




Lowering your IT Costs

Oracle's Database Strategy 2000-2009 and Beyond

ORACLE



Design Goals 2000 – 2009 and Beyond

- Take advantage of HW price/performance curve
- Improve utilization rates for servers/processors, storage, memory
- Increasingly automate repeatable, labor intensive tasks
- Reduce or eliminate the risk of change or user error
- Simplify the Information Architecture




Oracle Database 11g Release 2

Specific Areas of Cost Reduction


- Reduce hardware capital costs by factor of 5x
- Improve performance by at least 10x
- Reduce storage costs by factor of 12x
- Eliminate downtime AND unused redundancy
- Considerably simplify your software portfolio
- Raise DBA productivity by at least 2x
- Reduce upgrade costs by a factor of 4x

Is This Your Data Center?



- Lots of servers
- Separate storage silos
- Multiple vendors products (and versions)
- Different operating systems (and versions)
- Poor CPU utilization Rates(higher software costs)
- High operational expenses(more numbers of things and kinds of things to manage)


How Are Real Dollars Saved?



- Consolidation
 - Fewer **Number** of Things
 - Improved Utilization
- Rationalization
 - Fewer **Kinds** of Things
 - Simplified Utilization
- High-Level of Management Process Maturity
 - Documented
 - Repeatable
 - Automated

Real Application Clusters

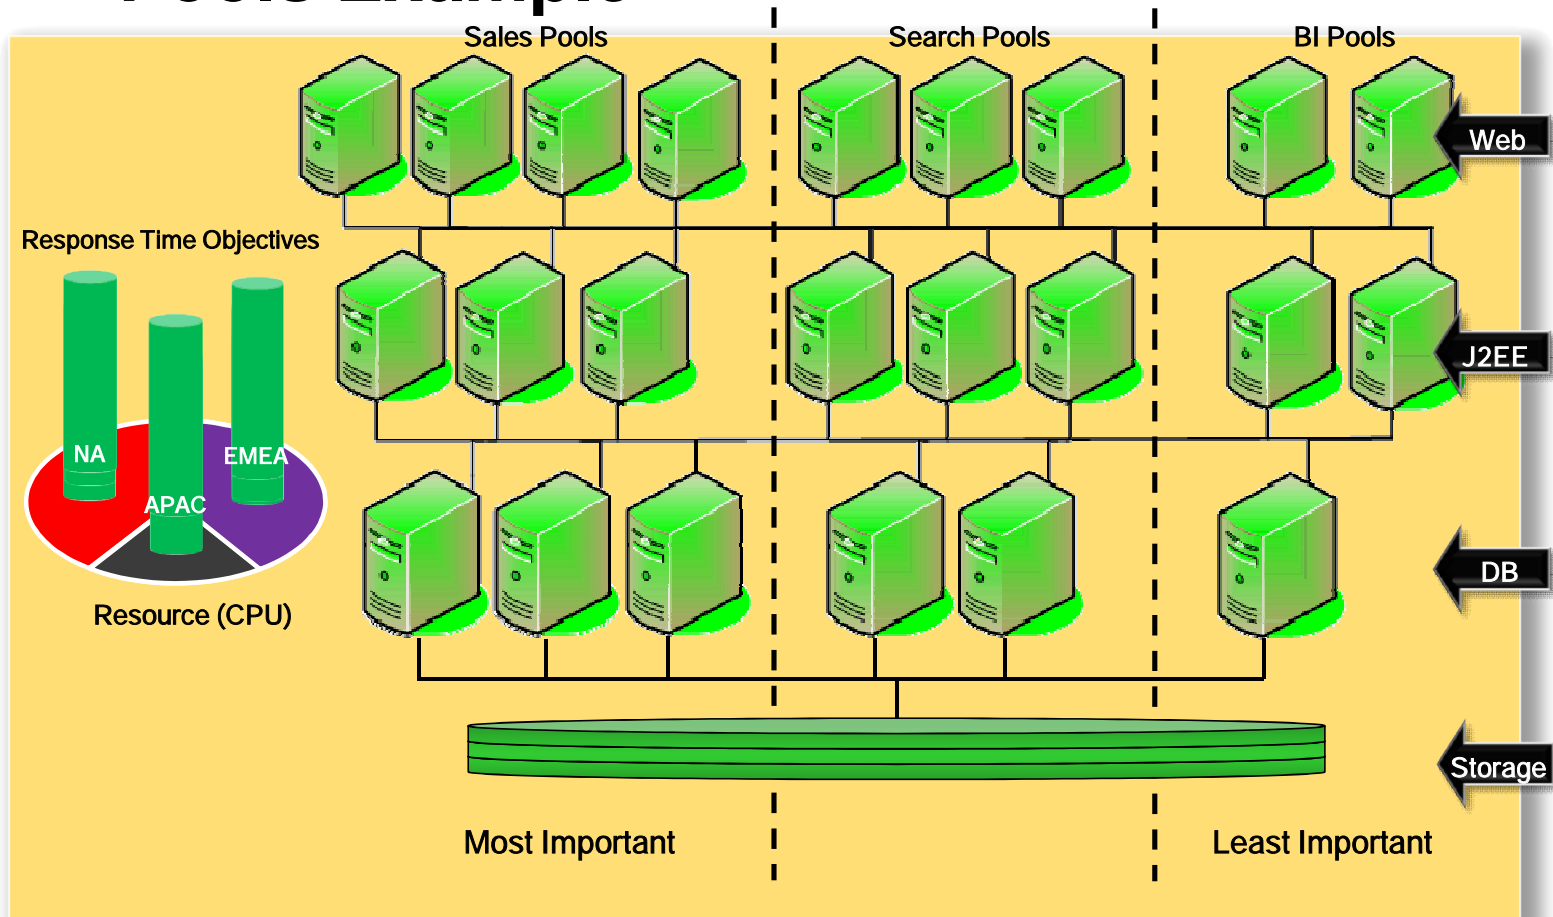
Virtualizes server resources



- Runs all Oracle database applications
- Highly available and scalable on demand
- Adapts to changes in workloads

ORACLE


Dynamic Cluster Partitioning via Server Pools Example



ORACLE

Oracle Database 11g Release 2

Simplified Grid Provisioning




- New intelligent installer - 40% fewer steps to install RAC
- SCAN - Single cluster-wide alias for database connections
- Nodes can be easily repurposed

Oracle Database 11g Release 2

Intra-Node Instance Caging


- Flexible alternative to server partitioning
- Wider platform support than operating system resource managers
- Lower administration overhead than virtualization
- Set CPU_COUNT per instance and enable resource manager

Sum of cpu_counts




Instance Caging Approaches

Partition – For critical workloads
Sum = Total




Over-provision - for Non-Critical/Test Instances
Sum > Total



Oracle Database 11g Release 2

ASM Cluster File System




- General purpose clustered or local file system built on ASM
- Optimized disk layout, Online disk add/drop/rebalance, Integrated mirroring
- Dynamic Volume Management, Read-Only Snapshots

