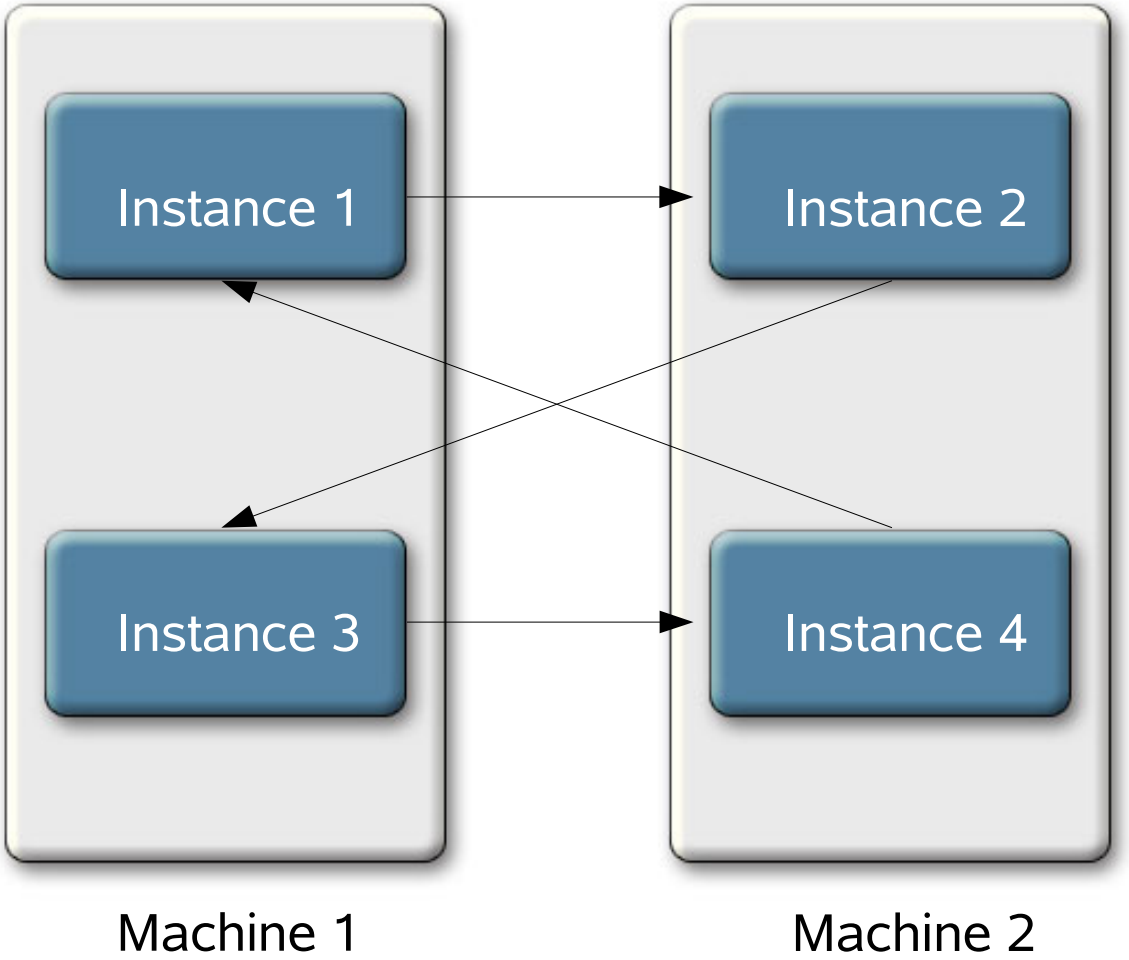

HTTP Session State Failover

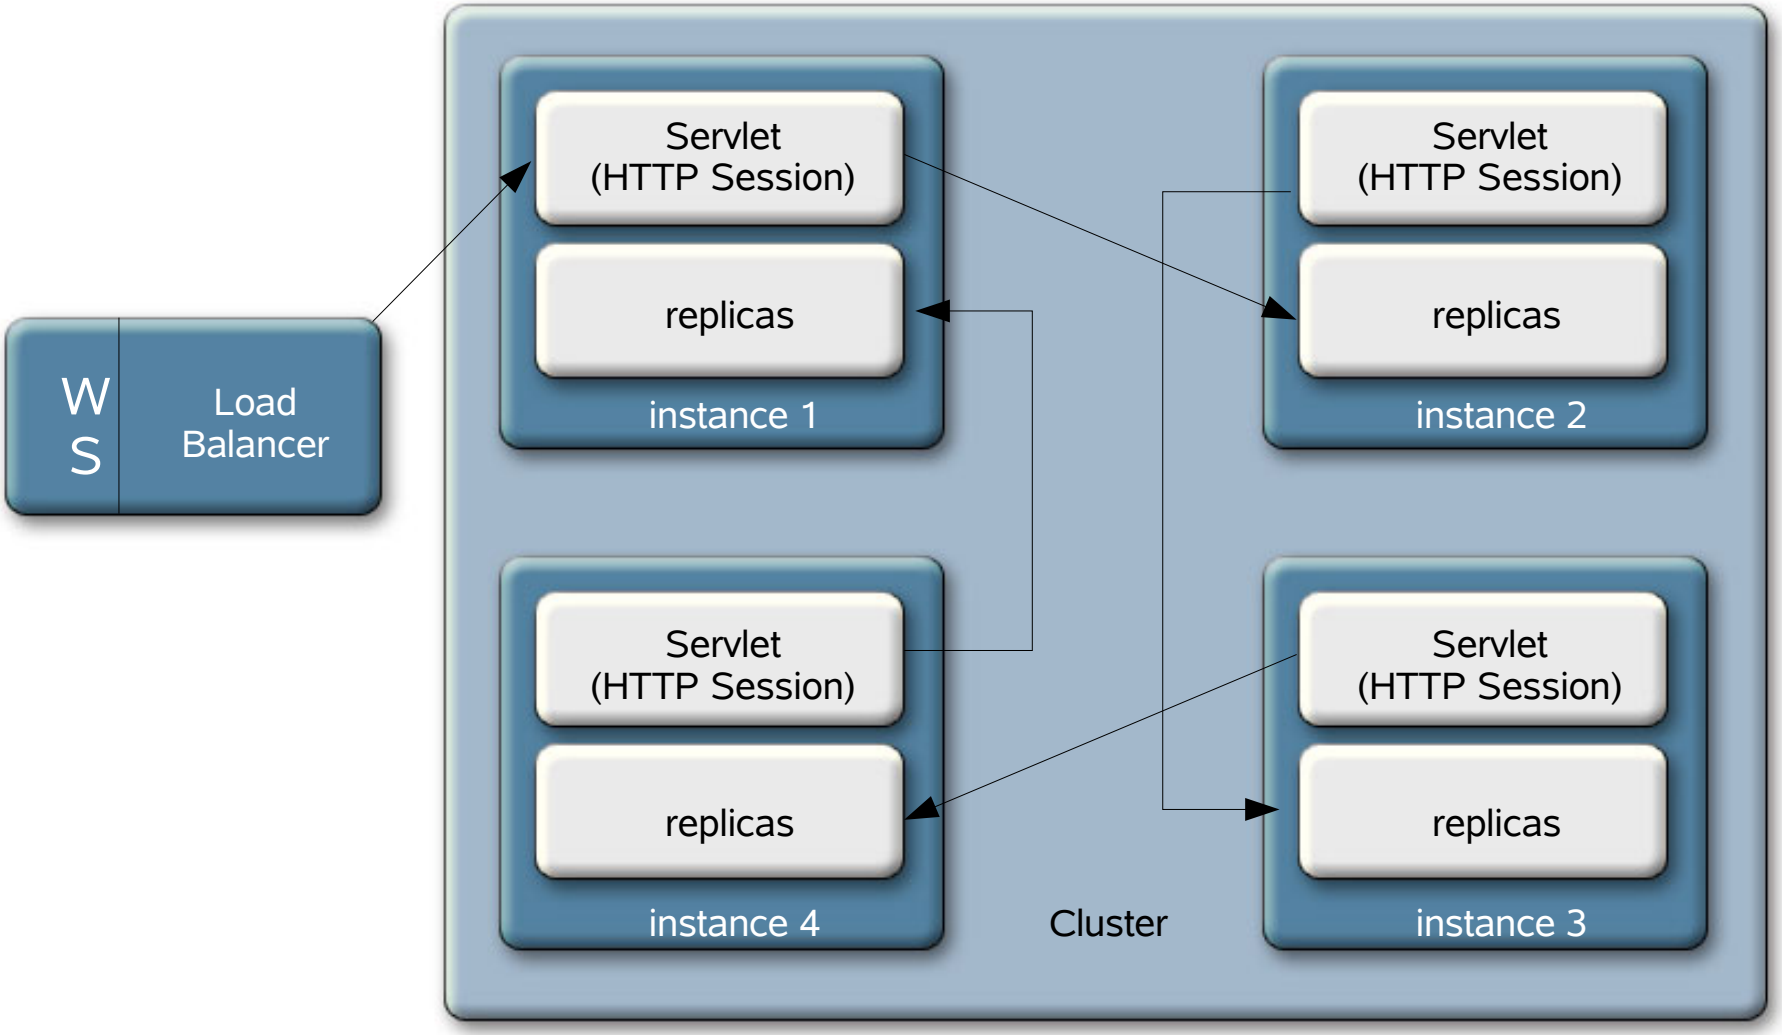

HTTP Session State Failover

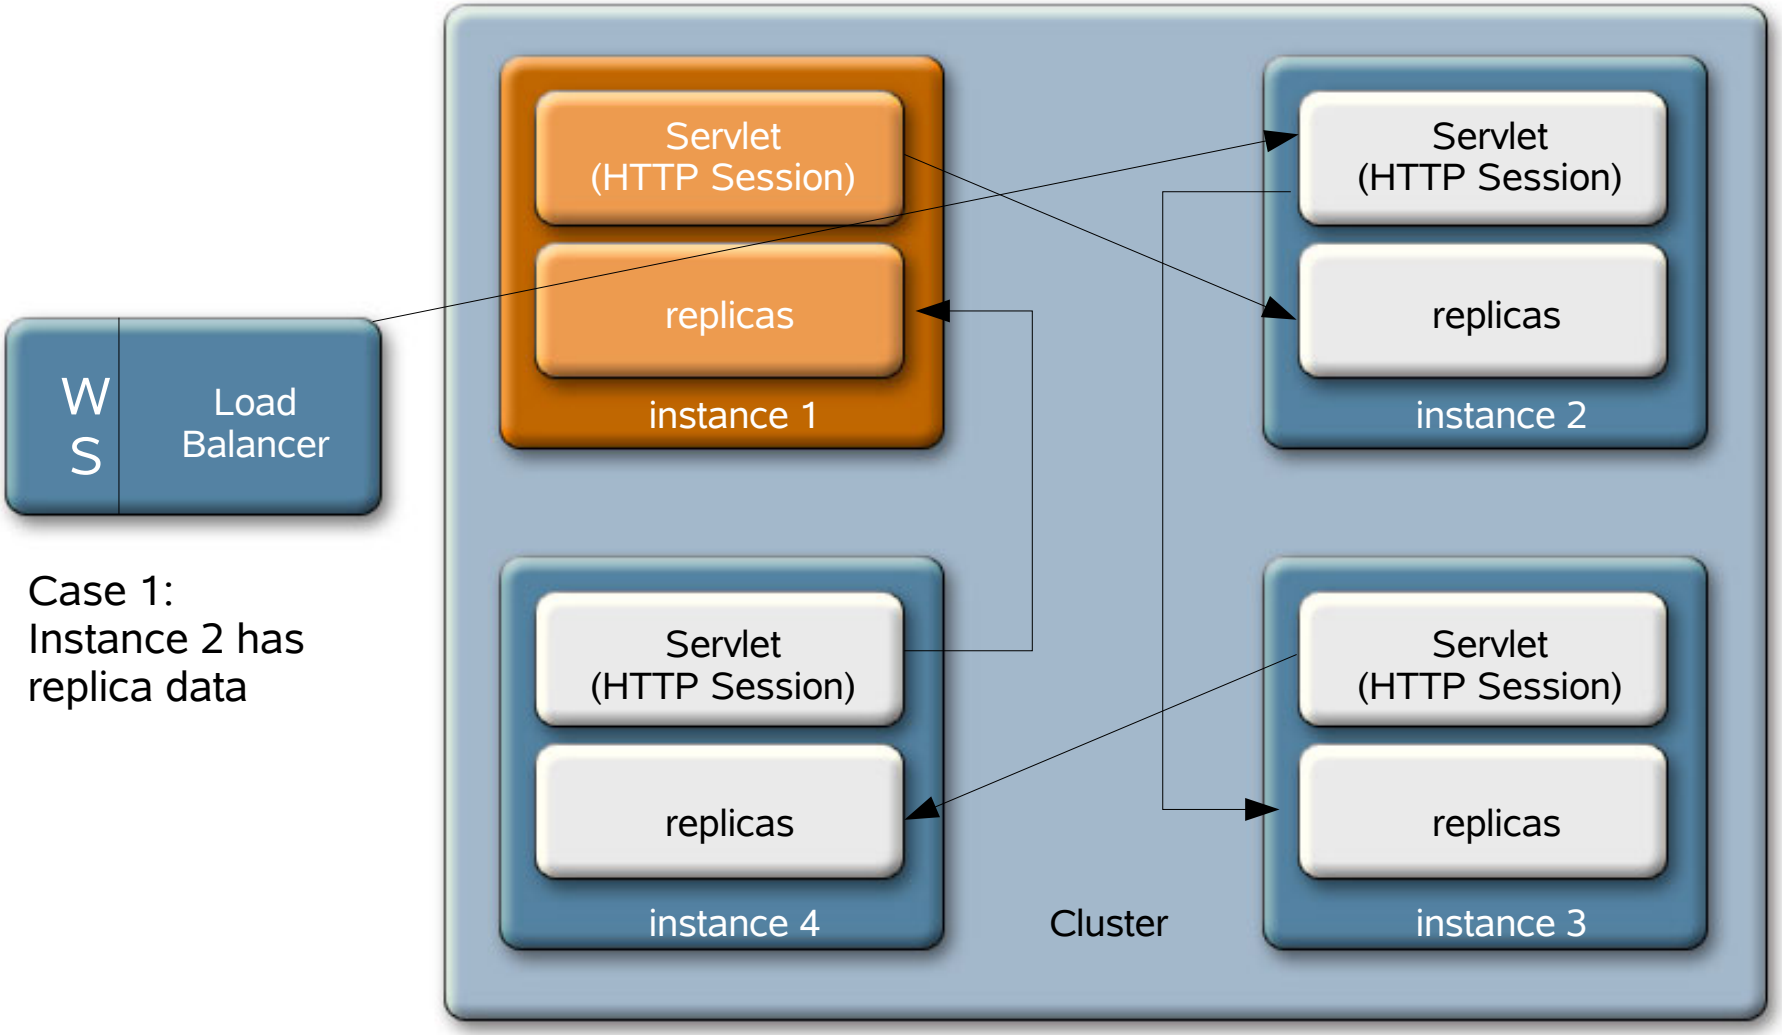

Case 1:
Instance 2 has
replica data

HTTP Session State Failover

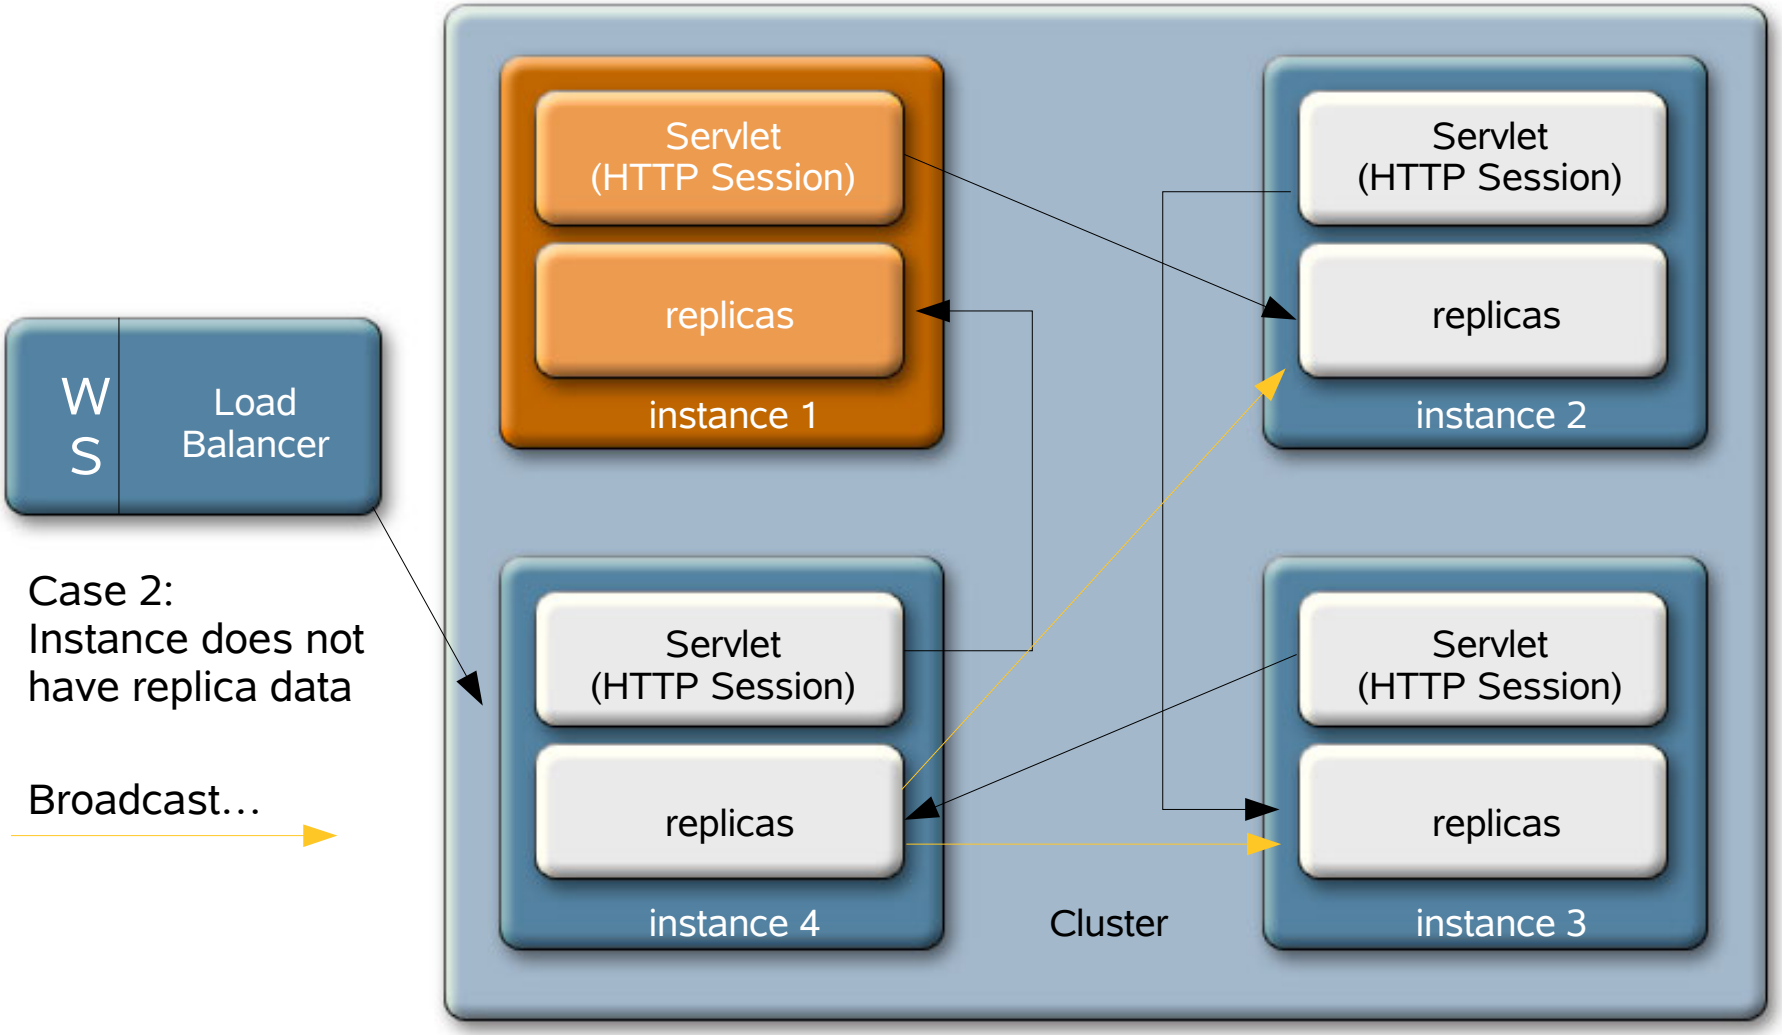

HTTP Session State Failover

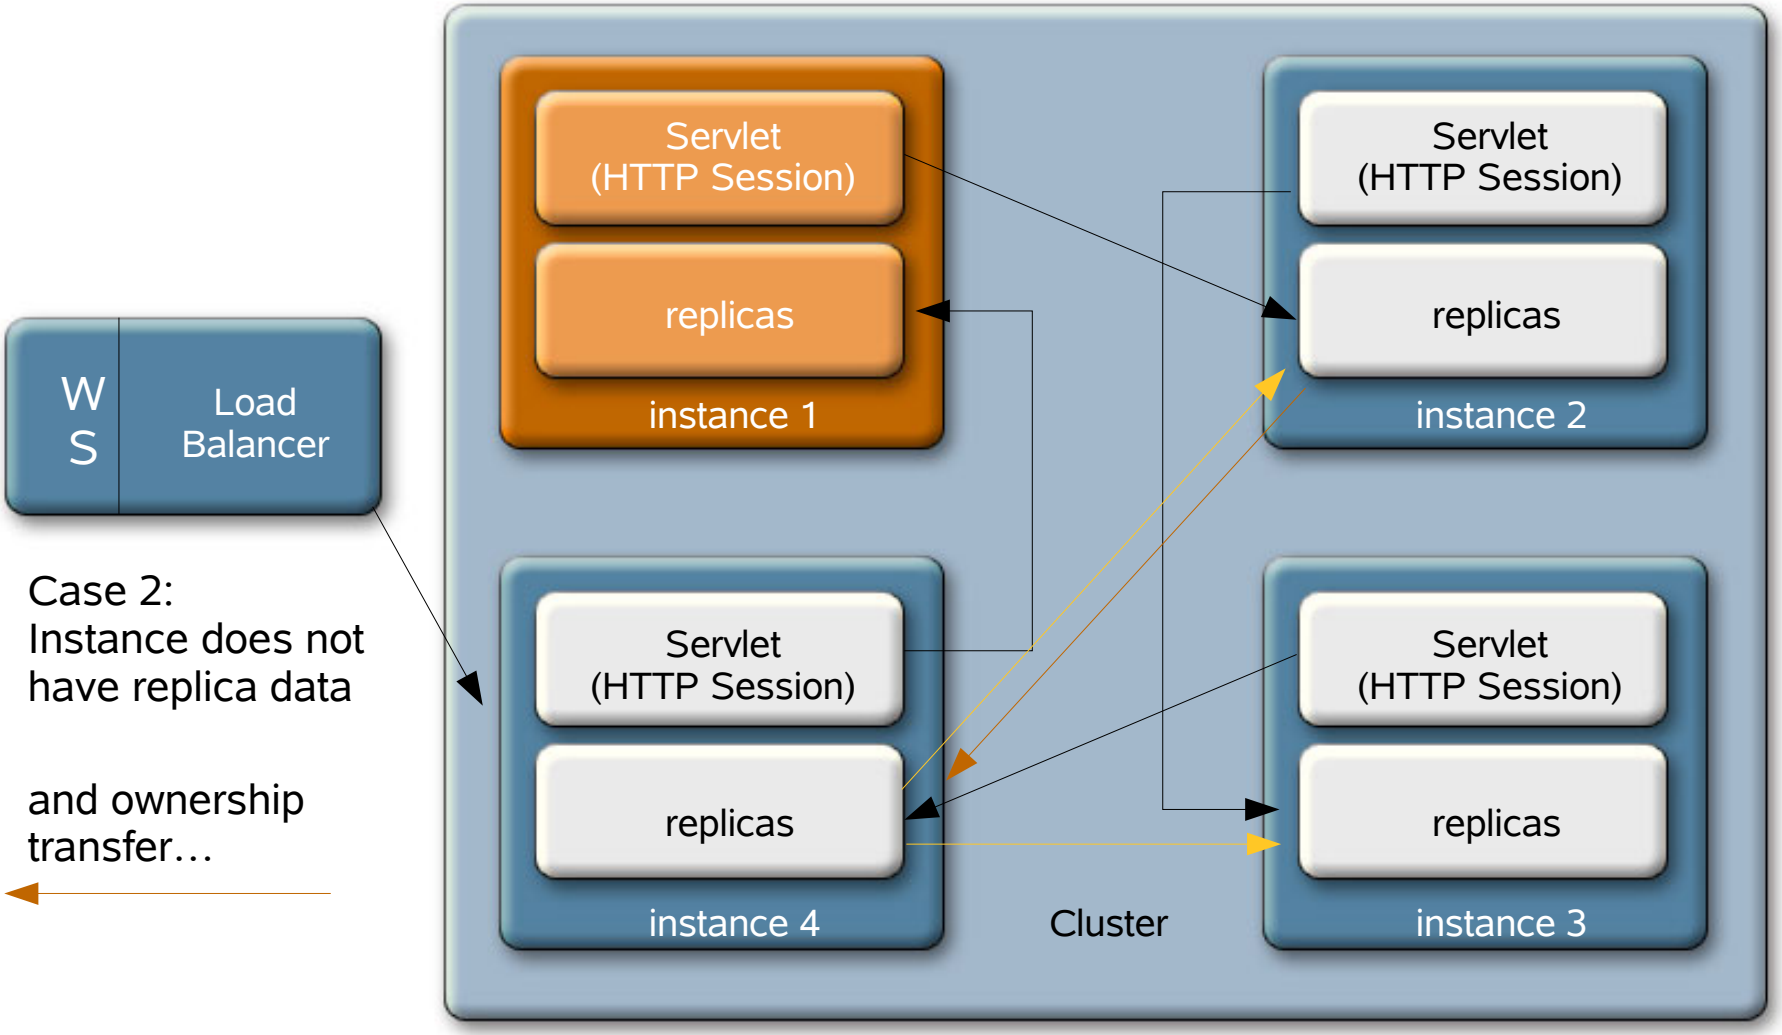

Case 2:
Instance does not have replica data

and ownership transfer…

Cluster Dynamic Shape Change

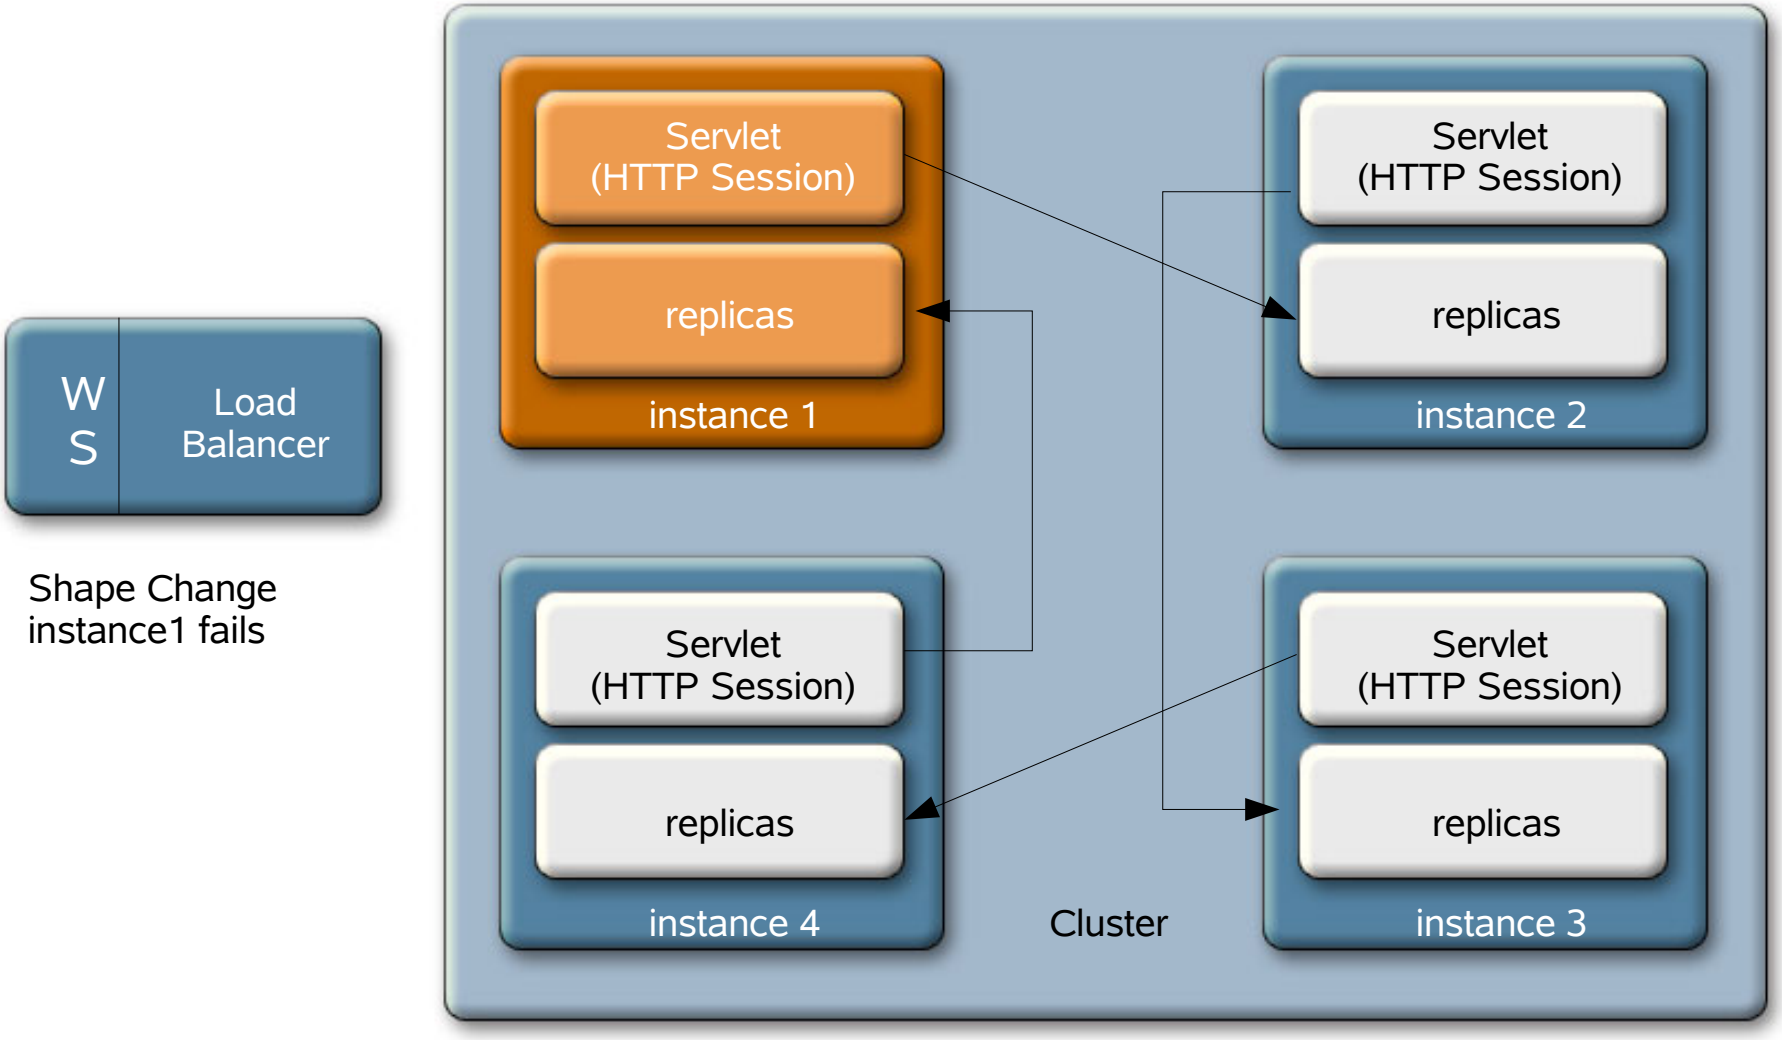

Cluster Dynamic Shape Change

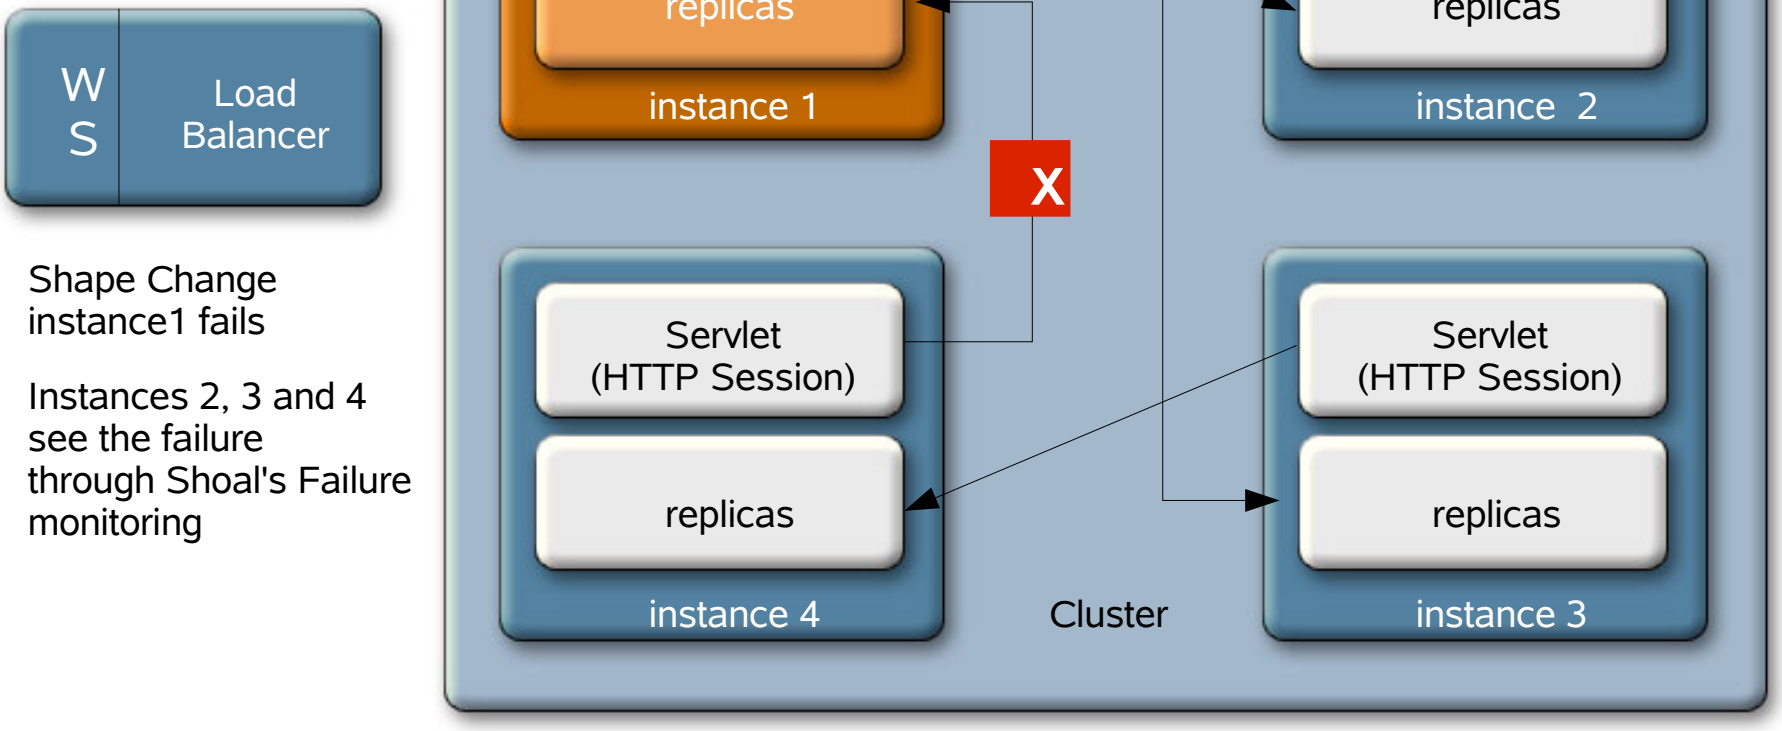

Shape Change
