ORACLE®

Oracle Coherence Technical Overview

Rob Misek

Senior Product Manager | 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.

Oracle Cloud Application Foundation

Core Capabilities for Modern Cloud Infrastructure

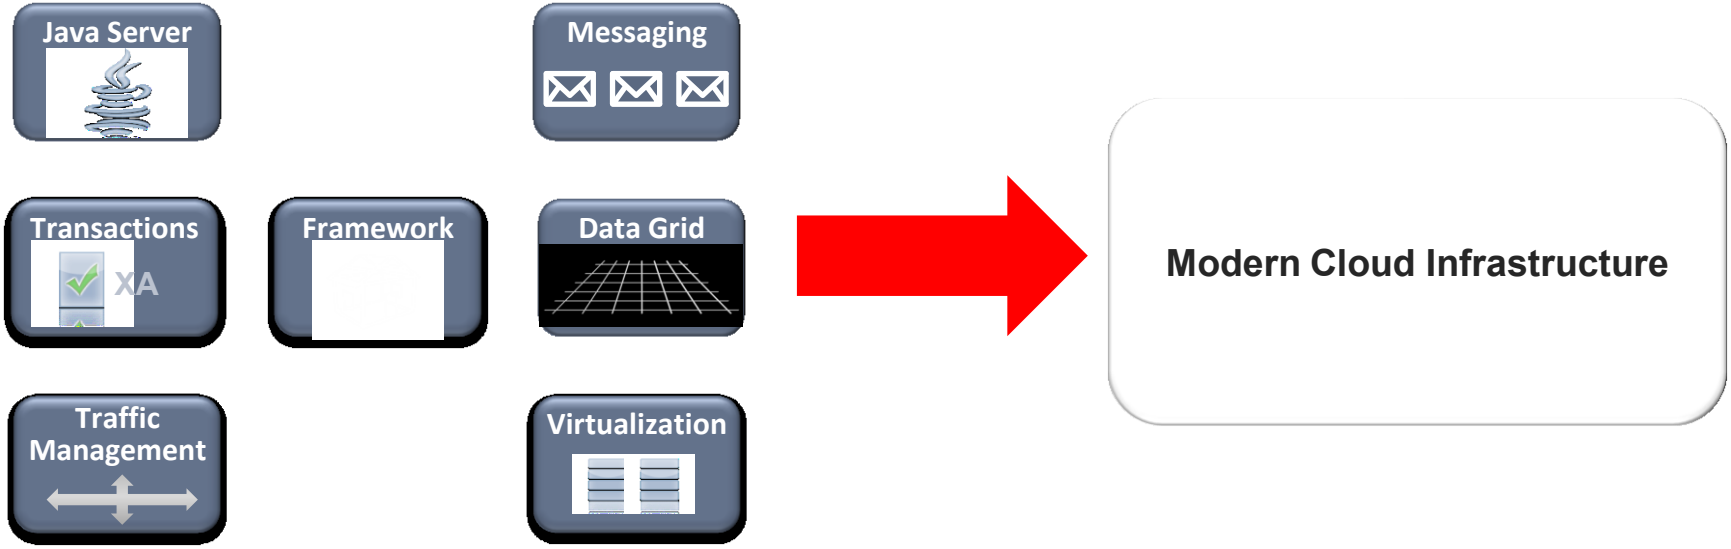

You Can Roll Your Own …

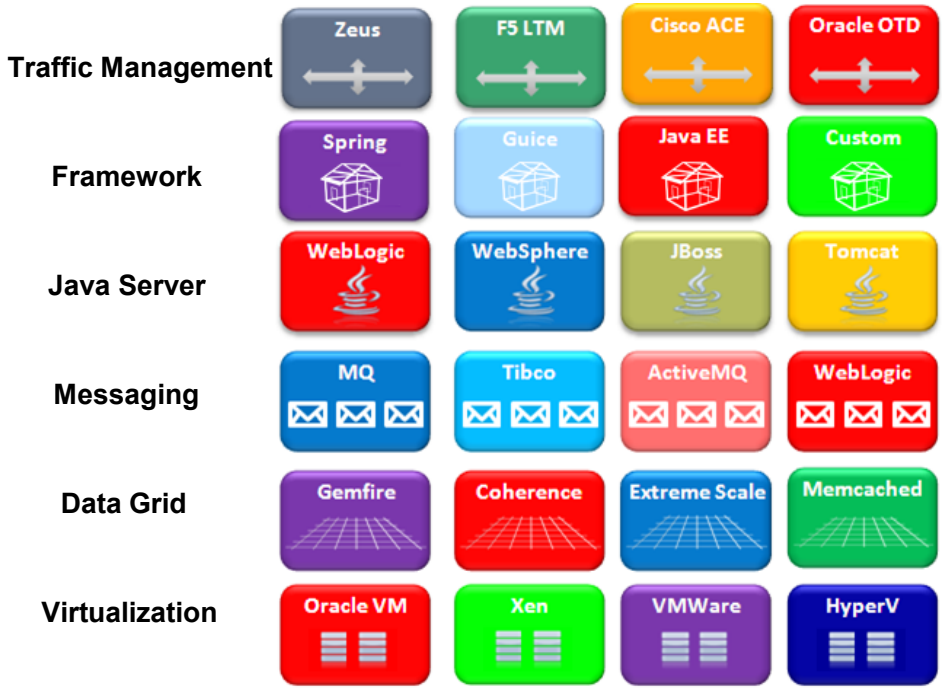

- Complex
- Unintegrated
- Multiple Vendors
- Multiple Consoles
- How Elastic?

- Hard to Maintain
- Non Standard
- Mission Critical?
- Provisioning?
- Consistent Security?

