

ORACLE[®]



MySQL Web Reference Architectures

Building Massively Scalable Web Infrastructure

Mario Beck (mario.beck@oracle.com)
Principal Sales Consultant MySQL

Session Agenda

- Requirements for Innovating on the Web
- Reference Architectures
 - Sizing & Topologies
 - Enabling Technology
 - Best Practices
- Server Sizing
- Value Added Components & Services
- Resources to Get Started



MySQL:

The World's Leading On-Line Database



Powering 9 of the top 10 most trafficked sites on the web*

* <http://www.alex.com/topsites>

Innovating on the Web

Translating Business Requirements to Technology Implementation

- Top Challenges
 - Fast Time to Market
 - High availability
 - High performance
 - Cost-effectively scale to meet rapidly growing demands
 - Open, customizable and repeatable
 - Protection of customer data
 - Low total cost of ownership
- Leverage “Best Practices”
 - Accelerating Time to Market with greater agility
 - Reducing cost, risk and complexity
 - Eliminate expensive “trial & error”

MySQL Web Reference Architectures

- Repeatable best practices developed with leading web properties
 - Recommended architectures & topologies
- 4 x Reference Architectures
 - Small
 - Medium
 - Large
 - Extra Large (Social Networking)
- 4 x common platform components
 - User Authentication & Session Management
 - Content Management
 - eCommerce
 - Analytics

Reference Architecture Sizing

Social Network

	Small	Medium	Large	Extra Large
Queries/Second	<500	<5,000	10,000+	25,000+
Transactions/Second	<100	<1,000	10,000+	25,000+
Concurrent Read Users	<100	<5,000	10,000+	25,000+
Concurrent Write Users	<10	<100	1,000+	2,500+

Database Size

Sessions	<2 GB	<10 GB	20+ GB	40+ GB
eCommerce	<2 GB	<10 GB	20+ GB	40+ GB
Analytics	<10 GB	<500 GB	1+ TB	2+ TB
Content Management	<10 GB	<500 GB	1+ TB	2+ TB