Oracle Database 11g Release

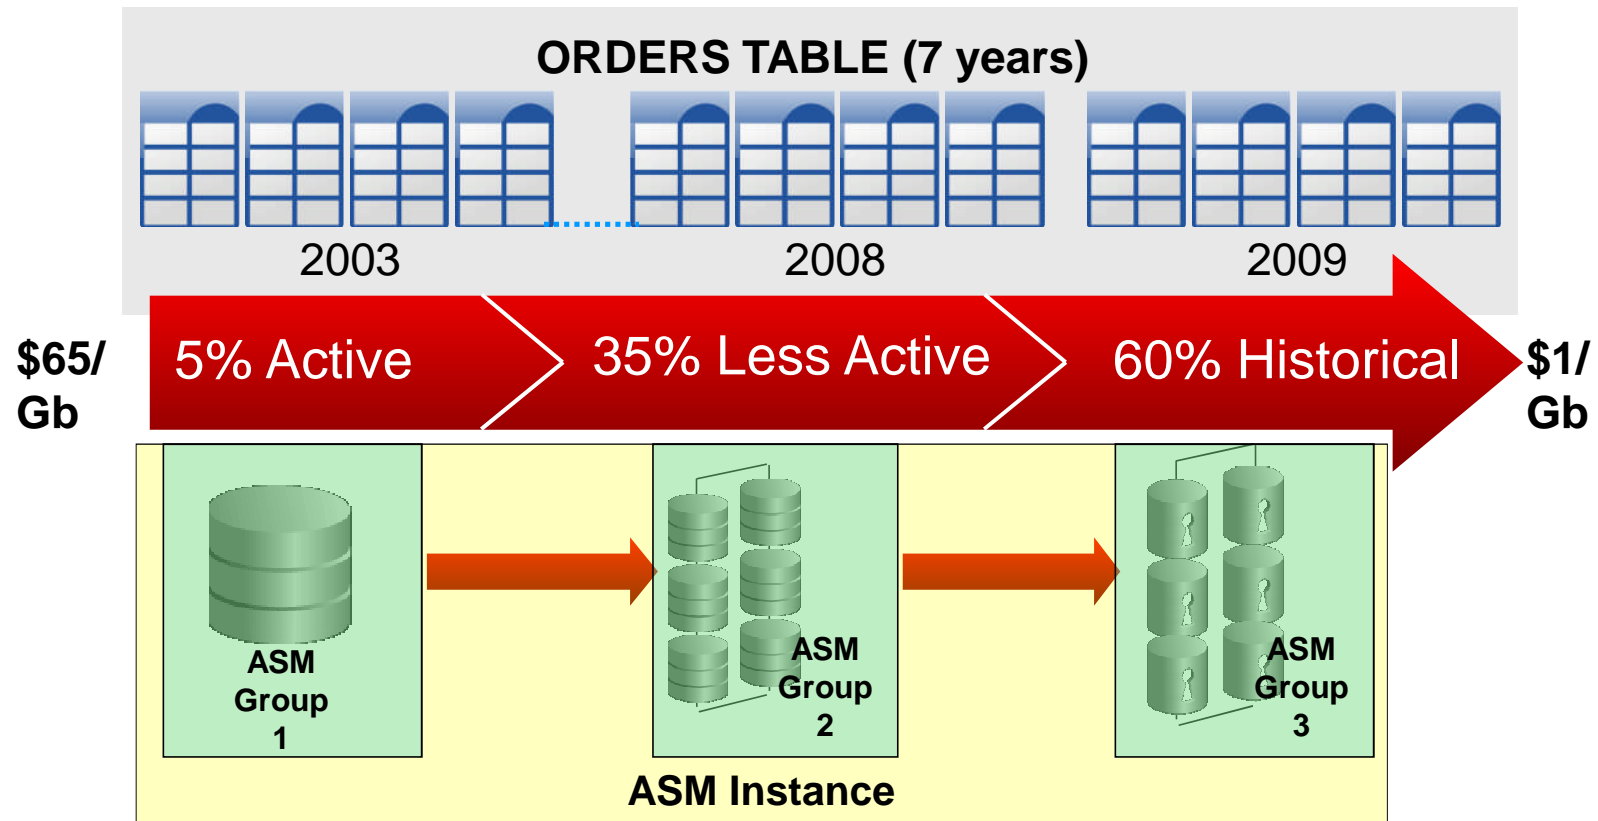
ASM Enhancements

- Improved Management
 - ASM Install & Configuration Assistant (ASMCA)
 - Full Featured ASMCMD
 - ASM File Access Control
 - ASM Disk Group Rename
 - Datafile to Disk Mapping
- Tunable Performance
 - Intelligent Data Placement