Cloud Application Foundation

Out of the Box, Best of Breed, Pre-Integrated Cloud Infrastructure

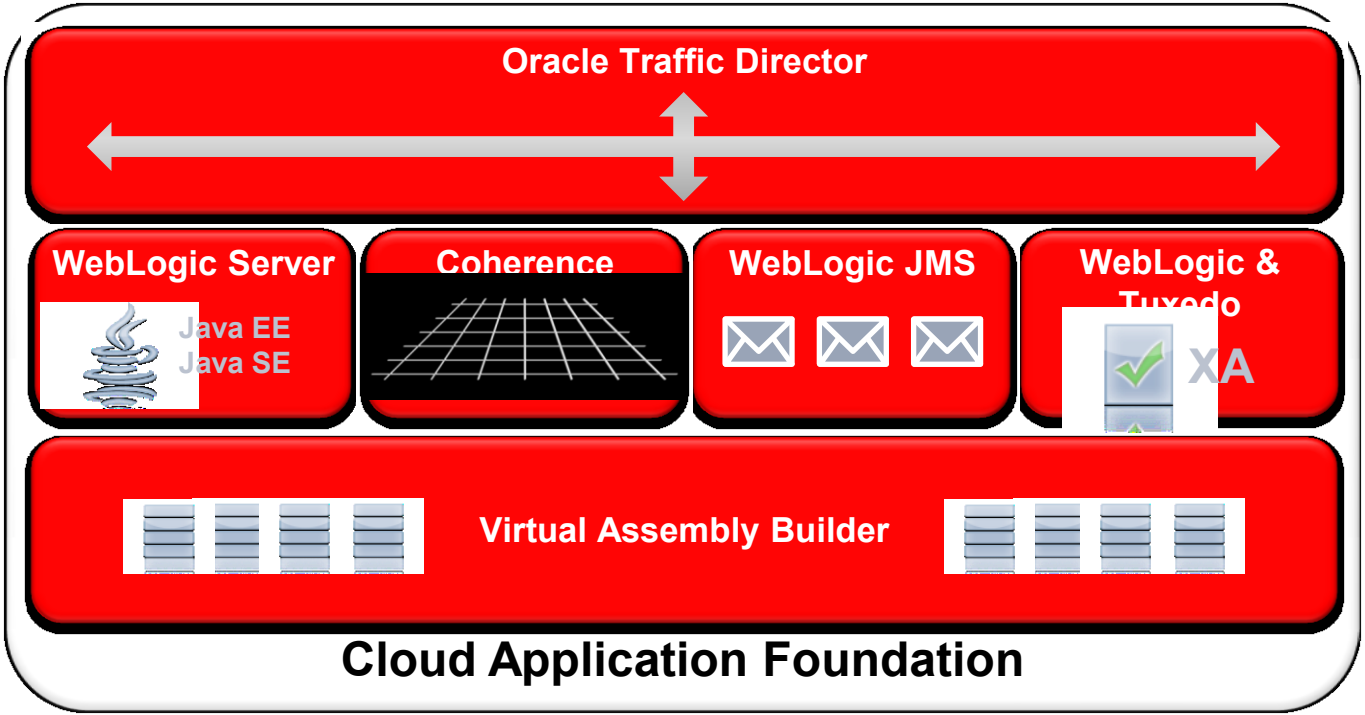

Oracle Cloud Application Foundation

Out of the Box, Best of Breed, Pre-Integrated Cloud Infrastructure

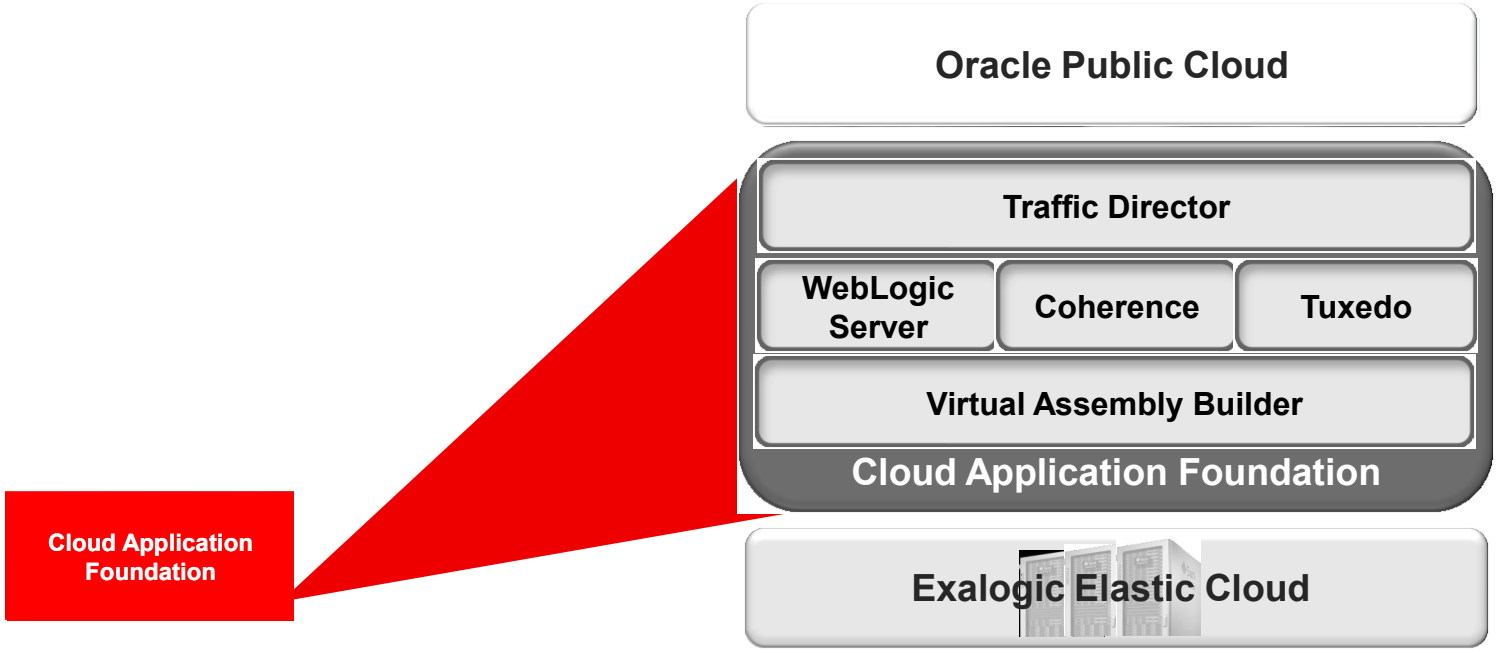

Cloud Application Foundation

For Conventional, Engineered and Cloud Environments

Industry's top choice for Conventional Systems

WebLogic Server and Coherence

Deploy/Run on Conventional Systems

Optimized for Engineered Systems

WebLogic Server and Coherence

Deploy/Run on High Performance Engineered Systems

Built for the Cloud

WebLogic Server and Coherence

Deploy/Run on Private and Public Clouds

ORACLE

Coherence

Problem Domain of Coherence

High load, stateful, middle tier services using shared services

Service Clients
Browsers, supplier and partner clients, application clients, mobile apps, and other service consumers.

- Grow and scale naturally

Service Implementation
Custom applications, BPM processes, service bus endpoints, reporting and batch services …

- **Facing rapidly increasing load**
- **Commonly stateful**
- **Highly dependent on shared services**

Shared Services
RDBMS, cloud services, supplier and partner services, mainframe applications, and so on.

