

ORACLE®

ORACLE®

Coherence 12.1.2 – Hidden Gems

Harvey Raja

Principal Member Technical Staff

Oracle Coherence



The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remain at the sole discretion of Oracle.

Coherence Roadmap



Apr 2011

Coherence 3.7.0

- Automatic Proxy Discovery for Clients
- Auto Re-connect for Clients
- Dynamic Load Balancing for Clients
- XML Schema for Config
- Load Balancer Integration (F5)
- Native Coherence*Web Glassfish Integration
- Query Monitoring
- Partition-Level Transactions
- Elastic Data

Sept 2011

Coherence 3.7.1

- POF Enhancements
- Query Explain Plan
- REST API
- Pluggable Partitioning Schemes
- Elastic Data Improvements (Journal-based Flash Storage)
- Delta Backups
- Leverage Exalogic Exabus Technology

NOW!!

Coherence 12c (12.1.2)

- Golden Gate Adapter for Coherence
- REST security and usability improvements
- Live Events
- Configuration Modernization
- Asynchronous Backups
- Backup Management Improvements
- Maven Support
- Exalogic performance optimizations
- Coherence Container
- Dynamic Thread Pooling for Proxy Servers
- OUI/Opatch Integration
- ECID Support
- OSGi Support



ORACLE

Coherence 12.1.2

- New major release of Coherence 12c
- Key Themes
 - Container Management with WLS
 - Continued Investment in Exalogic
 - Database Synchronization
 - Configuration and Usability Improvements
 - Oracle Fusion Middleware Convergence

2013

Coherence 12c (12.1.2)

- Golden Gate Adapter for Coherence
- REST security and usability improvements
- Live Events
- Configuration Modernization
- Asynchronous Backups
- Backup Management Improvements
- Maven Support
- Exalogic performance optimizations
- Coherence Container
- Dynamic Thread Pooling for Proxy Servers
- OUI/Opatch Integration
- ECID Support
- OSGi Support

A Deeper Look

- Asynchronous Backups

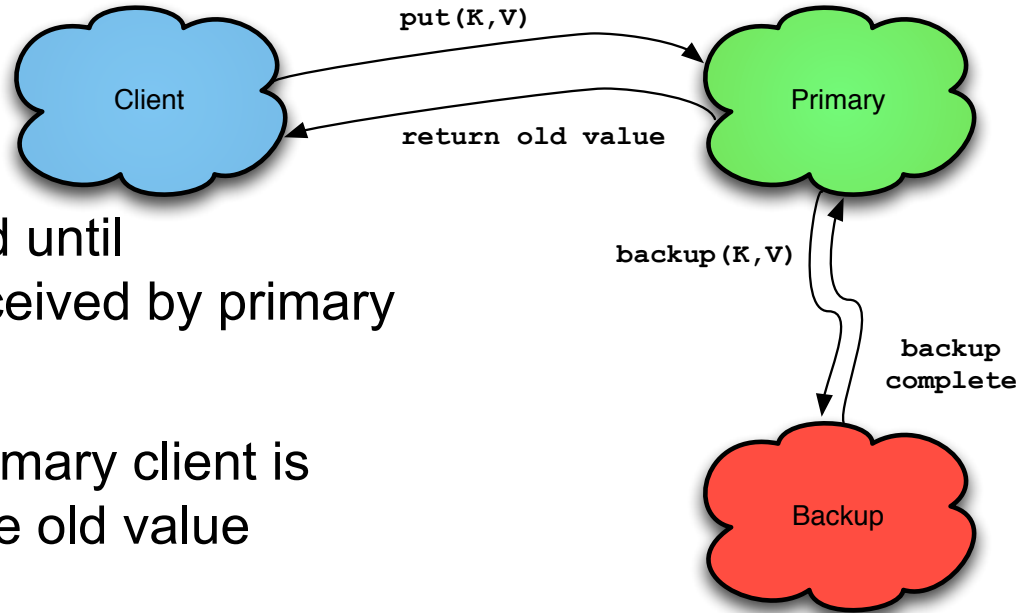
- Backup Management Improvements

Hidden Gems

- NameService
- TransactionEvent in Live Events
- POF Configuration Generator
- BinaryEntry synthetic ops
- Poll logging
- TcpRing improvements
- SLF4J native support
- WKA address resolution carried out on separate thread
- Preprocessing
- NearCache invalidation strategy