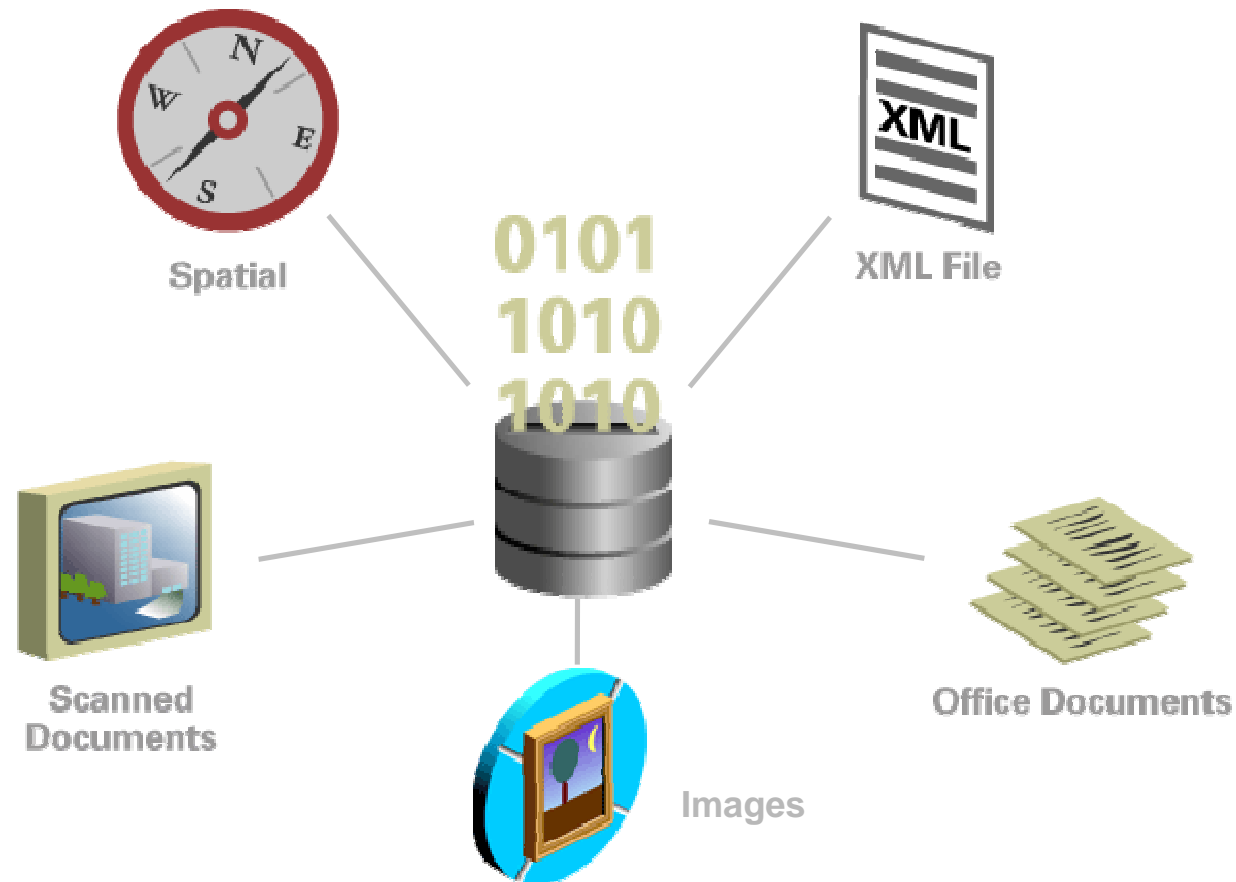
Manage Data Growth Costs

Partition for performance, management and cost



ORACLE

Consolidating All Your Data




ORACLE®


Oracle Secure Files

Consolidate Unstructured Data On the Grid

Read Performance

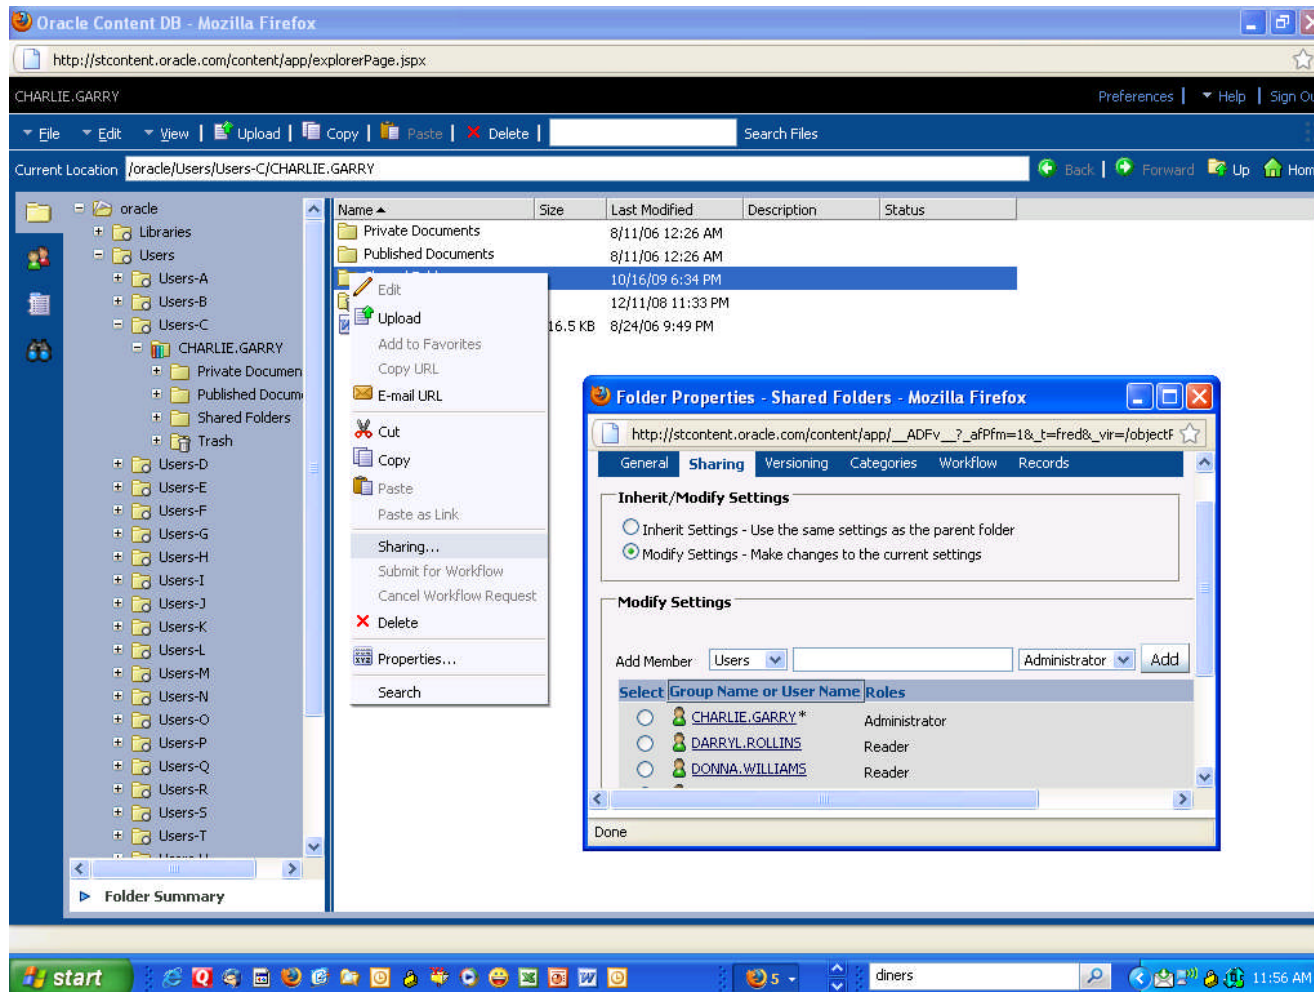


Write Performance



ORACLE

Unstructured Content in the Database



Optimize I/O Performance

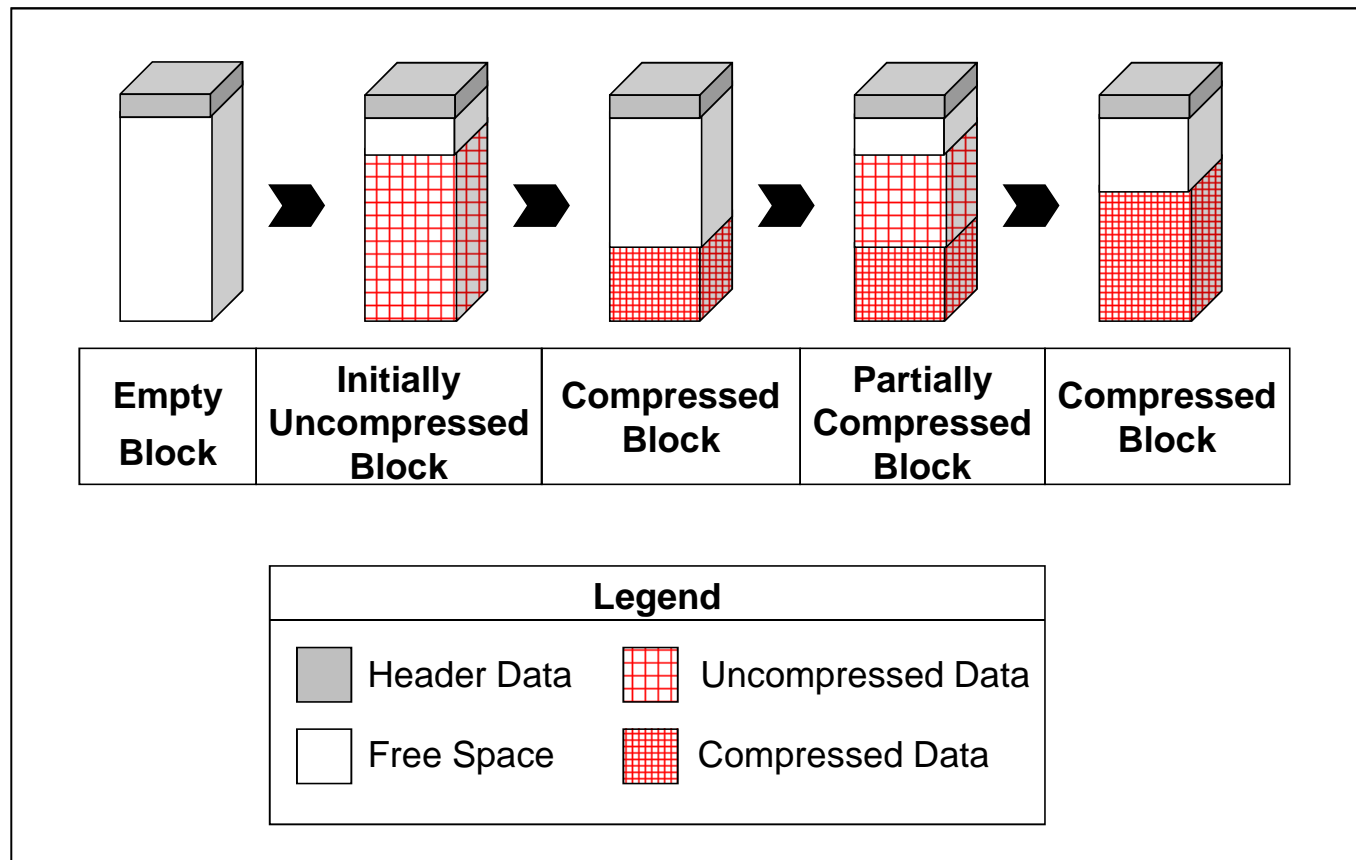
Advanced Compression Option

- Compress large application tables
 - Transaction processing, data warehousing
- Compress all data types
 - Structured and unstructured data types (de-duplication)
- Improve query performance
 - Increased physical I/O efficiency
 - Increased SGA efficiency reduces need for physical read
- Compressed log-redo stream to standby
 - Reduce bandwidth, increase distance between DCs




Up To **4X** Compression

OLTP Table Compression Process



Block-Level *Batch* Compression




- Patent pending algorithm minimizes performance overhead and maximizes compression
- Individual INSERT and UPDATES do not cause recompression