- Complex and expensive to scale
- High latency SPOB and SPOF

Coherence In Memory Data Grid

Scaling, Performance and Reliability of Application Infrastructure

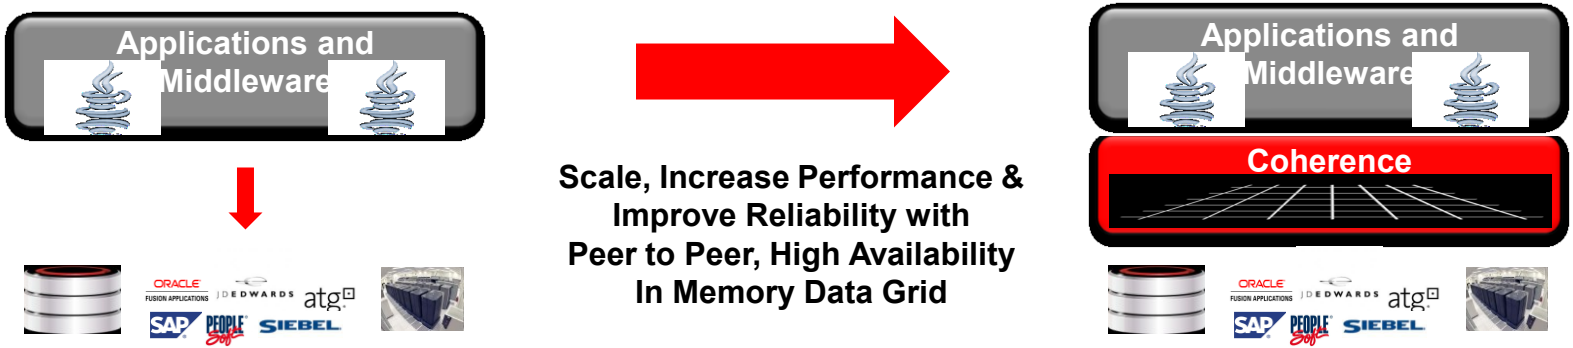

- **Application Scalability**
  - Reducing backend load with in memory caching – database, applications, mainframe
- **Application Performance**
  - In memory data access – information close in memory, not in slow backend systems
- **Application Reliability**
  - Reliable in memory data grid providing high-availability to applications

Oracle Coherence In-Memory Data Grid

Secret Sauce for Scaling, Performance and Reliability at Internet Scale

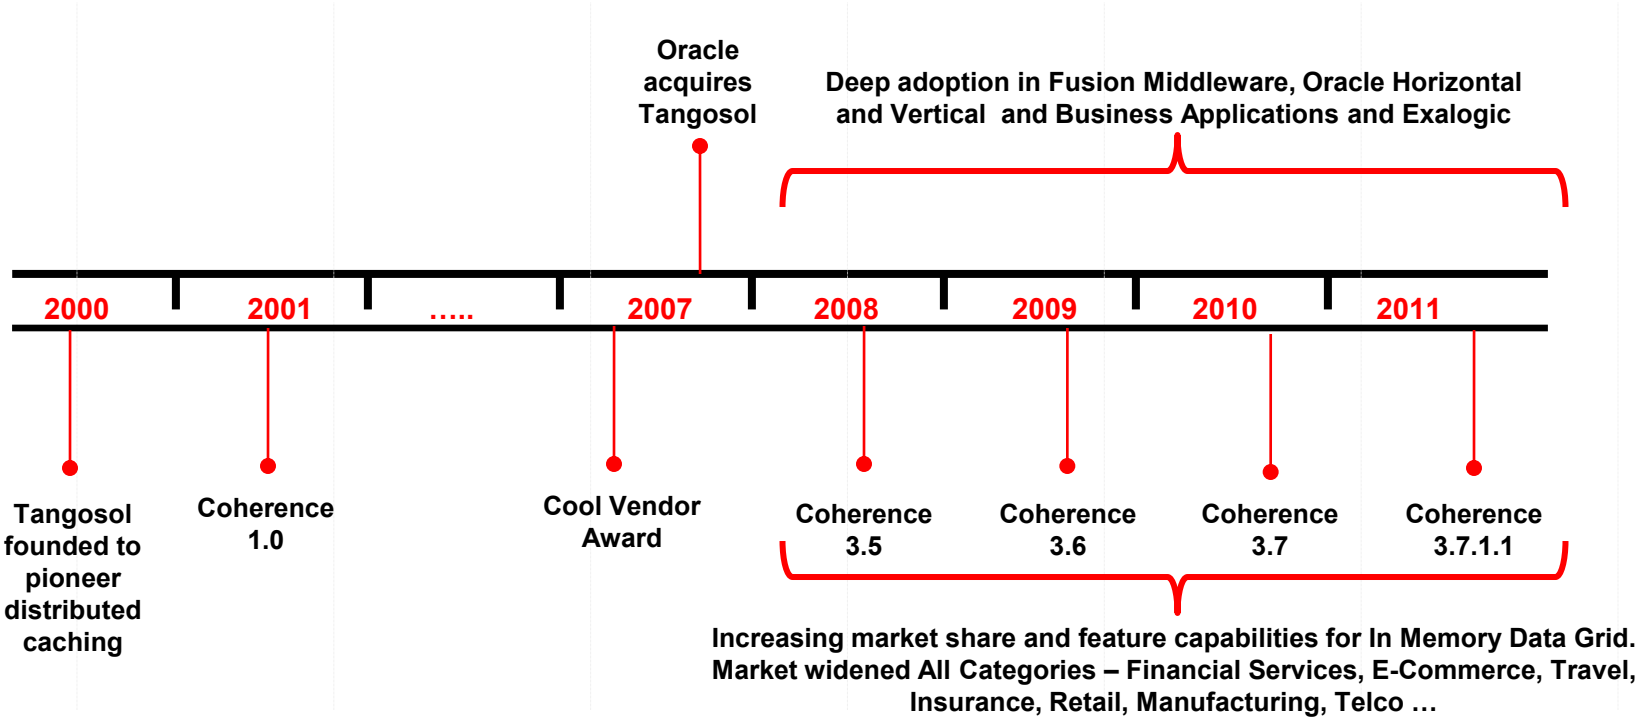

Coherence Clustered Caching Explained

Partitioned, Fault Tolerant, Self-Healing Cache

- Cluster of nodes holding % of primary data locally
- Back-up of primary data is distributed across all other nodes
- Logical view of all data from any node

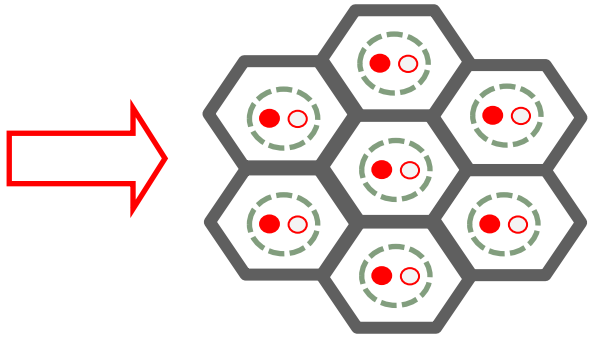

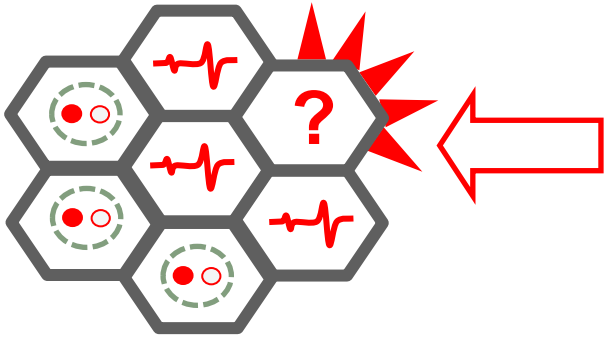

- All nodes verify health of each other
- In the event a node is unhealthy, other nodes diagnose state