Reference Architectures

Small: Web Reference Architecture

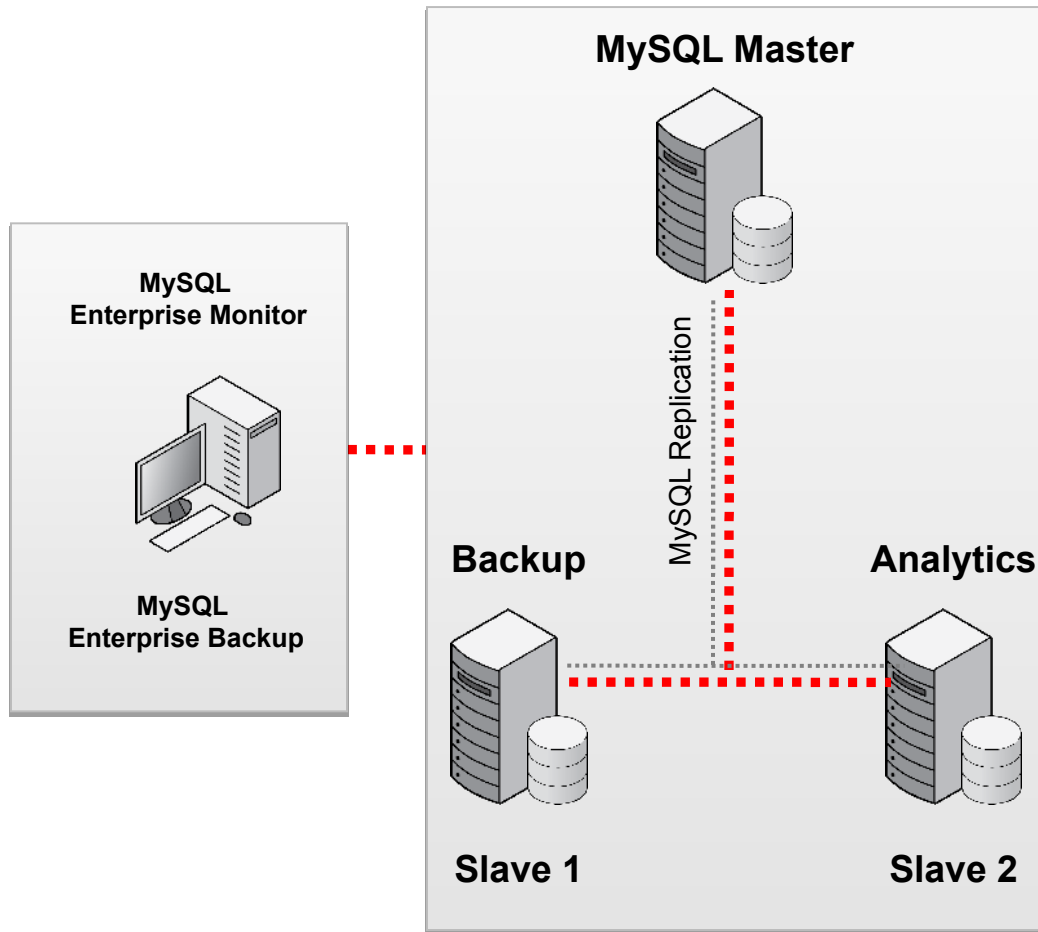
Social Network

	Small	Medium	Large	Extra Large
Queries/Second	<500	<5,000	10,000+	25,000+
Transactions/Second	<100	<1,000	10,000+	25,000+
Concurrent Read Users	<100	<5,000	10,000+	25,000+
Concurrent Write Users	<10	<100	1,000+	2,500+

Database Size

Sessions	<2 GB	<10 GB	20+ GB	40+ GB
eCommerce	<2 GB	<10 GB	20+ GB	40+ GB
Analytics	<10 GB	<500 GB	1+ TB	2+ TB
Content Management	<10 GB	<500 GB	1+ TB	2+ TB

Small: Web Reference Architecture



- Single server supporting all workloads

Members/Authentication
eCommerce
Content Management
Search

- Data replicated to slaves for back-up & analysis
- 1st scaling step: move session management to dedicated MySQL Server
- Complex to tune

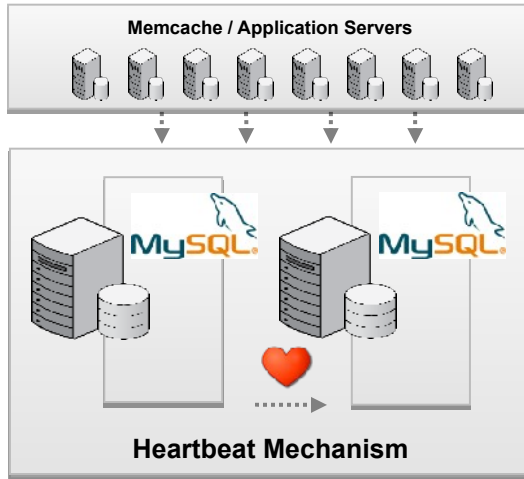
Only deploy when future traffic growth is very limited

Medium: Web Reference Architecture


	Small	Medium	Large	Extra Large
				Social Network
Queries/Second	<500	<5,000	10,000+	25,000+
Transactions/Second	<100	<1,000	10,000+	25,000+
Concurrent Read Users	<100	<5,000	10,000+	25,000+
Concurrent Write Users	<10	<100	1,000+	2,500+
Database Size				
Sessions	<2 GB	<10 GB	20+ GB	40+ GB
eCommerce	<2 GB	<10 GB	20+ GB	40+ GB
Analytics	<10 GB	<500 GB	1+ TB	2+ TB
Content Management	<10 GB	<500 GB	1+ TB	2+ TB