Current Backup Approach

- A put request from the client is entirely synchronous
- Client thread blocked until backup message received by primary
- Once received by primary client is responded to with the old value
- Provides consistency guarantees



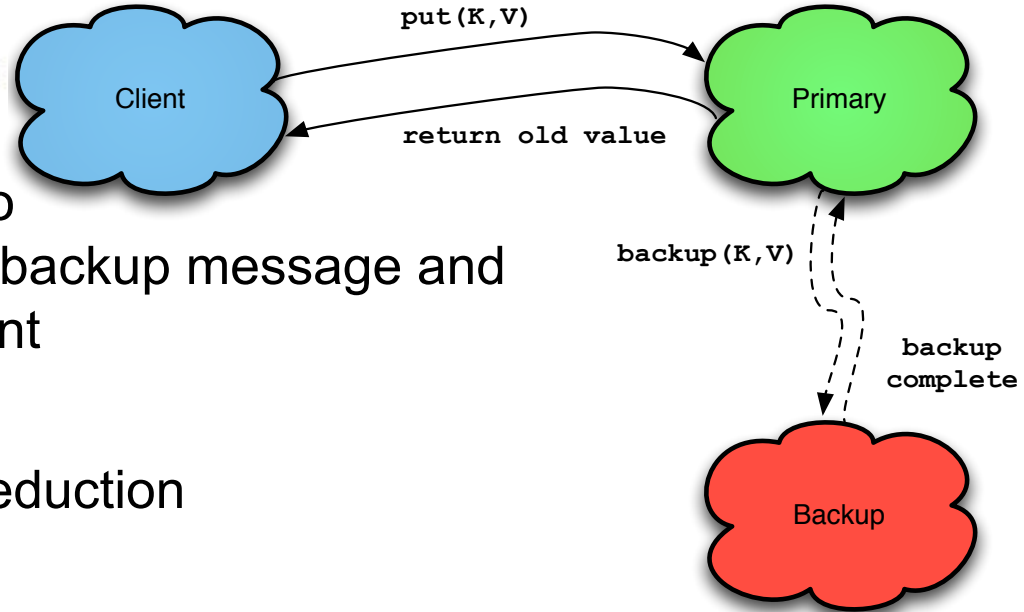
12c Backup Approach (Asynchronous)

- Opt-in strategy

```
<distributed-scheme>  
  <async-backup>true</async-backup>  
</distributed-scheme>
```

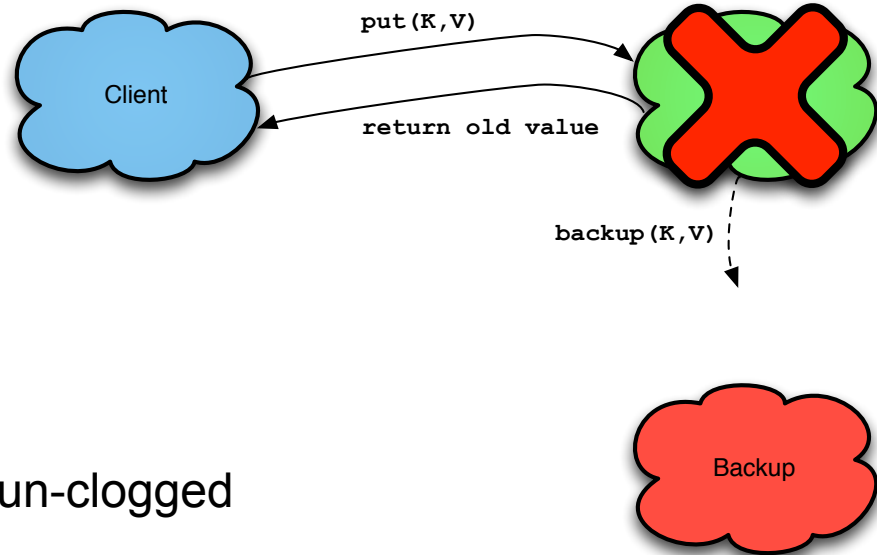
- Client request sent to primary & in parallel backup message and client response is sent