- Unhealthy node isolated from cluster
- Remaining nodes redistribute primary and back-up responsibilities to healthy nodes

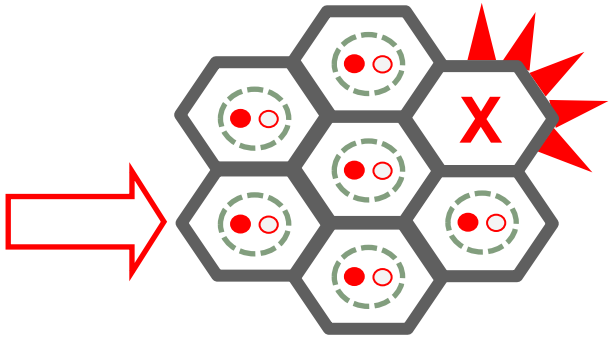

ORACLE

Example Uses Coherence In-Memory Data Grid

- **Caching**
  - Client applications request data objects and state from the Data Grid rather than backend data sources
- **Analytics**
  - Client applications ask the Data Grid questions from simple queries to scenario modeling
- **Compute**
  - Applications execute compute – business logic and transactions - across the Data Grid
- **Events**
  - Applications take action based on events from the Data Grid

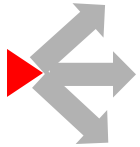

Data Management Options

- **Partitioned Caching**
  - Automatically partitions data set across members of grid.
- **Partitioned Fault Tolerance**
  - Automatically manages the fault tolerance of your data.
  - Backups guaranteed to be on separate machine from primary.
  - Backup for one node's data is shared across nodes.
- **Cache Client/Cache Server**
  - Member is responsible for a partition of the data or not ("storage enabled" vs. "storage disabled")
  - Cache Client – typically an application instance
  - Cache Server – typically for storage and data processing only.
- **Near caching**
  - Blend of replicated and partitioned topologies. Most recently/frequently used data is stored locally.

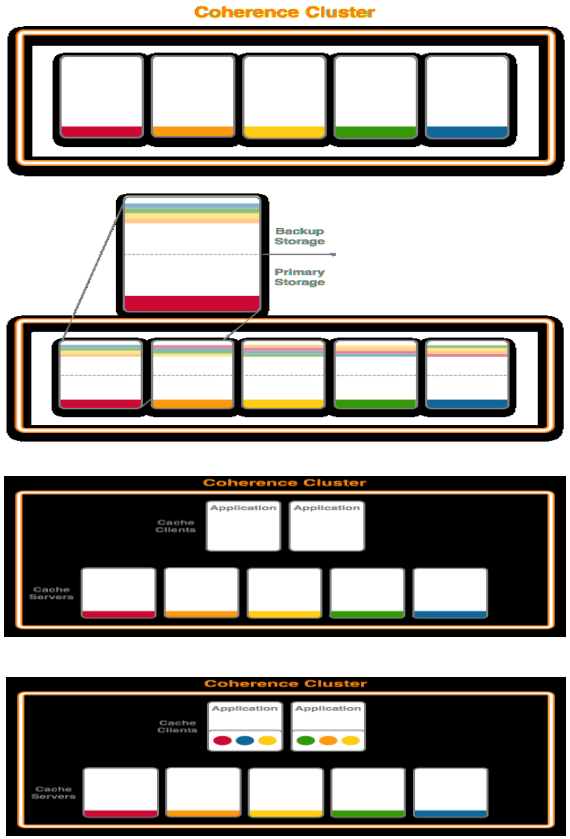

Coherence Caching Patterns

- **Cache Aside - Developer manages cache**
  - Check the cache before reading from data source
  - Put data into cache after reading from data source
  - Evict or update cache when updating data source
- **Read Through/Write Through**
  - All data reads/writes occur through cache
  - Cache miss causes load from data source automatically
  - Updates to cache written synchronously to the data source
- **Write Behind**
  - All data writes occur through cache
  - Updates to cache written asynchronously to the data source

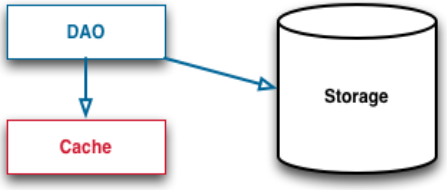

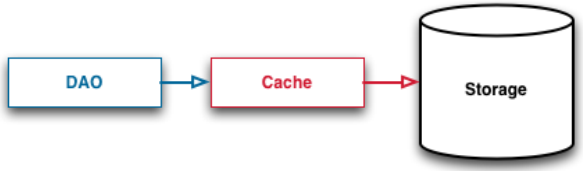

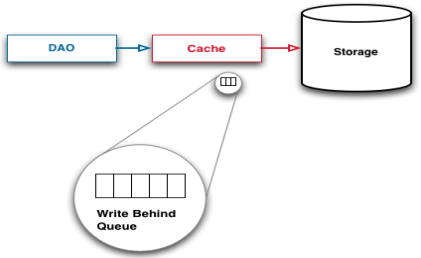

Querying The Data Grid

- **Query**
  - Programmatic query mechanism
  - Queries performed in parallel across the grid
  - Standard indexes provided out-of-the-box and supports implementing your own custom indexes
- **Continuous Query Cache**
  - Automatically, transparently and dynamically maintains a view locally based on a specific criteria

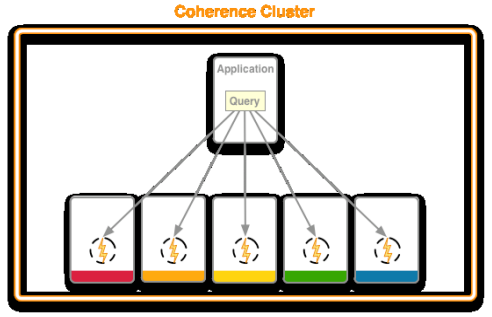

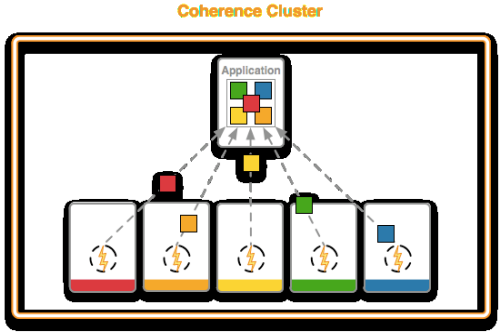

Eventing

- **Complex Event-Driven Apps**
  - Java Bean Model
  - Key-Based
  - Filter-Based
- **“Live Objects”**
  - Objects can respond to own state changes
  - State always recoverable
  - Build complex Staged Event Driven Architectures