instance1 fails

Instances 2, 3 and 4
see the failure
through Shoal's Failure
monitoring

Cluster Dynamic Shape Change

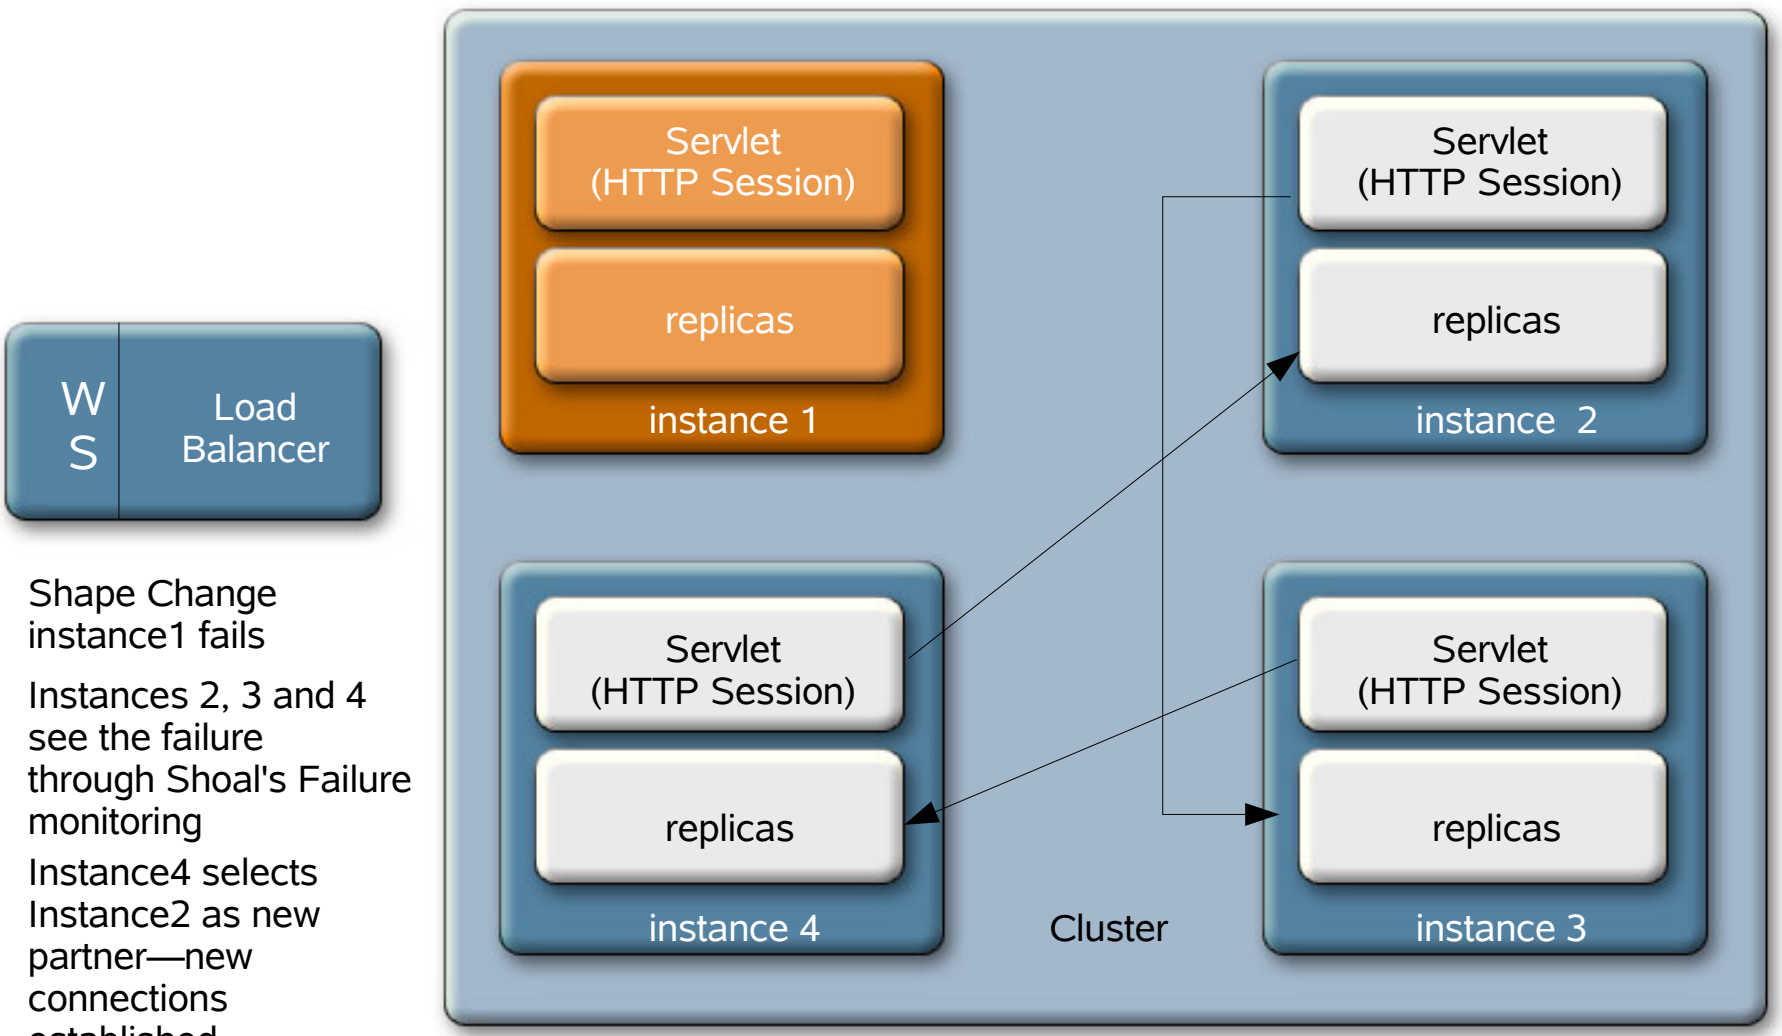

Shape Change
instance1 fails

Instances 2, 3 and 4 see the failure through Shoal's Failure monitoring

Instance4 selects Instance2 as new partner—new connections established

the reverse happens when an instance joins or re-joins the cluster

Memory Replication Configuration

Our hope was to say…

- "This page left intentionally blank" ;-)
  - Meaning "zero configuration required"
- We came close to that goal…