- ~40 – 50% latency reduction



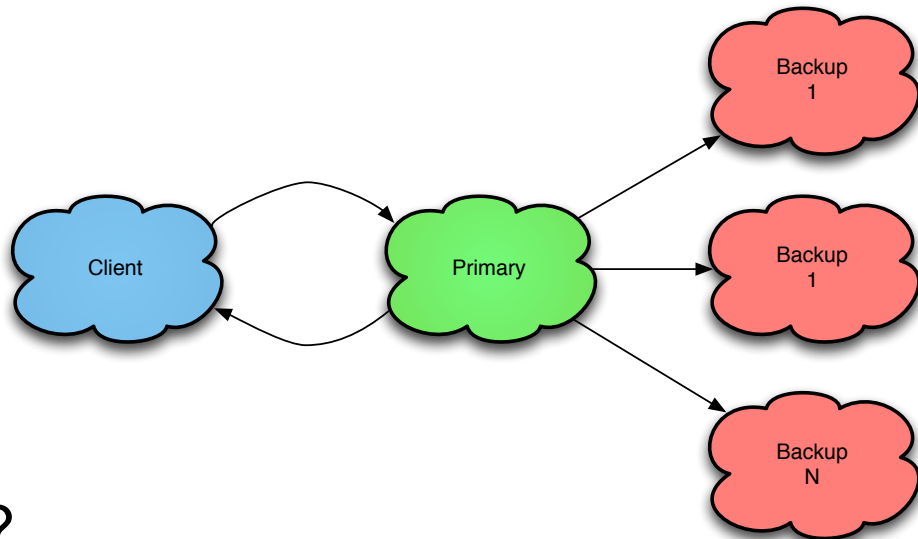
12c Backup Approach (Asynchronous)

- Speed vs Consistency
- Adaptive switch back to synchronous
 - Based on partition load
 - Falls back to async once un-clogged

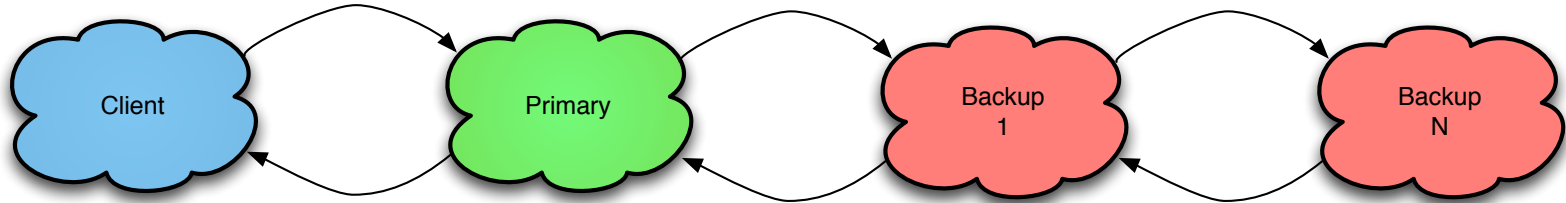


Pre-12c Backup Management

- Fan-out approach
- Requires coordination
- What if coordinator leaves?
 - Which backup received the message?