Medium: Web Reference Architecture

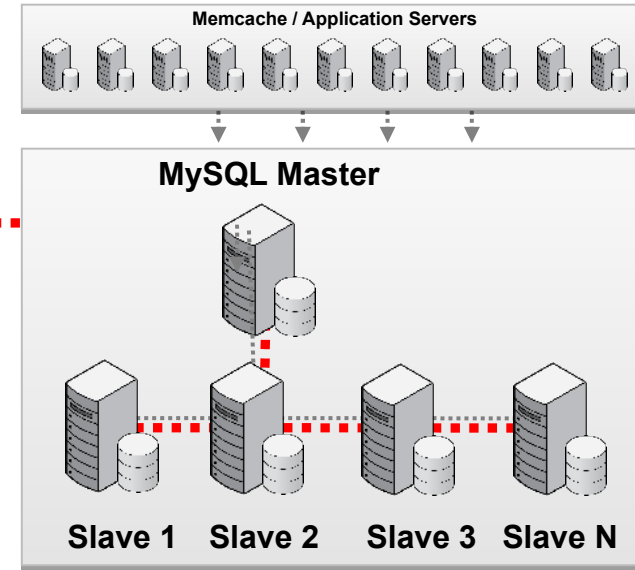
Session Management



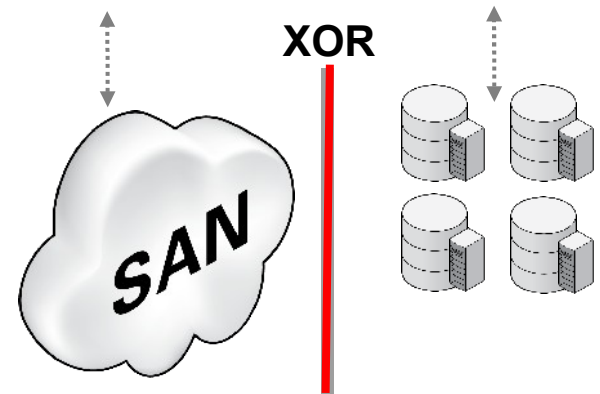

eCommerce



Content Management



Analytics



Best Practices (1)

Medium Web Reference Architecture

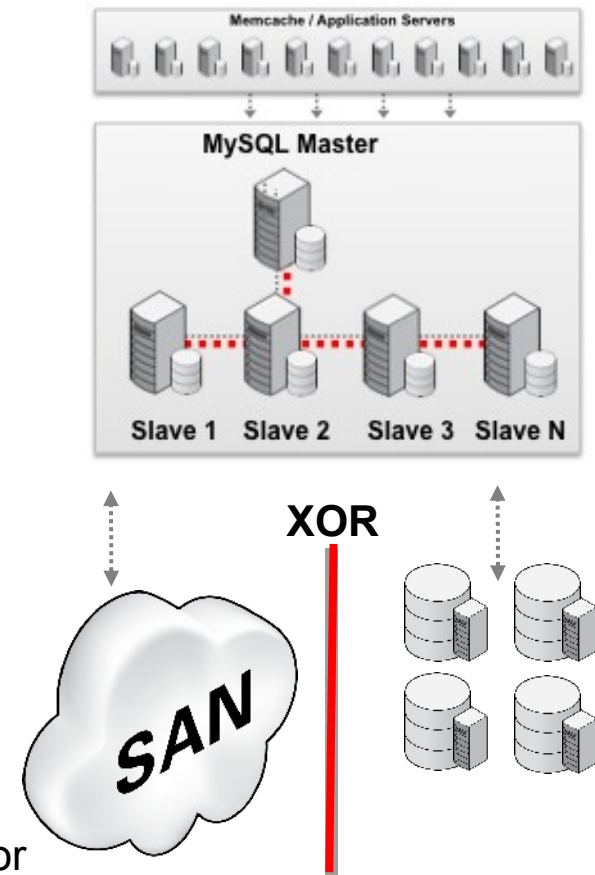
- Each component deployed onto dedicated server & storage infrastructure
 - Deployed, managed and scaled independently
- Server ratio: 8 application servers to each MySQL Server
 - More for PHP applications, less for Java
 - Add more slaves as the application tier scales
- Memcached deployed in session & content management components
 - Distributed memory caching layer
 - Reads fulfilled from cache, relieving load on the source database servers

Best Practices (2)

Medium Web Reference Architecture

- Content Management

- Each slave can handle around 3,000 concurrent users
- Each master can handle up to 30 slaves
- MySQL Replication for high availability
 - Can include Heartbeat, depending on application failover requirements
- Meta-data & indexing of content assets managed by MySQL
- File System & physical storage manage content assets
 - Distributed file system (i.e. MogileFS)
 - High quality SAN (redundancy for HA)
 - Distributed across local storage with OS level clustering for HA



Best Practices (3)

Medium Web Reference Architecture

- Session Management & eCommerce
 - Both deployed onto InnoDB storage engine
 - Session data maintained for up to 1 hour in a dedicated partition, rolling partitions used to delete aged data
 - More users persist session data to provide greater personalization for repeat visitors
 - Data is captured in Analytics Database
 - MySQL Replication with Heartbeat for HA
 - Configure semi-synchronous replication or OS-level DRBD for eCommerce
 - If web traffic grows, move Session Management to MySQL Cluster
 - HA and in-memory data management can reduce need for external HA mechanisms & memcached