Memory Replication Configuration

Out of the box…

- Create a domain
  - Use the ‘cluster’ admin profile—defaults for replication are handled
    - Enables Shoal GMS — heartbeat enabled
    - persistence-type = “replicated”
- Create a cluster and instances
- Deploy your application with availability-enabled=true
- That’s it
- Caveat:
  - make sure your cluster is all in the same subnet (Cross subnet support coming in next release)
  - Synchronize clocks on all machines involved. Important for expiration logic.

Memory Replication Configuration

Beyond “out of the box” admin cluster profile

- Increase heap size
  - Default is 512MB for cluster admin profile
    - To accommodate cluster demos on laptops, etc.
    - Too small for serious replication
  - Increase to 1GB recommended
    - <jvm-options>-Xmx1000m</jvm-options>
      <jvm-options>-Xms1000m</jvm-options>

Memory Replication Configuration

Making your app distributable

- <distributable/> element
  - Required in web.xml
    - Indicates you believe your application is ready to run in a cluster
- Serializable objects required
  - HTTP Session state
  - EJB technology Stateful Session Bean state

Memory Replication Configuration

3 Main Axes of Configuration – memory replication or HADB

- Persistence Type
- Persistence Frequency
- Persistence Scope

Memory Replication Configuration

3 Main Axes of Configuration

- persistence-type “where are you storing”
  - quality of Replication Storage
- Options
  - memory
    - No replication (default, but understand effect of admin profiles)
  - replicated
    - In-memory replication
  - ha
    - HADB replication
  - file
    - Store to a file. (developer can verify serializes correctly)
    - Not supported for production environments