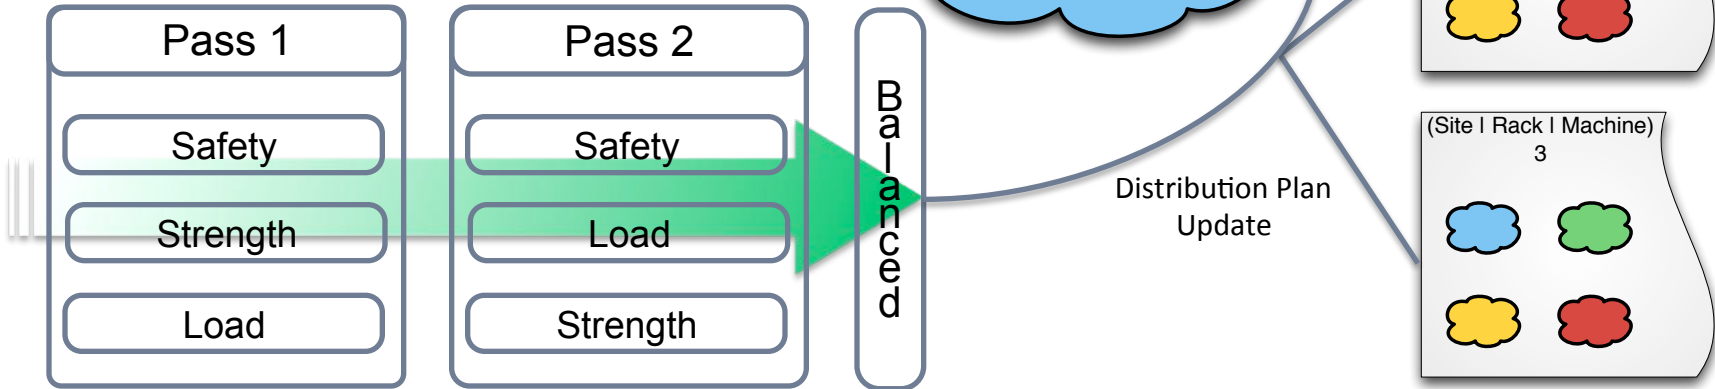
12c Backup Management



- Chained mechanics ensures prior node in the chain is responsible for ensuring delivery
- Departure of node in the chain ensures they are skipped
- Pre-process response from forwarded message

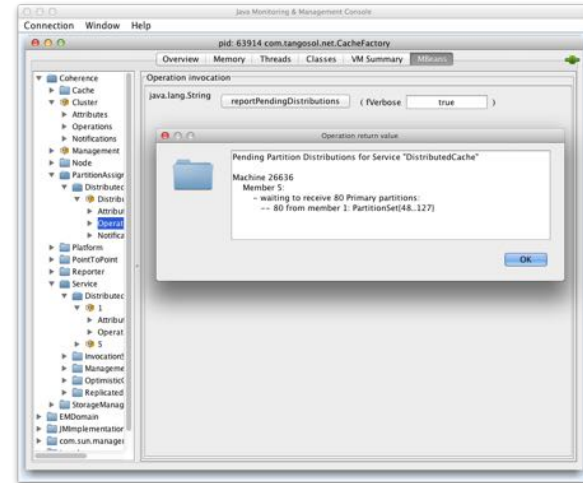
Backup Management Improvements

- As well as Machine, 12c now supports Rack & Site safety
- Goal is to reach highest level of safety



Backup Management Improvements

- Mirroring Strategy
 - Attempt to mirror assignments of another service
- JMX Enhancements
 - Pending distributions
 - PartitionLost JMX Notification



Coherence *Extend - NameService



Apr 2010

Coherence 3.6

- List of InetAddresses on the client

Apr 2011

Coherence 3.7.0

- Automatic Proxy Discovery for Clients
- Dynamic Load Balancing for Clients

2013

Coherence 12c (12.1.2)

- NameService

```
<remote-cache-scheme>
  <initiator-config>
    <tcp-initiator>
      <remote-addresses>
        <socket-address>
          <address>host1</address>
          <port>9000</port>
        </socket-address>
        <socket-address>
          <address>host1</address>
          <port>9001</port>
        </socket-address>
        <socket-address>
          <address>host2</address>
          <port>9000</port>
        </socket-address>
        <socket-address>
          <address>host2</address>
          <port>9001</port>
        </socket-address>
      </remote-addresses>
    </tcp-initiator>
  </initiator-config>
</remote-cache-scheme>
```