Large: Web Reference Architecture

Social Network


	Small	Medium	Large	Extra Large
Queries/Second	<500	<5,000	10,000+	25,000+
Transactions/Second	<100	<1,000	10,000+	25,000+
Concurrent Read Users	<100	<5,000	10,000+	25,000+
Concurrent Write Users	<10	<100	1,000+	2,500+

Database Size

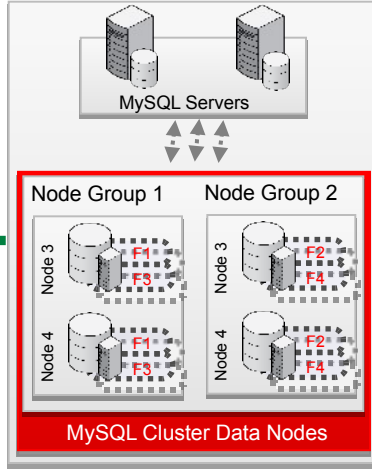
Sessions	<2 GB	<10 GB	20+ GB	40+ GB
eCommerce	<2 GB	<10 GB	20+ GB	40+ GB
Analytics	<10 GB	<500 GB	1+ TB	2+ TB
Content Management	<10 GB	<500 GB	1+ TB	2+ TB

Large: Web Reference Architecture

Session Management




eCommerce



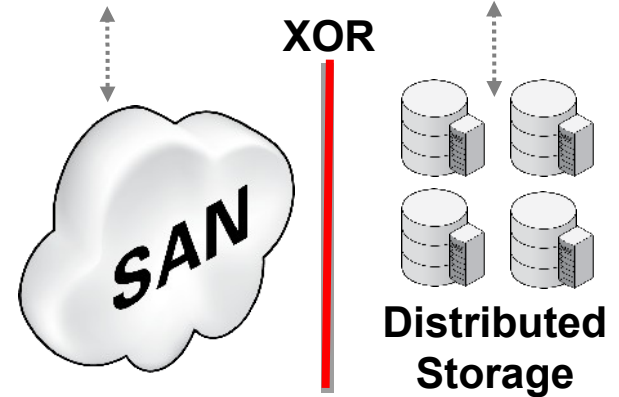
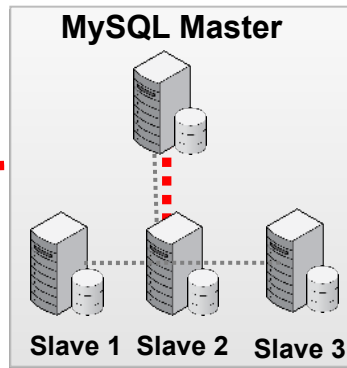
Data Refinery