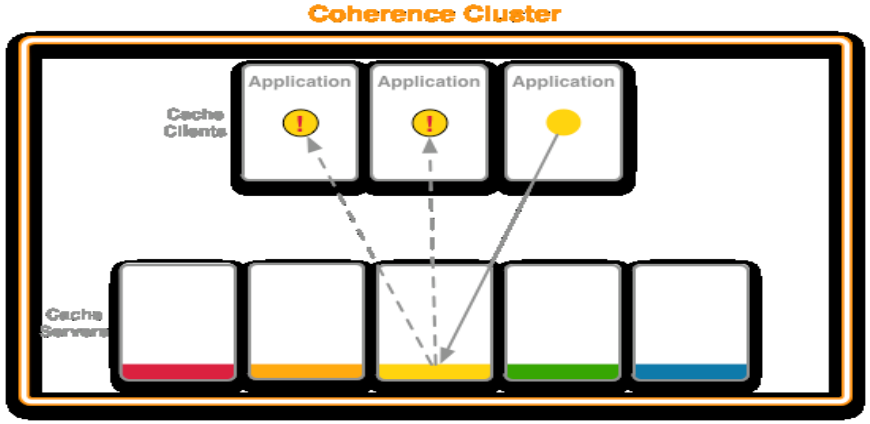

Data Processing and Map/Reduce Aggregation

- **Processing on the grid**
  - The **inverse of caching**
  - Sends the processing to where the data is in the grid
    - Once and only once guarantees
    - Processing is automatically fault-tolerant
  - Processing can be
    - Targeted to a specific key or collection of keys
    - Targeted to any object that matches a specific criteria
- **Parallel Aggregation**
  - Aggregates at each cluster node
  - Final “reduce” at client node

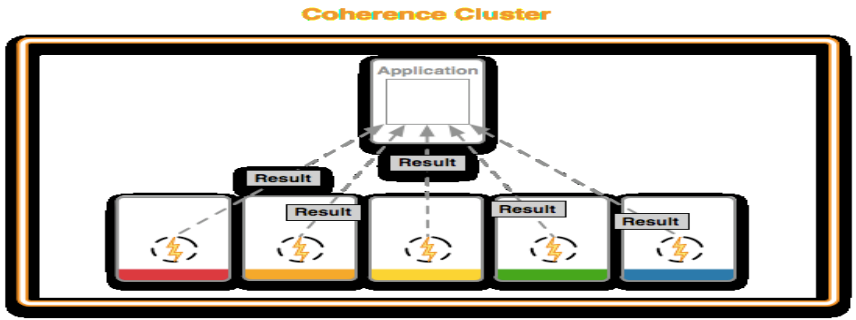

WebLogic with Coherence in WebLogic Suite

Scale Session, Cache and Data Grid as Independent Architectural Tier

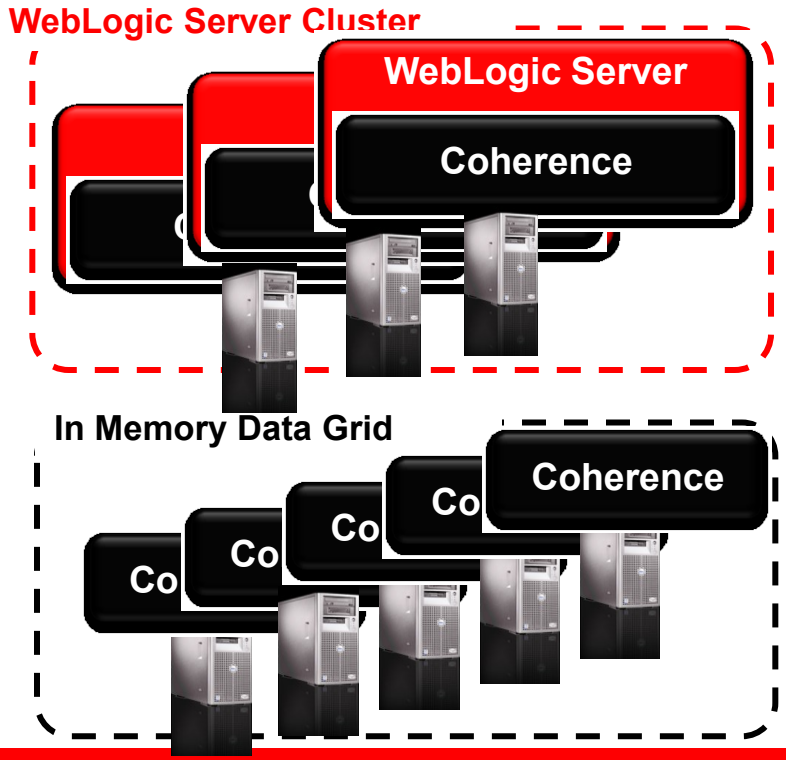

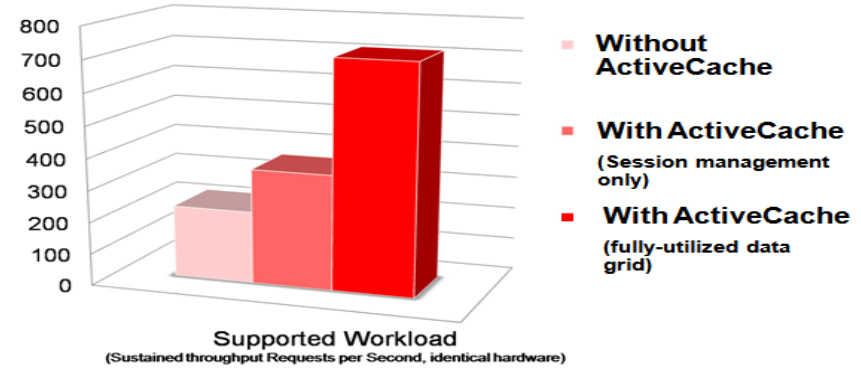

WebLogic Server with Coherence

WebLogic and Coherence Integration Summary

Built in Out of the Box

- Administration, operations and management built into WebLogic
- Declarative scale out session management
- Cache data access with synch/asynch read /write through
- Analytics, events and compute