- Compression cost is amortized over several DML operations
- Block-level (Local) compression keeps up with frequent data changes in OLTP environments
 - Others use static, fixed size dictionary table thereby compromising compression benefits
- Extends industry standard compression algorithm to databases
 - Compression utilities such as GZIP and BZ2 use similar adaptive, block level compression

Consolidate onto the Grid

Oracle's Grid Computing Architecture


In-Memory Database Cache (AKA TimesTen)




Real Application Clusters



Automatic Storage Management



**Enterprise
Manager**



Oracle Times Ten 11.2

Released May 2009

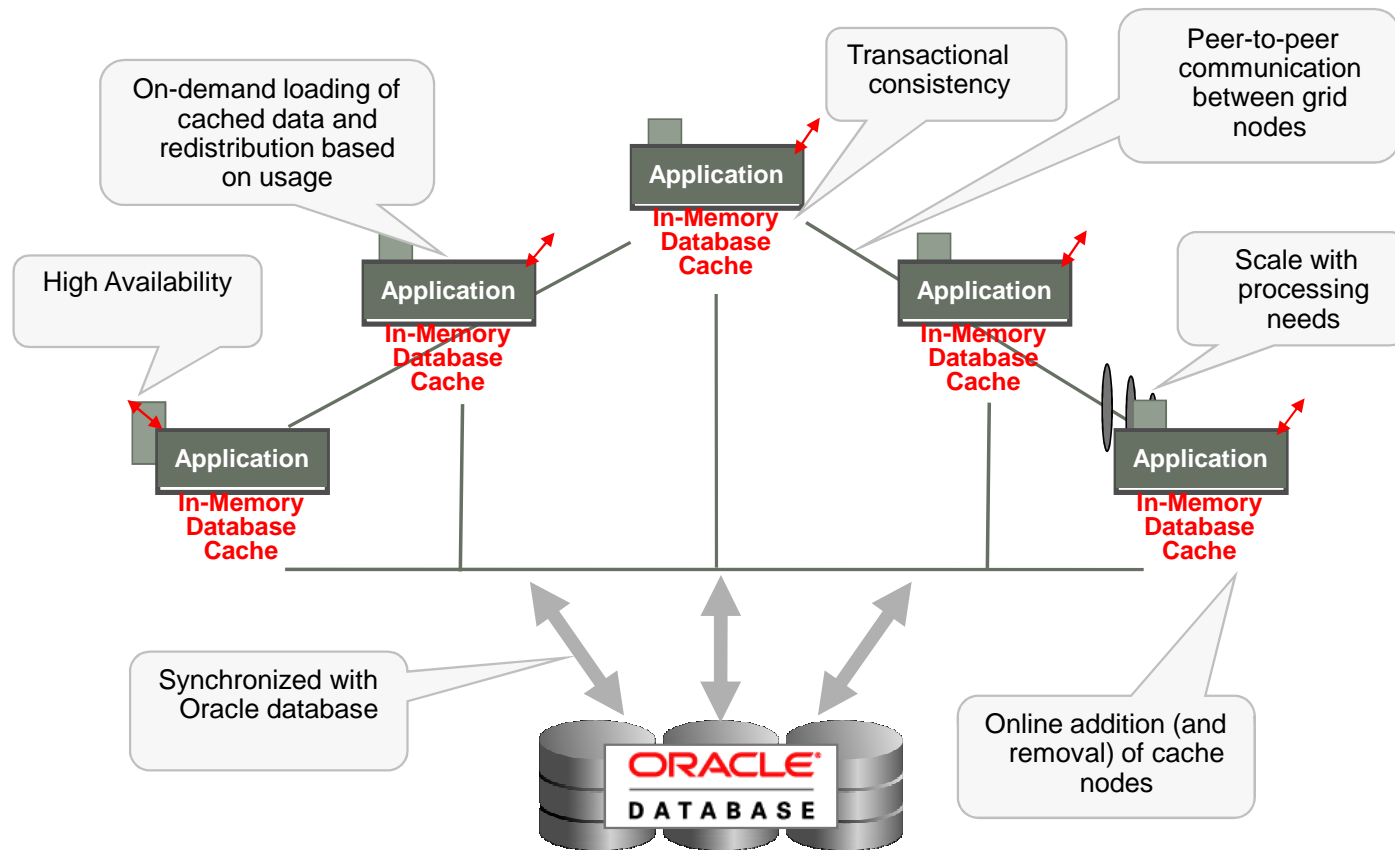


- Improved performance
 - Improved SQL optimization and support for bitmap indexes
 - Improved write throughput and scalability
- Automatic failover
 - Automatic database failover (integration with CRS)
 - Automatic client connections failover and notification
- Enhanced compatibility with Oracle Database
 - Support for Oracle Call Interface (OCI) and Pro*C
 - PL/SQL