Content Management

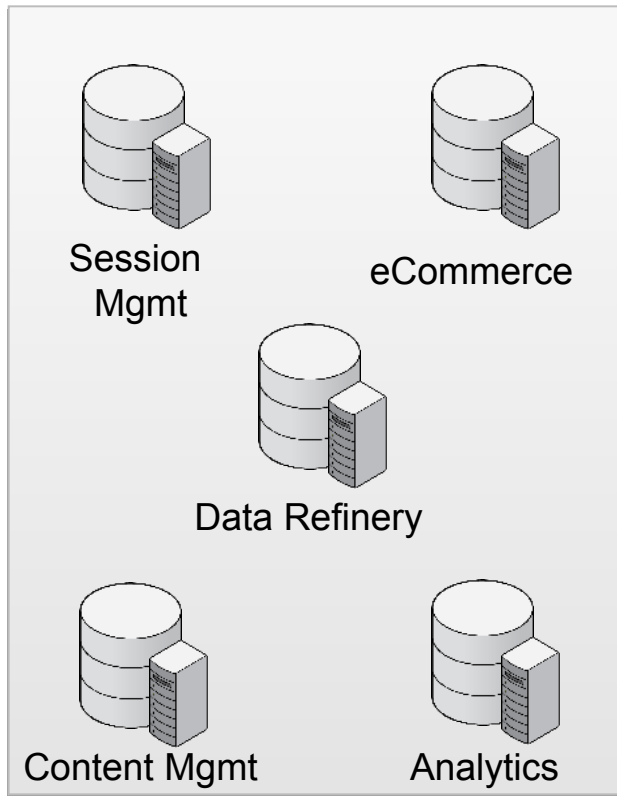


Analytics

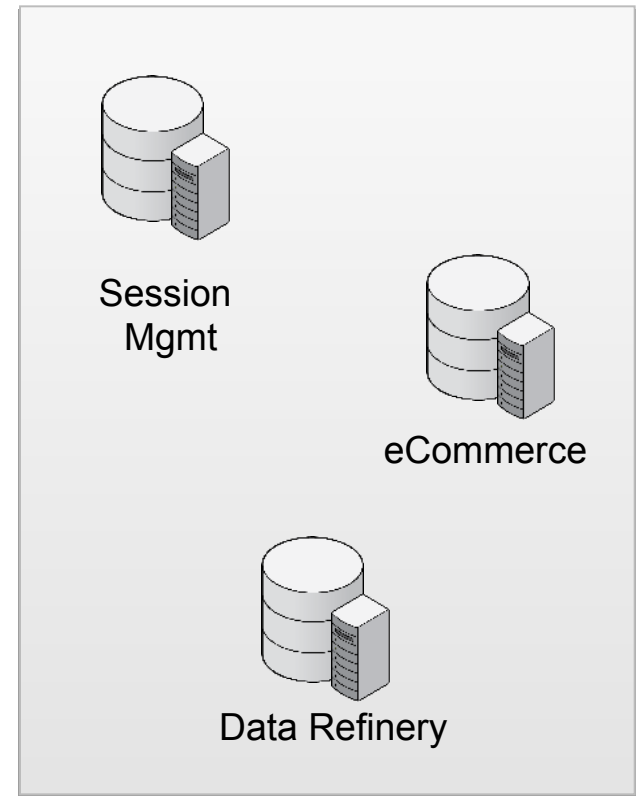



Large: Web Reference Architecture

Conceptual View



Wien Data Center



Sidney Data Center

Best Practices

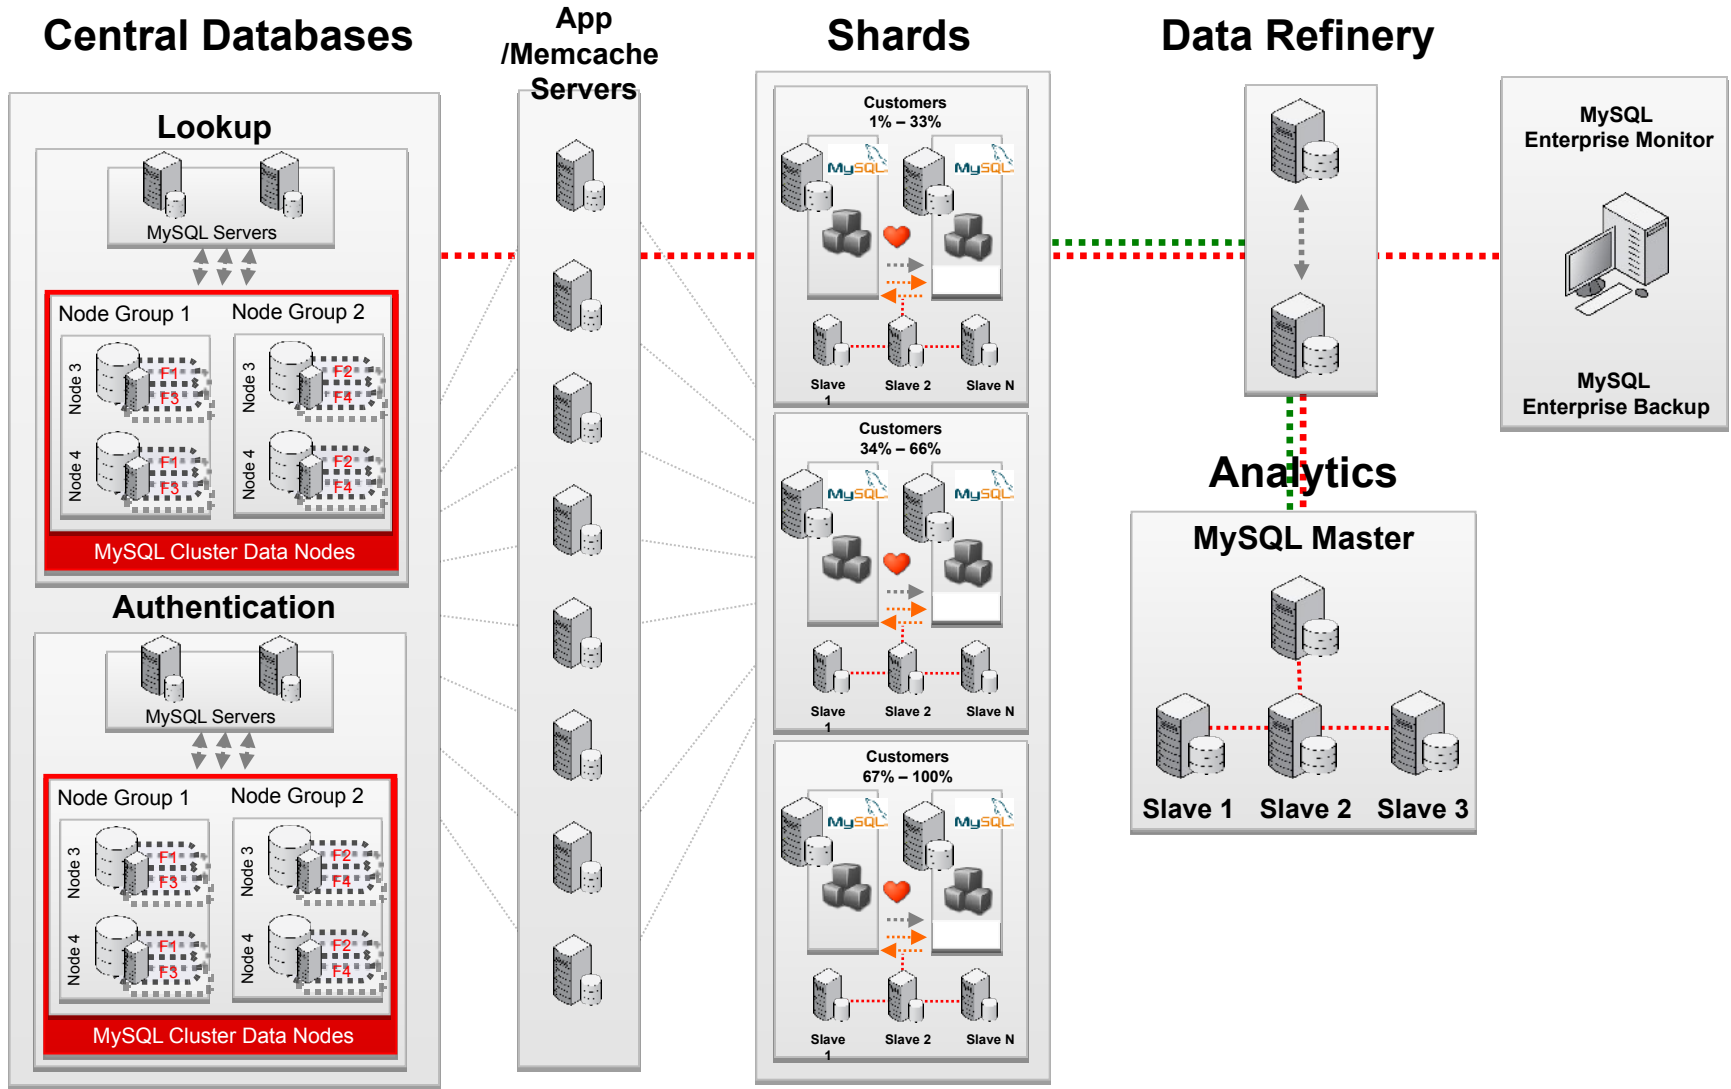
Large Web Reference Architecture

- Builds on best practices of Medium Web Ref Arch
 - Dedicated infrastructure for each workload, MySQL Replication, Memcached, etc.
- Introduces Data Refinery
 - Aggregate data across the web components
 - Data cleansing
 - Supports higher volume content management and analytics
 - Builds Data Warehouse Dimensions
- Introduces MySQL Cluster
 - Session Management and eCommerce
 - 4 x Data Nodes support 6k page hits per second
 - Each page hit generating 8 – 12 database operations