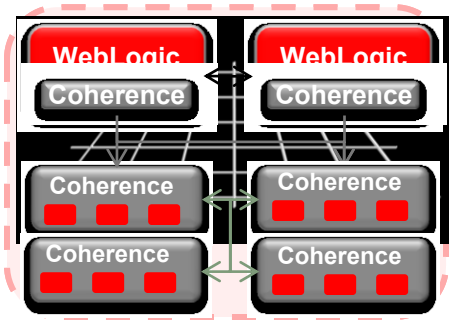

Declarative Session Management

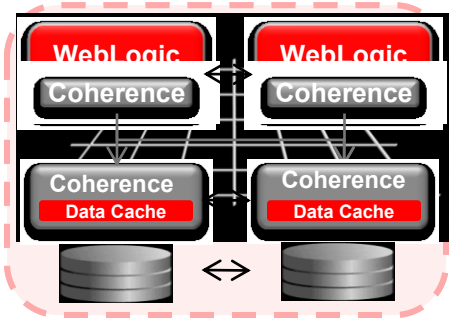

Persistence Caching with Read and Write Through

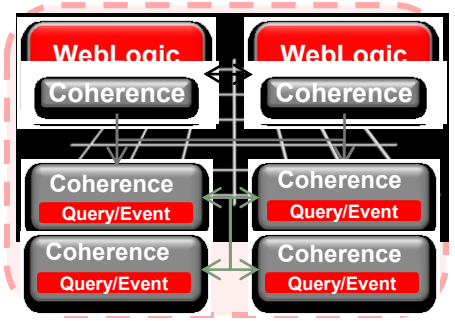

Query, compute and event

Real Customers, Real Business Value

Oracle Coherence Customer Case Studies

- Application supports 25,000 events each year, including webcasts, event portals, virtual events, and virtual environments
- **Eliminated service interruptions for events of 30,000 attendees**

- Provides timely risk and P&L data to traders, researchers, and financial controllers in a growing and volatile trading market
- **Replaced legacy spreadsheets to provide resilient, scalable risk platform**

- Real-time game show requires powerful, reliable, and scalable back-end technology for best customer experience
- **Provided sub-second response time to quiz show users with 100% uptime.**

- Gains Edge in Intraday Credit Risk Management with In-memory, Distributed Data Grid Solution
- **Achieve 5 millisecond response times for pre-transaction credit checks against counterparties globally and scale by up to a factor of ten without risk of increase in latency**

ORACLE

Packaging and What's New Summary

Coherence Editions

Coherence Standard Edition	Coherence Enterprise Edition	Coherence Grid Edition*
• Fault-tolerant data caching • Unlimited Data Clients	• Fault-tolerant data caching • Data management, including write-behind, transactions, analytics and events • Unlimited Data Clients	• Fault-tolerant data caching • Data management including write-behind, transactions, analytics and events • Real-time clients • WAN Support • Elastic Data

*** Customers who have purchased WL Suite and want to upgrade to from Coherence EE to GE can now do so via a WebLogic Suite option.**

WebLogic Product Packaging

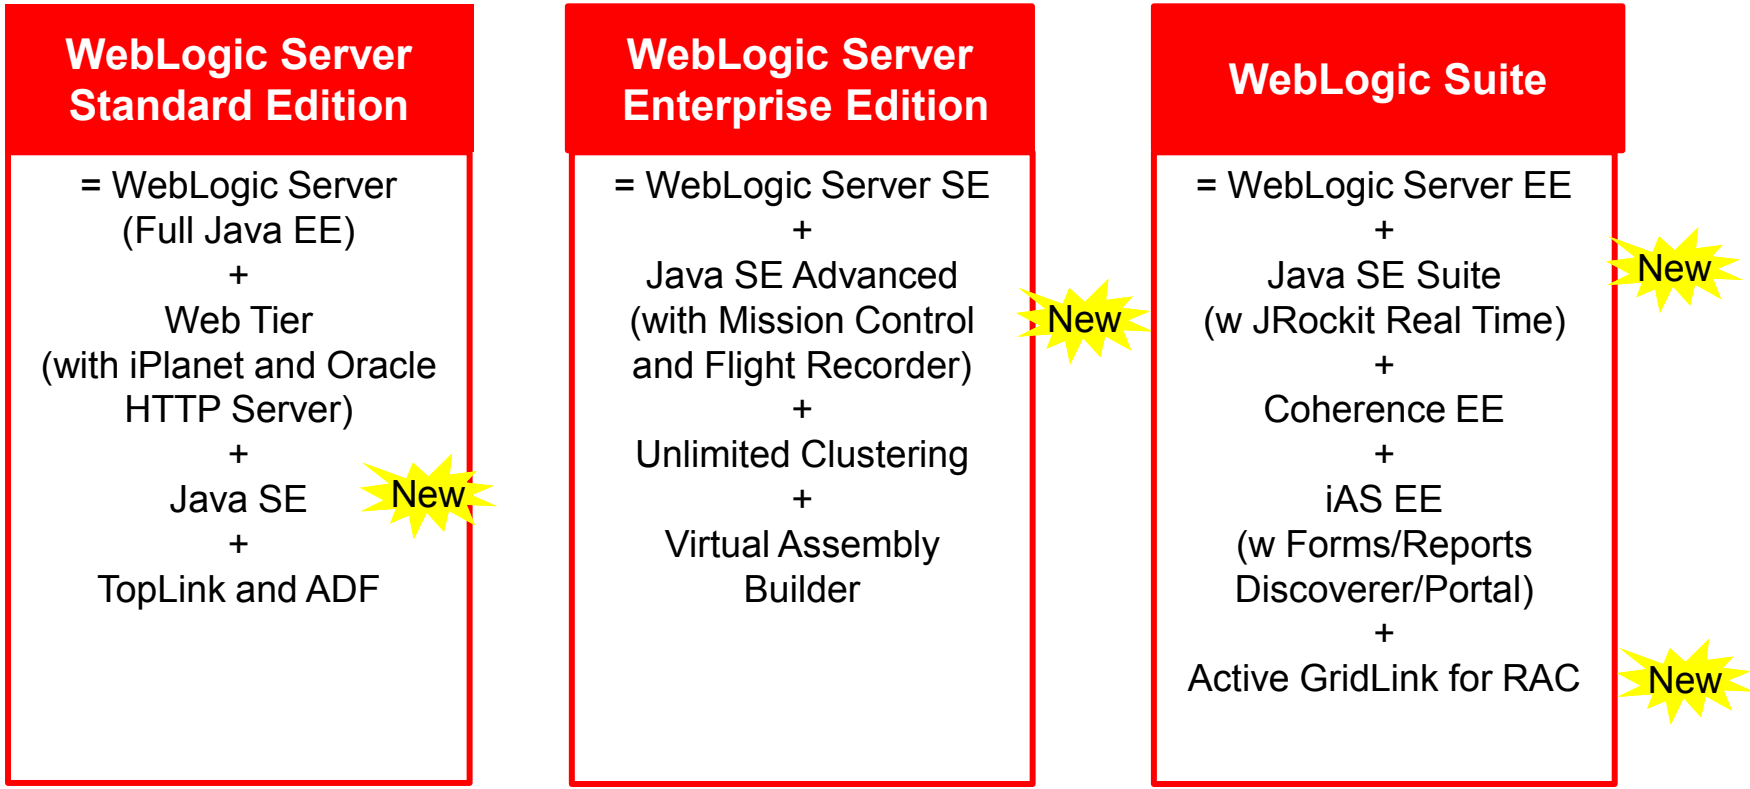

Over 40 Complete/Planned Integrations

Generally Available

- Oracle IGBU
- Oracle CGBU
- Oracle WebLogic Server
- Oracle Access Manager
- Oracle CEP
- Real-Time Decision
- TopLink
- WebLogic Portal
- Service Delivery Platform
- Oracle GlassFish Server
- Oracle PeopleSoft
- Oracle Service Bus
- SOA Infrastructure
- Oracle Data Integration