ORACLE

In-Memory Database Cache Grid

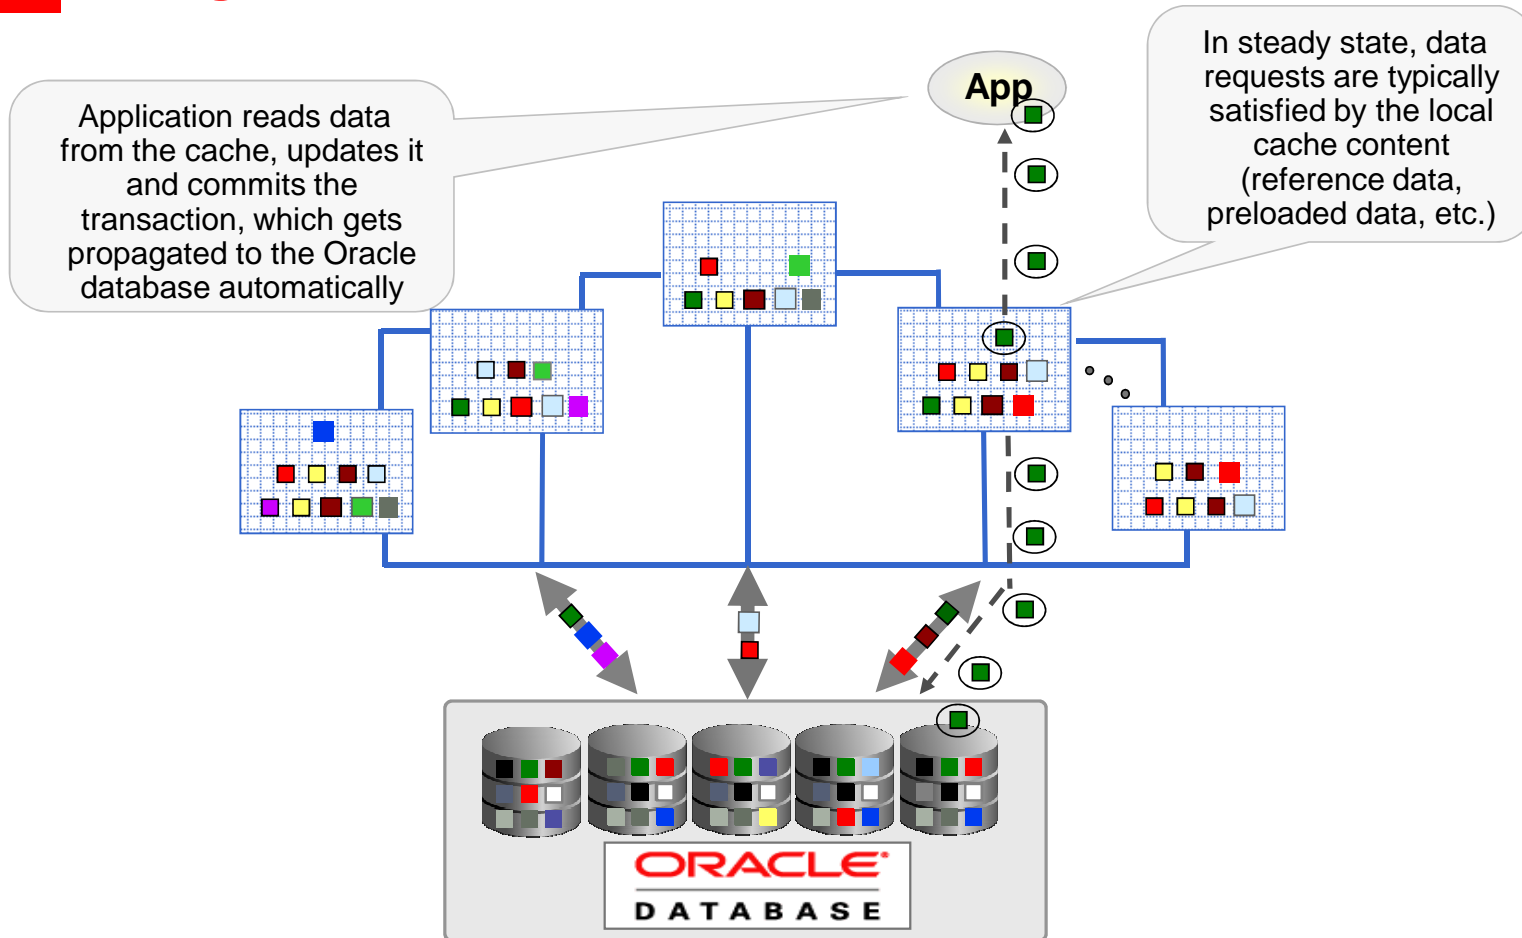
Scaling Middle Tier Resources



ORACLE

In-Memory Database Cache Grid

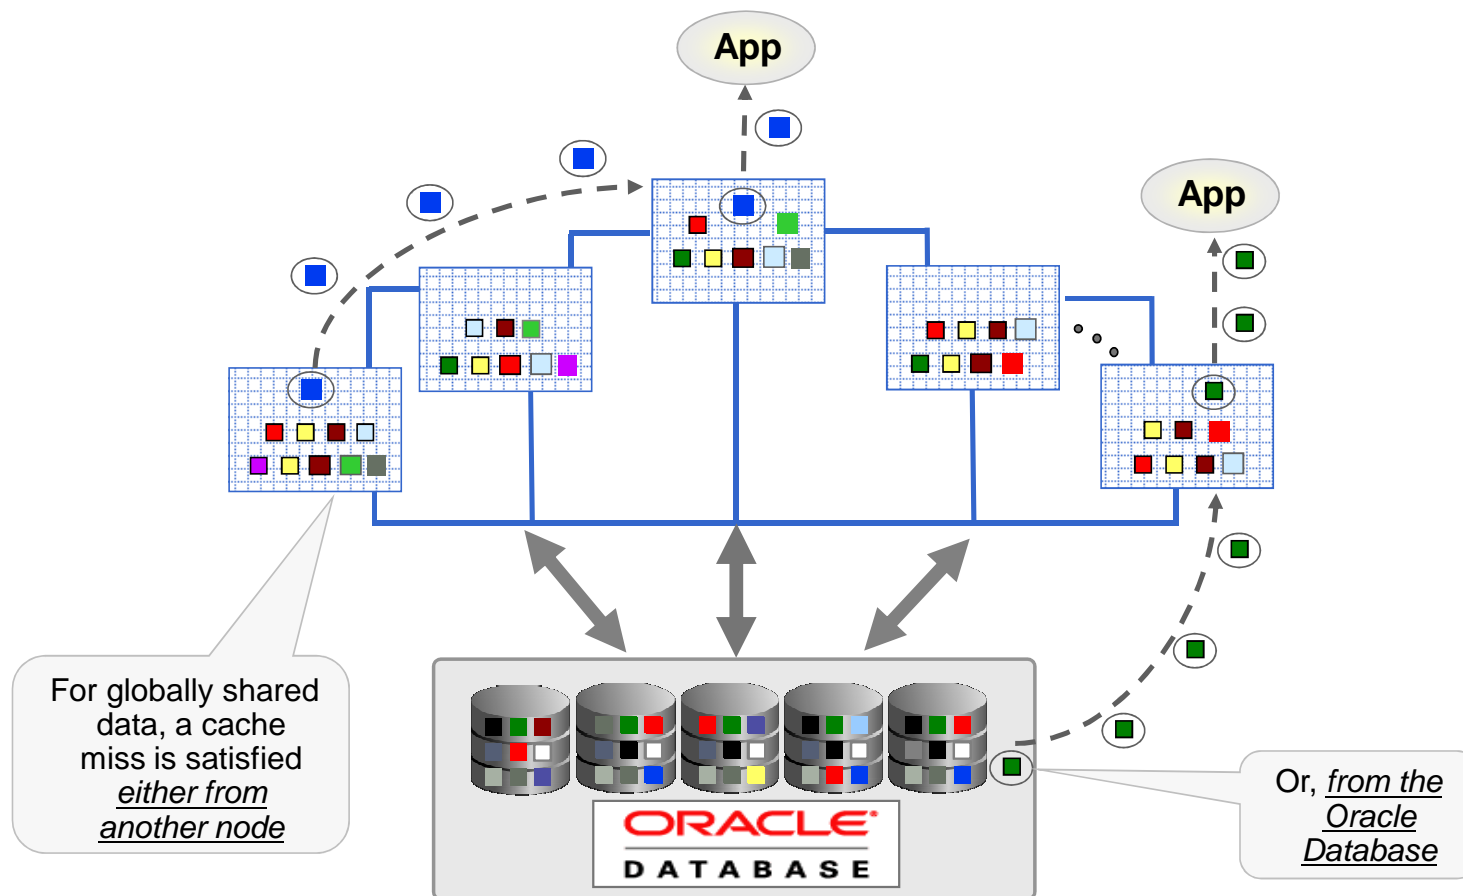
Using Cache Data in the Grid



ORACLE

In-Memory Database Cache Grid


Concurrency and Coherency Across Nodes



ORACLE

In-Memory Database Cache Grid