Memory Replication Configuration

3 Main Axes of Configuration

- persistence-frequency WHEN to replicate
- Options
  - web-event
    - Replicate when returning http-response
  - time-based
    - Replicate every ***T*** seconds
    - Default of 1 minute
- Discuss tradeoffs between these options

Memory Replication Configuration

3 Main Axes of Configuration

- Persistence-scope - WHAT granularity to replicate
- Options
  - session - i.e.full-session
    - Always requiring replication. Simplest but least efficient.
  - modified-session
    - Replicate full session when session set/remove attribute called
    - Otherwise only update timestamp
  - modified-attribute
    - Fine grain. Replicate at attribute granularity.
      - Sequence of Add, Delete, Update attribute(s)
    - Each “bucket” must be separately serializable with no x-references

Memory Replication Configuration

EJB stateful session replication story

- Again, feature compatible for HADB or replication
- replication called “checkpointing”
  - Passivation also causes replication
- Done at end of transaction – when there is a transaction
  - Container-managed or bean-managed
- Only for non-transactional beans
  - Particular methods can be marked for checkpointing using sub-ejb-jar.xml

Reference Resources

- GlassFish High Availability Guide
http://docs.sun.com/app/docs/doc/819-3679
- JMS Availability article of Interest
- GlassFish users mailing list
https://glassfish.dev.java.net/servlets/ProjectMailListList
- Clustering blogs http://blogs.sun.com/main/tags/clustering
- GlassFish Cluster/Developer Profile Downloads :
https://glassfish.dev.java.net/downloads/v2ur1-b09d.html
- GlassFish Enterprise Profile Downloads at sun.com

GlassFish High Availability Overview

Shreedhar Ganapathy