Near Term

- Oracle Golden Gate
- ATG
- SOA BPEL
- Oracle ADF
- Secure Token Service
- Oracle Identity Federation
- Meta-data Services (MDS)

Advanced Capabilities
Data Grid
Data Cache
Simple Clustering

Medium Term

- WebCenter
- SOA Human Workflow
- SOA Business Rules
- SOA Coherence Adapter
- SOA Technology Adapters
- BI Publisher
- Oracle Entitlement Services
- Oracle BEAM
- Oracle Web Services Manager
- ATG Hosted Services

Exalogic

Oracle Exalogic Elastic Cloud

Cloud Application Foundation

Exalogic Configurations | Scalability

Eighth Rack	Quarter Rack	Half Rack	Full Rack	Multi-Rack
4 Nodes	**8** Nodes	**16** Nodes	**30** Nodes	**240** Nodes
384 GB RAM	**768** GB RAM	**1.5** TB RAM	**2.8** GB RAM	**23** TB RAM
400 GB SSD	**1.6** GB SSD	**1.6** TB SSD	**3** TB SSD	**24** TB SSD
60 TB ZFSSA	**60** TB ZFSSA	**60** TB ZFSSA	**60** TB ZFSSA	**480** TB ZFSSA

Exalogic Elastic Cloud Software 2.0

Improved Consolidation, Management, Performance

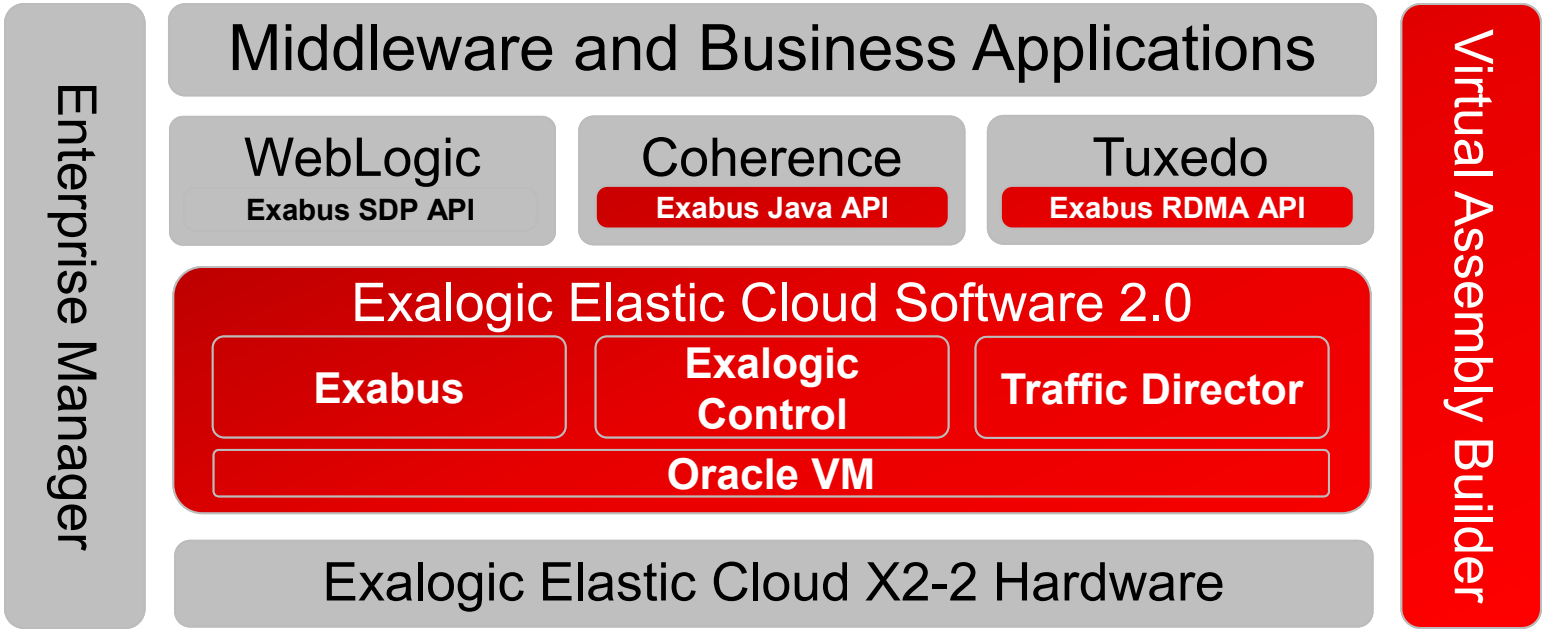

Exabus - High Speed Network Virtualization

Direct Memory I/O for Java and C++

Standard Hardware I/O

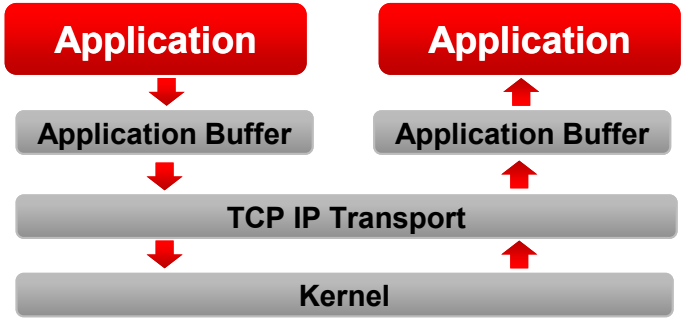

20% Buffer Copies
40% Transport Processing
40 % Kernel Context Switches

Exabus

Zero Buffer Copy
Direct Memory Access
Kernel Bypass

4X Throughput, 6X Lower Latency

Coherence Optimizations for Exabus

4x bandwidth, 1/6 latency

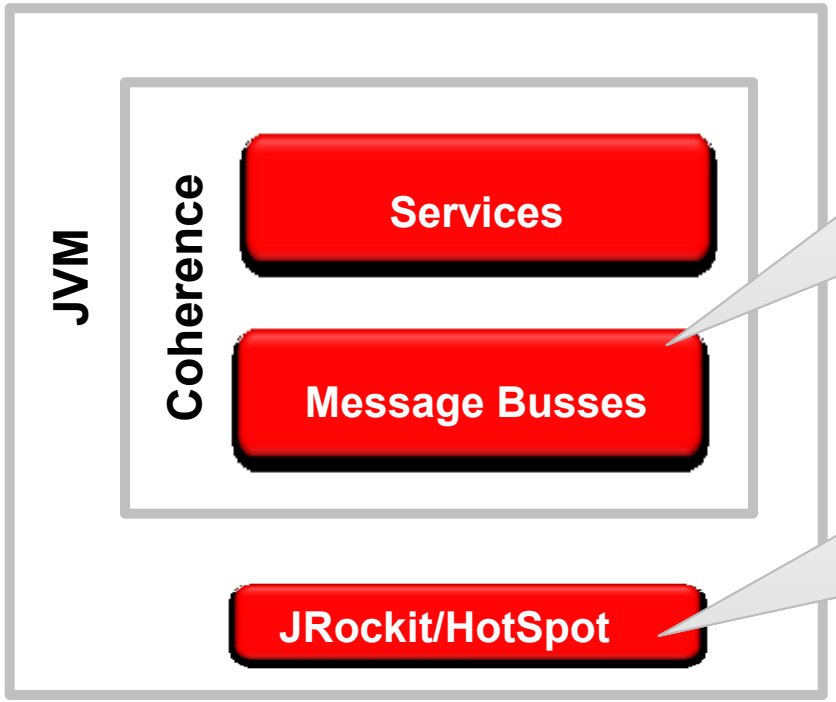

Provider-based Transport Layer