Concurrency and Coherency Across Nodes



ORACLE

Traditional High Availability


Expensive, idle redundancy



ORACLE


Oracle Maximum Availability Architecture

Fully Utilizing Redundancy



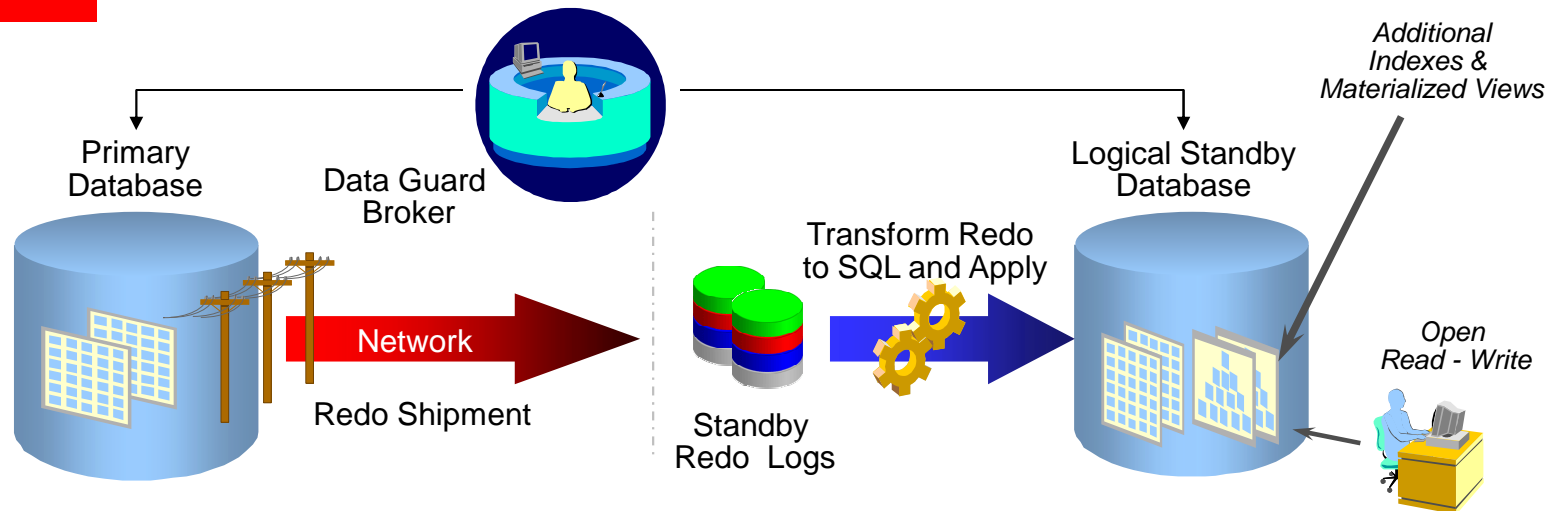
ORACLE

Oracle Data Guard 11g Architecture



ORACLE

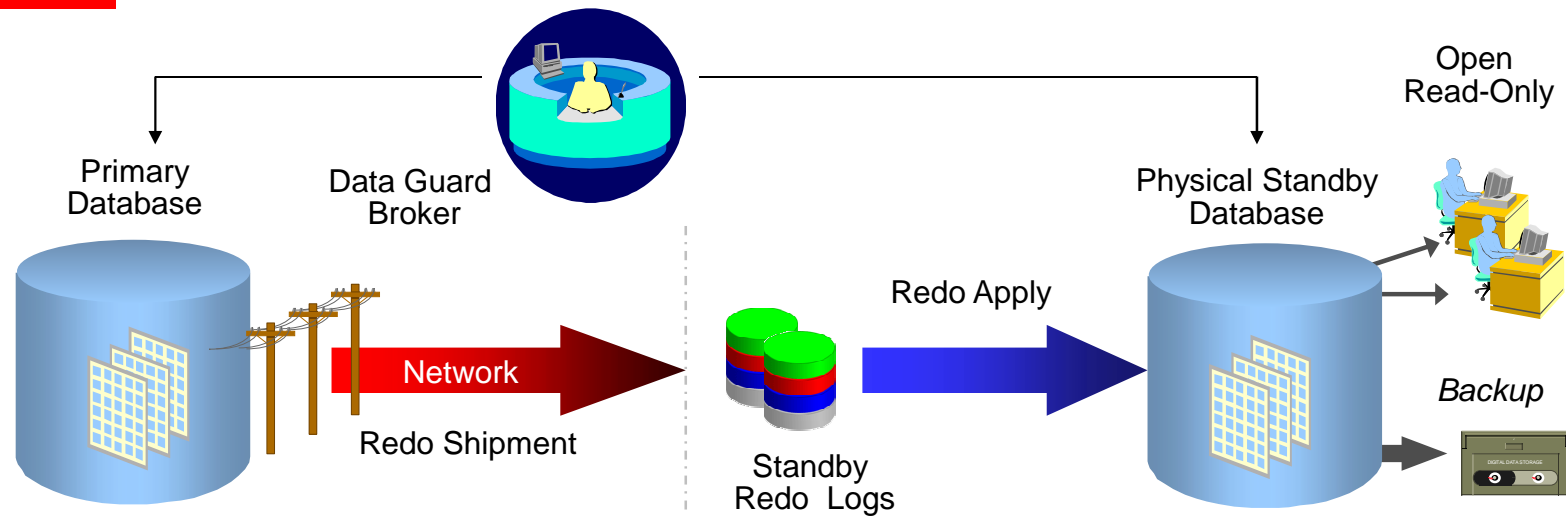
Data Guard SQL Apply (Logical Standby)