Extra Large: Social Network Reference Architecture

	Social Network			
	Small	Medium	Large	Extra Large
Queries/Second	<500	<5,000	10,000+	25,000+
Transactions/Second	<100	<1,000	10,000+	25,000+
Concurrent Read Users	<100	<5,000	10,000+	25,000+
Concurrent Write Users	<10	<100	1,000+	2,500+
Database Size				
Sessions	<2 GB	<10 GB	20+ GB	40+ GB
eCommerce	<2 GB	<10 GB	20+ GB	40+ GB
Analytics	<10 GB	<500 GB	1+ TB	2+ TB
Content Management	<10 GB	<500 GB	1+ TB	2+ TB

Extra Large: Social Network



Best Practices


Social Networking Reference Architecture

- Introduces Sharding
 - Implemented at the application layer for scaling very high volume of writes
 - Data divided into smaller sets, distributed across low-cost hardware
 - Shards based on Hash of a single column – ie. User ID
- Sharding is complex
 - Recommend the Architecture and Design Consulting Engagement
- Sharding only used in a small percentage of workloads
 - Most Web 2.0 workloads are still read-intensive, ie record is read before updates applied



Server Sizing

The Perfect MySQL Server



- MySQL 5.1: 4-8 Cores, MySQL 5.5: 16 – 24 cores for 5.5
- x86_64 - 64 bit for more memory is important
 - Data/Memory ration 1/3 to 1/10 to good (rule of thumb)
 - The more the better
- Linux or Solaris best, Windows and Unix also fine.
- RAID 10 for most, RAID 5 OK if very read intensive
- Hardware RAID battery backed up cache critical!
 - More disks are always better!
 - 4+ recommended, 8-16 can increase IO performance if needed
- At least 2 x NICs for redundancy
- Slaves should be as powerful as the Master
- Oracle Sun X4170 for example