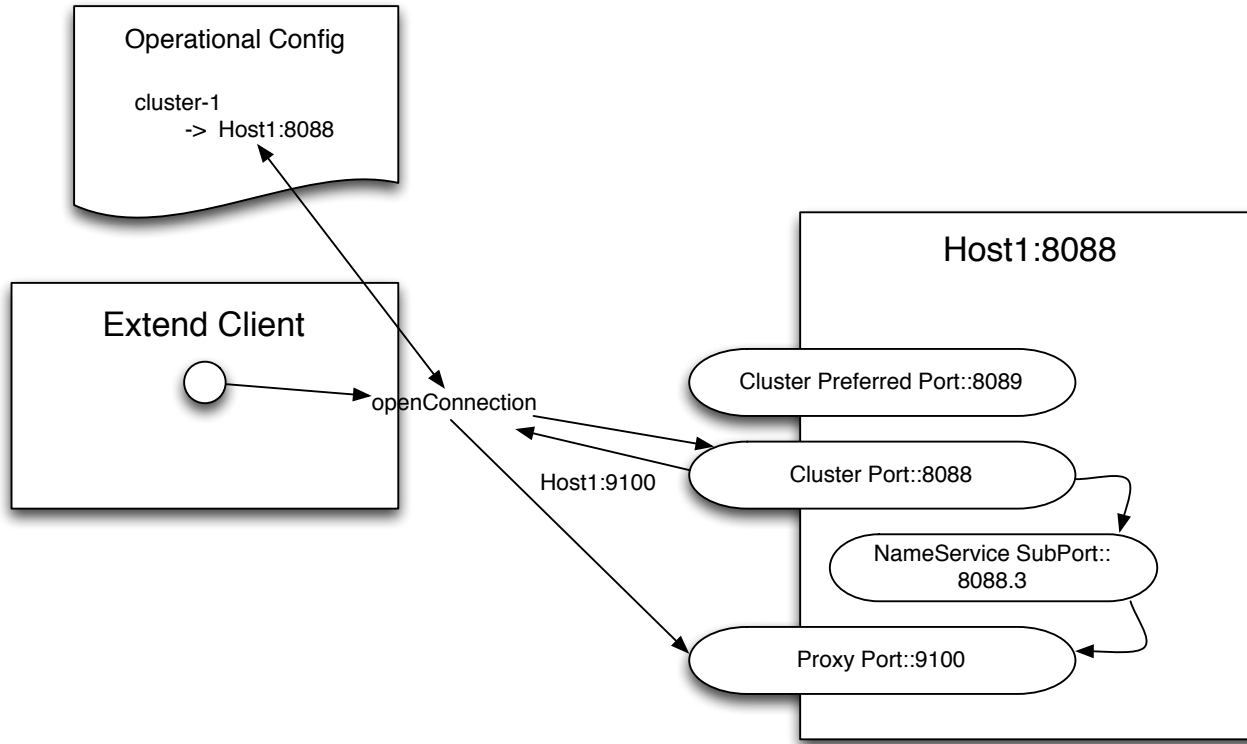
```
<remote-cache-scheme>
  <initiator-config>
    <tcp-initiator>
      <remote-addresses>
        <socket-address>
          <address>host1</address>
          <port>9000</port>
        </socket-address>
      </remote-addresses>
    </tcp-initiator>
  </initiator-config>
</remote-cache-scheme>
```

```
<remote-cache-scheme>
  <initiator-config>
    <tcp-initiator>
      <name-service-addresses>
        <address-provider>cluster-1</address-provider>
      </name-service-addresses>
    </tcp-initiator>
  </initiator-config>
</remote-cache-scheme>
```

Coherence *Extend - NameService

- Allows discovery of proxies using logical names (service names)
- Handshake takes place during connection establishment
- Uses well known cluster sub-port to discover proxy end points
- Could be considered as ClusterNameService (CNS \approx DNS)

Coherence *Extend - NameService



Coherence *Extend - NameService

```
<cluster-config>
  <address-providers>
    <address-provider id="cluster-1">
      <socket-address>
        <address>localhost</address>
        <port>8088</port>
      </socket-address>
    </address-provider>
  </address-providers>
</cluster-config>
```

```
<remote-cache-scheme>
  <service-name>MyProxy</service-name>
  <initiator-config>
    <tcp-initiator>
      <name-service-addresses>
        <address-provider>cluster-1</address-provider>
      </name-service-addresses>
    </tcp-initiator>
  </initiator-config>
</remote-cache-scheme>
```

```
<proxy-scheme>
  <service-name>MyProxy</service-name>
</proxy-scheme>
```

- Service name must be the same
- Discovers all proxy addresses for the same service name
- Proxy service can use ephemeral ports

LiveEvents – TransactionEvent

- Ability to intercept events as they occur in the grid
 - Fine grained events with logical causality
- 12.1.2 introduces a TransactionEvent
 - Partition Lite transaction event
- Receives all enlisted entries in a single event
- Can enlist more entries