- Logical Standby Database is an open, independent, active database
 - Contains the same logical information (rows) as the production database
 - Physical organization and structure can be very different
 - Can host multiple schemas
- Can be queried for reports while redo data is being applied via SQL
- Can create additional indexes and materialized views for better query performance

Data Guard Redo Apply (Physical Standby)

Active Data Guard Option



- Physical Standby Database is a block-for-block copy of the primary database
- Uses the database recovery functionality to apply changes
- While apply is active can be opened in read-only mode for reporting/queries*
- Supports Fast incremental backups, further offloading the production database
- Support for all features and data types
- Support for new “Snapshot Standby” state
- Eliminates the need for separate standby for reporting

ORACLE



Redo Apply or SQL Apply?

Redo Apply

- Physical, block-for-block copy of the primary
- Can be open for read-only queries – supports real-time reporting in 11g
- At role transition, offers assurance that the standby database is an exact replica of the old primary
- Can be used for fast backups
- Higher performance
- No data type restrictions

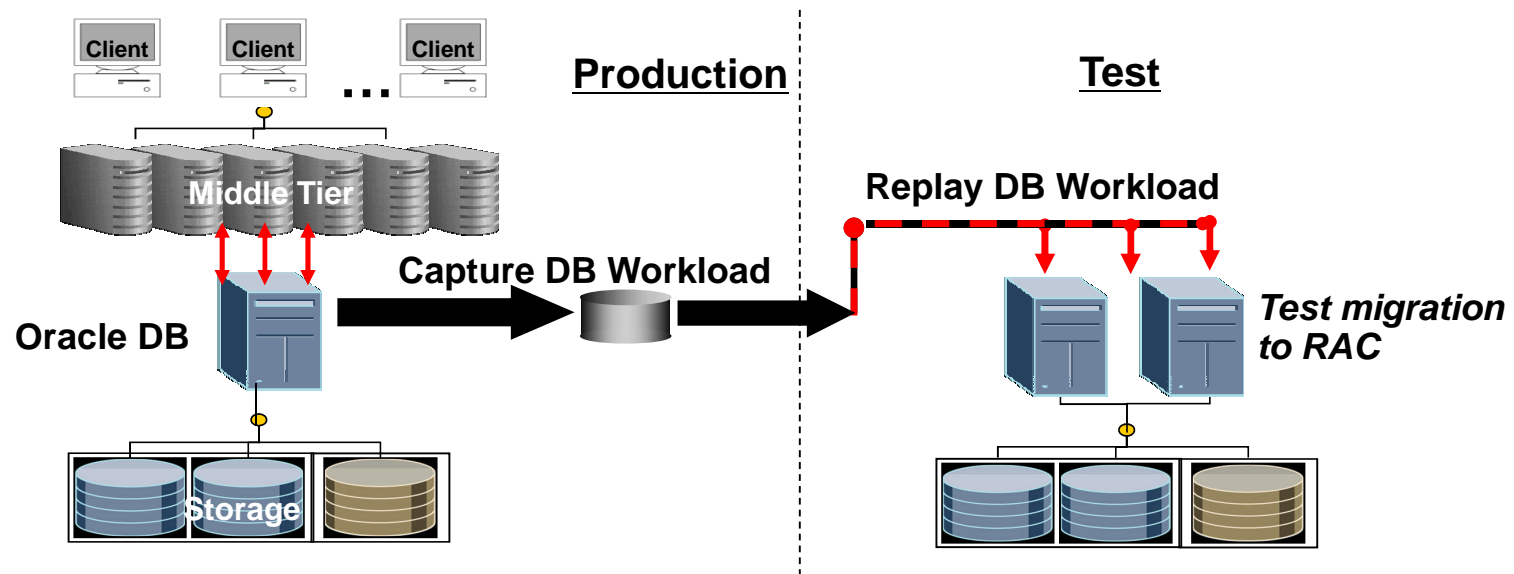
SQL Apply

- Logical, transaction-for-transaction copy of the primary
- Allows creation of additional objects, modification of objects
- Able to skip apply on certain objects
- Is open read-write (data in tables maintained by SQL Apply can not be changed)
- Supports real-time reporting
- Has datatype restrictions

Make Change Safe -

Real Application Testing: Database Replay

- Recreate actual production database workload in test environment
- No test development required
- Replay workload in test with production timing
- Analyze & fix issues before production



ORACLE

Make Change Safe – Real Application Testing: SQL Performance Analyzer


SQL Replay Analysis Result: STE_PAT

Task Name **STE_PAT**
Task Owner **SYSTEM**
Task Description

SQL Tuning Set Name **STE_123**
STS Owner **SYSTEM**


Global Statistics

Projected Workload Buffer Gets



Improvement Impact [+86.692%](#)
Regression Impact [-0.080%](#)
Overall Impact [+86.612%](#)

SQL Statement Count




Recommendations


Run SQL Tuning Advisor to tune regressed SQL statements.

[Run SQL Tuning Advisor](#)