- One per service for scalability to multiple cores
- Increased message parallelism to max network
- Context Switches reduced from 3 to 1
- Datagram still used for cluster protocol

Infiniband Message Bus Provider

- RDMA offloads host processor
- Zero copy and kernel bypass
- Predictive notifications avoid costly interrupts
- Custom off-heap DirectByteBuffer reduces GC

Exalogic: Simplified Coherence Deployments

Fewer, Larger JVMs in EECS 2.0

Number of JVMs Needed for Coherence

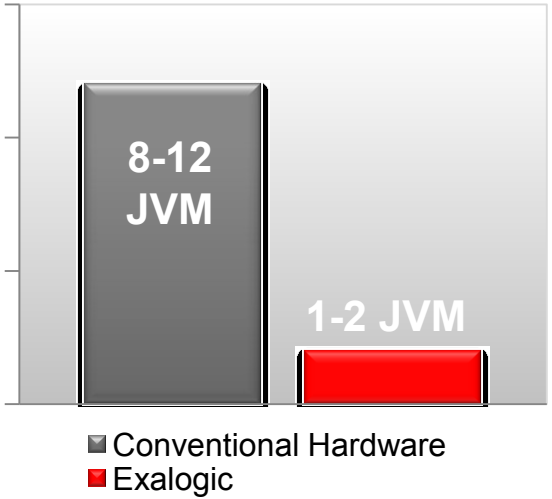

Conventional Infrastructure

Datagram Data Communications

Exalogic Infrastructure

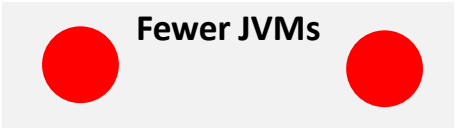

Message Bus Data Communications

- Scalability data communications to maximize network bandwidth with fewer JVMs.
- Memory Buffers allocated off-heap dramatically reduces object allocation and thus garbage collection pauses
- Scalable communications and reduced GC pressure favors fewer, larger JVMs for simplified deployment

Exalogic: Coherence Cluster Rebalancing

Exabus increases availability

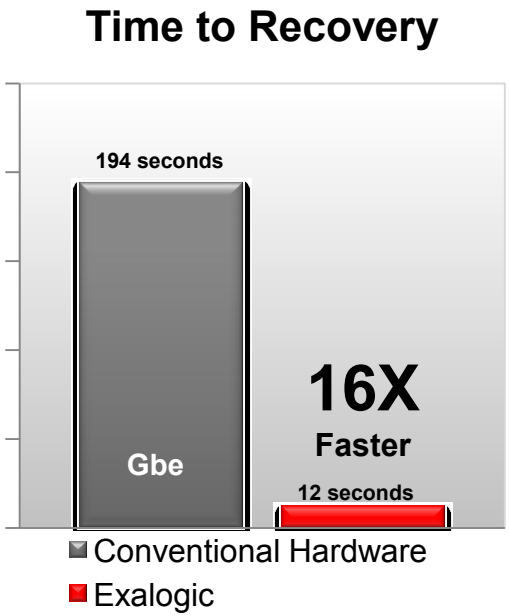

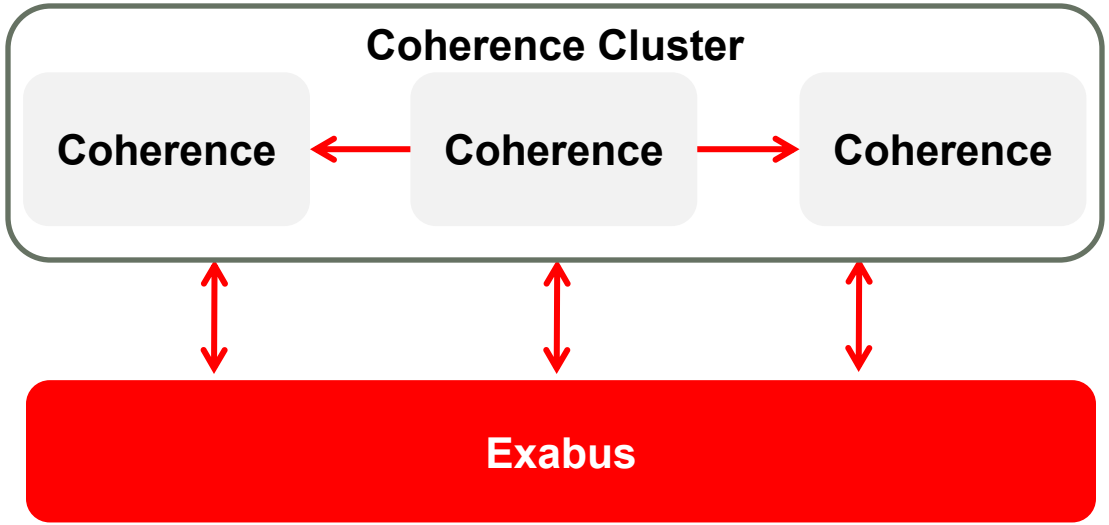

More Information

Join the Coherence Community

http://coherence.oracle.com


@OracleCoherence


facebook.com/OracleCoherence


blogs.oracle.com/OracleCoherence


Oracle Coherence Users


youtube.com/OracleCoherence


coherence.oracle.com/display/CSIG

Coherence Special Interest Group


ORACLE®