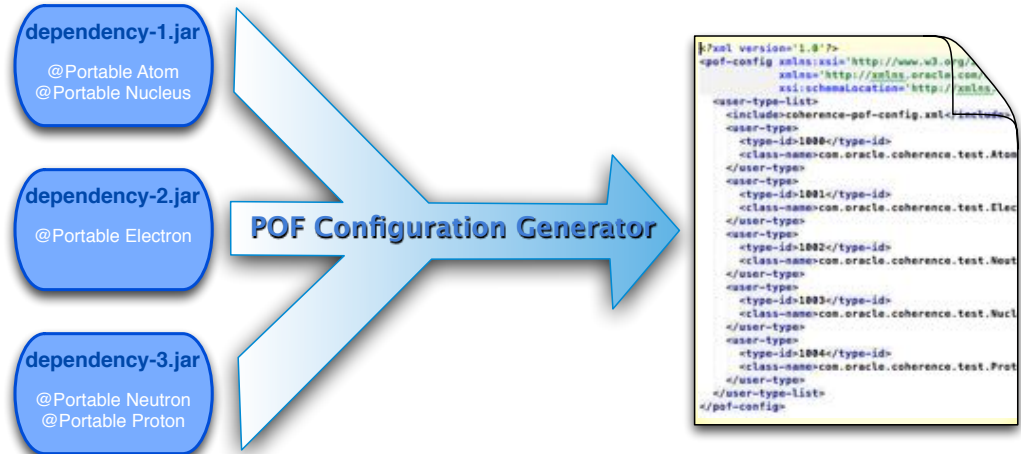
```
/**
 * A TransactionEvent captures information pertaining to all mutations
 * performed within the context of a single request. All modified
 * entries processed to the interceptor(s) of this event. All entries are
 * guaranteed to come from (glink.com.tangosol.net.PartitionedService), but may
 * also include entries from parent caches.
 *
 * @author rhl/Arjg 2012.09.21
 * @since Coherence 12.1.2
 */
public interface TransactionEvent
    extends Event<TransactionEvent.Type>
{
    /**
     * A set of (glink.BinaryEntry entries) enlisted within this
     * transaction.
     *
     * @return a set of entries enlisted within this transaction
     */
    public Set<BinaryEntry> getEntrySet();
}

/**
 * The TransactionEvent types.
 */
public static enum Type
{
    /**
     * A COMMITTING event is raised prior to any updates to the
     * underlying backing map. This event will contain all modified
     * entries which may span multiple backing maps.
     *
     * @see
     * @see The following holder:
     * @see
     * @see <code>The BinaryEntry instances passed for this event type are mutable.</code>
     * @see <code>A lock will be held for each entry during the processing of
     * this event, preventing concurrent updates.</code>
     * @see <code>Throwing an exception from this event will prevent the
     * operation from being committed.</code>
     * @see
     */
    COMMITTING,

    /**
     * A COMMITTED event is raised after any mutations have been
     * committed to the underlying backing maps. This event will contain
     * all modified entries which may span multiple backing maps.
     *
     * @see The BinaryEntry instances passed for this event type
     * are immutable.
     */
    COMMITTED
}
}
```

POF Configuration Generator

- Generates POF configuration file based on **@Portable** classes
- Predictable type-id generation
- Generational
 - Accepts previous POF configuration file
- Operates against a GAR and supported by maven GAR plugin

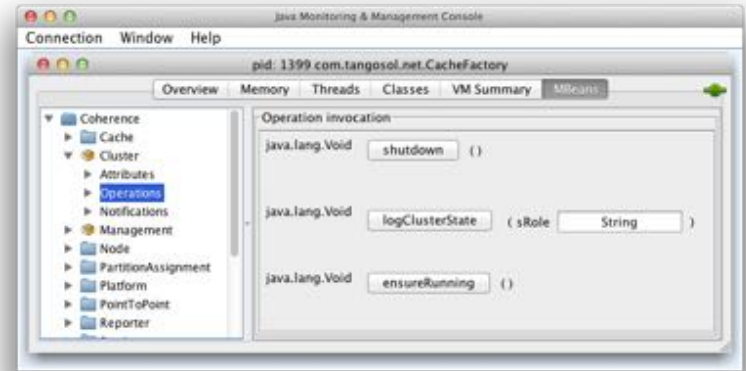


12c - Loose Change

- BinaryEntry synthetic operations
 - `updateBinaryValue(Binary, boolean) // new`
 - `setValue(Object, boolean)`
 - `remove(boolean)`
- SLF4J logger
 - SLF4J libraries must be on classpath and logger destination be `slf4j`
- WKA host name lookup now on a separate thread
- NearCache default invalidation strategy is now PRESENT

12c - Loose Change

- Log Cluster State
 - Role based
 - Distributed Thread Dump
 - Includes outstanding Polls
- Log Node State
 - Node Thread dump
 - Includes outstanding Polls
- Responsibility MBeans



12c - Loose Change

- TcpRing improvements
 - Death broadcast is communicated across the cluster
- Message preprocessing where possible

Demo



Q & A



Hardware and Software

ORACLE®

Engineered to Work Together

ORACLE®