MySQL Cluster Hardware Selection - RAM & CPU

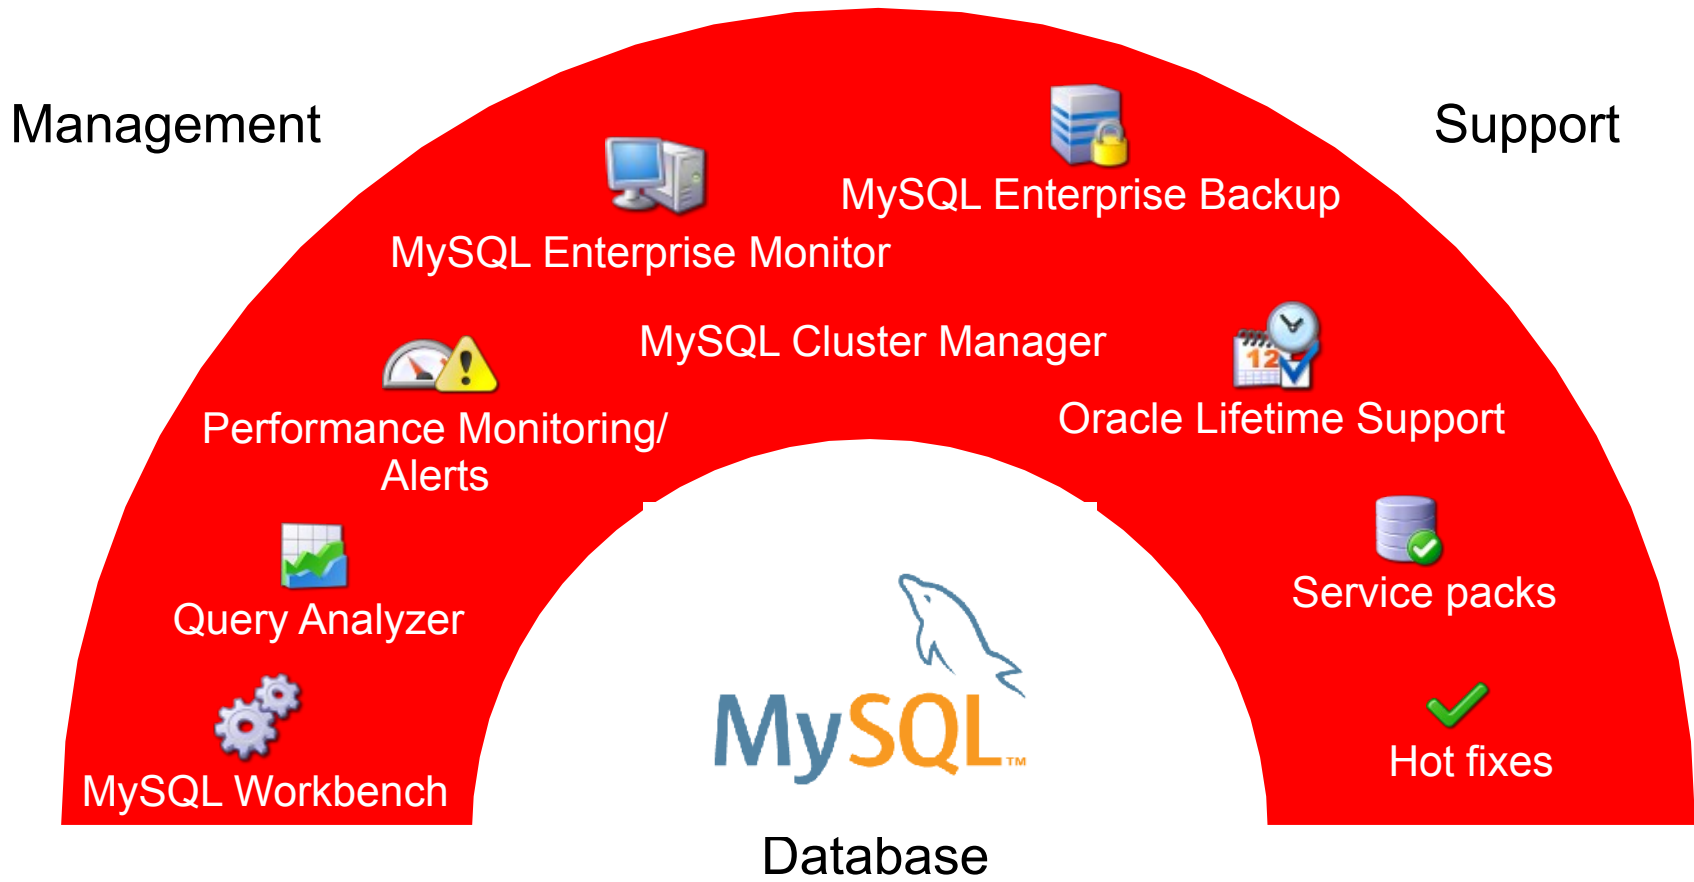
- Storage Layer (Data nodes)
 - One data node can (7.0+) use 8 cores
 - **CPU**: 2 x 4 core (Nehalem works really well). Faster CPU → faster processing of messages.
 - **RAM**: As much as you need
 - $\text{RAM per Server} = \text{Data Size} * \text{Replicas} * 1.25 / \# \text{ Data Nodes}$
 - (data redundancy + indexes drive overall memory requirement).
 - Example: 10GB database * 2 replicas * 1.25 / 2 Data Nodes = 12.5GB of RAM per data node.
- SQL Layer (MySQL Servers)
 - **CPU**: 2 – 16 cores
 - **RAM**: Not as important – 4GB enough (depends on connections and buffers)



Value-Added Components & Services

MySQL Commercial Editions

- Comprehensive offering of MySQL Database, Management tools, and Oracle Lifetime Support services



MySQL Services from Oracle

- MySQL Support
 - Global, 24 x 7 support coverage
- MySQL Consulting
 - Architecture and Design
 - Performance Tuning
 - High Availability
 - Migration
 - Remote DBAs
- MySQL Training
 - DBAs & Developers of all levels
 - Database and applications
 - Developing Dynamic Web Applications




Summary

- Reference Architectures designed as a springboard to innovating on the web
- Based on insight from most successful web properties
- Best practices & repeatable technologies for Scale & HA
- Next Step: Engage with MySQL Consulting from Oracle

Resources

- **MySQL Web Reference Architectures Whitepaper**
http://www.mysql.com/why-mysql/white-papers/mysql_wp_high-availability_webrefarchs.php
- **Designing and Implementing Scalable Applications with MySQL & Memcached**
http://www.mysql.com/why-mysql/white-papers/mysql_wp_cluster_ScalingWebServices.php
- **MySQL Enterprise Whitepaper**
http://www.mysql.com/why-mysql/white-papers/mysql_wp_enterprise_ready.php
- **MySQL Cluster for Web & eCommerce Applications, Whitepaper**
http://www.mysql.com/why-mysql/white-papers/mysql_wp_Cluster_For_OnlineApps.php



ORACLE[®]



MySQL Web Reference Architectures

Building Massively Scalable Web Infrastructure