Projected Workload Buffer Gets Distribution

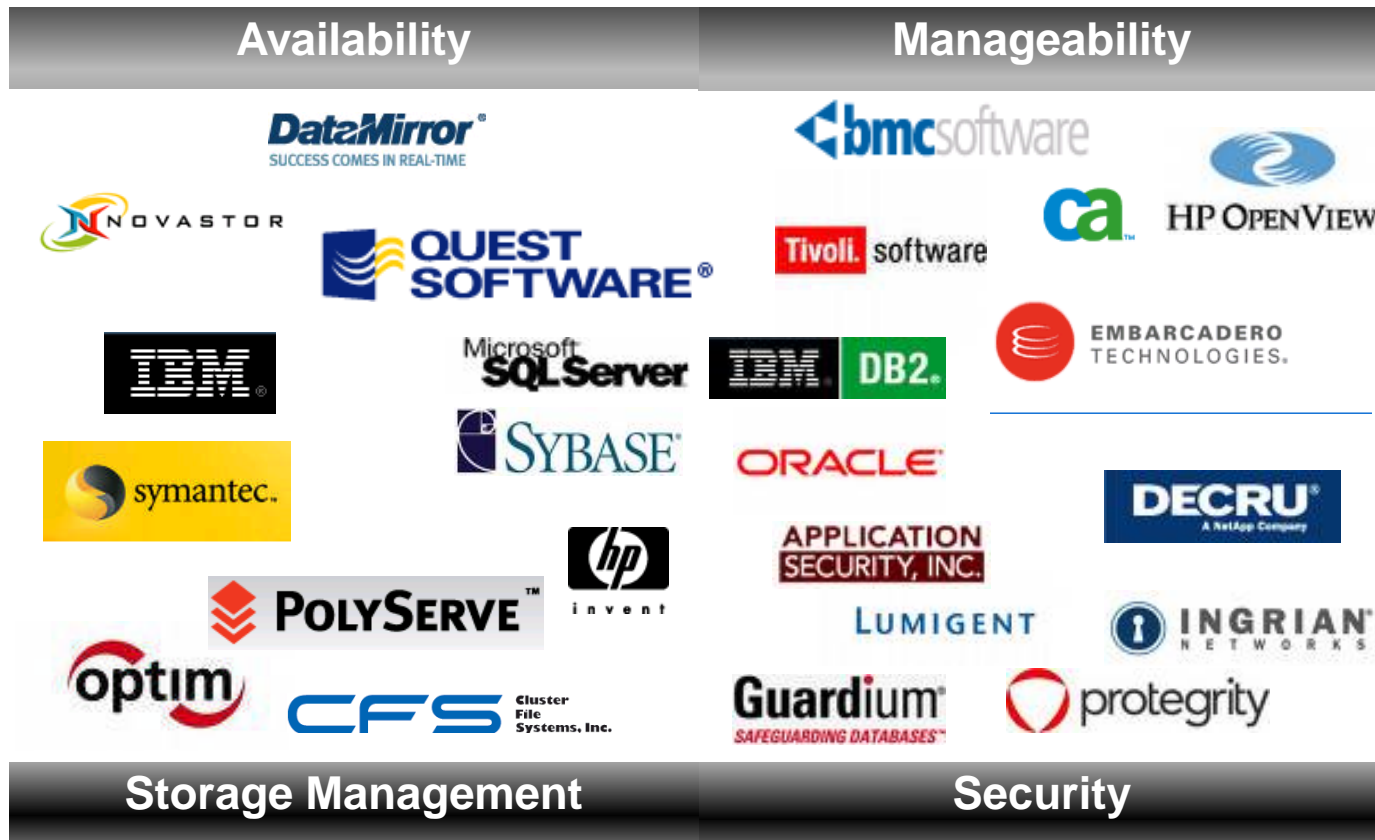


Single Execution SQL Statement Count Distribution



ORACLE

Is This Your Software Portfolio?



ORACLE®

Software Rationalization

Availability

Oracle Clusterware
Oracle Real Application Clusters
Oracle Secure Backup
Oracle Data Guard
Flashback Operations
Online Operations

ORACLE
DATABASE **11^g**

Automatic Storage Management
Automatic Space Management
Disk based Backup/Recovery
Compression
Partitioning
Exadata Storage

Manageability

Diagnostics Pack
Tuning Pack
Change Management Pack
Configuration Management Pack
Provisioning Pack

Fine Grained Access
Identity Management
Transparent Data Encryption
Data Masking Pack
Database Vault
Audit Vault

Storage Management

Security

ORACLE

Managing Complexity

Automated Self-management

Automated:

- Storage
- Memory
- Statistics
- SQL tuning
- Backup and Recovery

Advisory:

- Indexing
- Partitioning
- Compression
- Availability
- Data Recovery




ORACLE

Oracle Stack

Complete, Open and Integrated


- Standard components
- Certified configurations
- Comprehensive security
- Higher availability
- Easier to manage
- Lower cost of ownership
-



ORACLE®

Oracle Database Machine V2

Your Information Infrastructure In a Box



**Oracle Database 11g
Database Machine V2**

Sun Oracle Database Machine

Get on the Grid Faster - OLTP & Data Warehousing



Oracle Database Server Grid

- 8 Database Servers
 - 64 Cores
 - 400 GB DRAM

Exadata Storage Server Grid

- 14 Storage Servers
 - 5TB Smart Flash Cache
 - 336 TB Disk Storage


Unified Server/Storage Network

- 40 Gb/sec Infiniband Links
 - 880 Gb/sec Aggregate Throughput

Completely Fault Tolerant

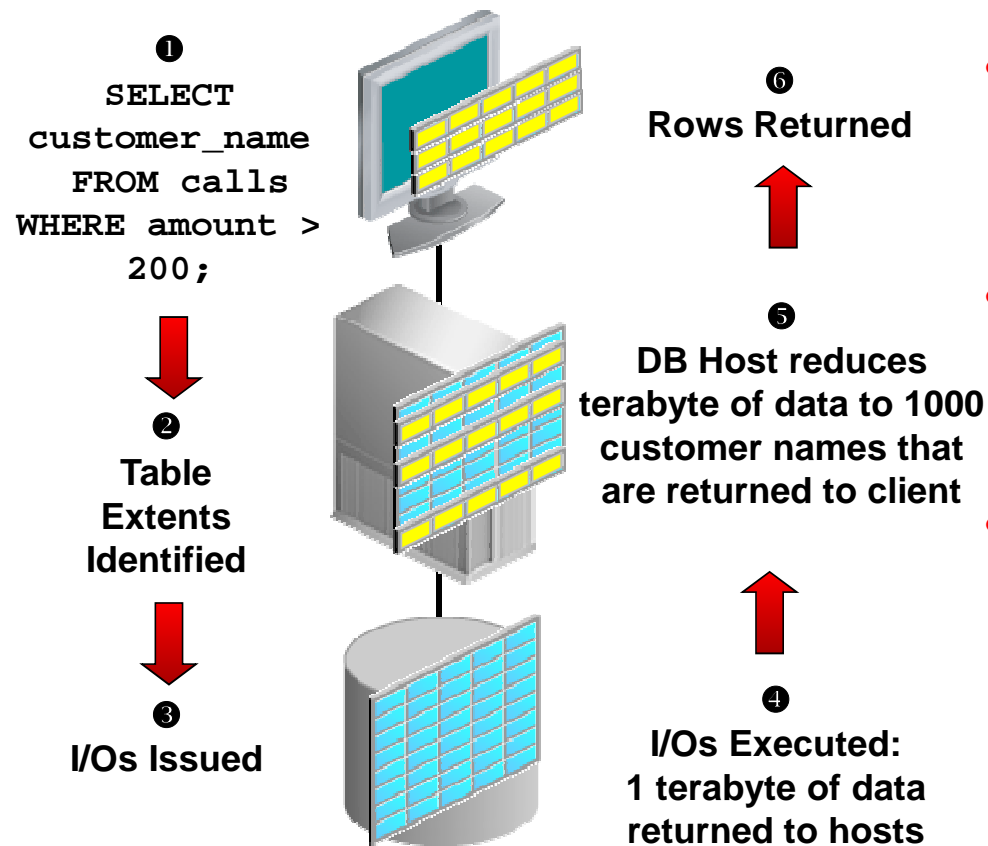
ORACLE

Solutions To Data Bandwidth Bottleneck



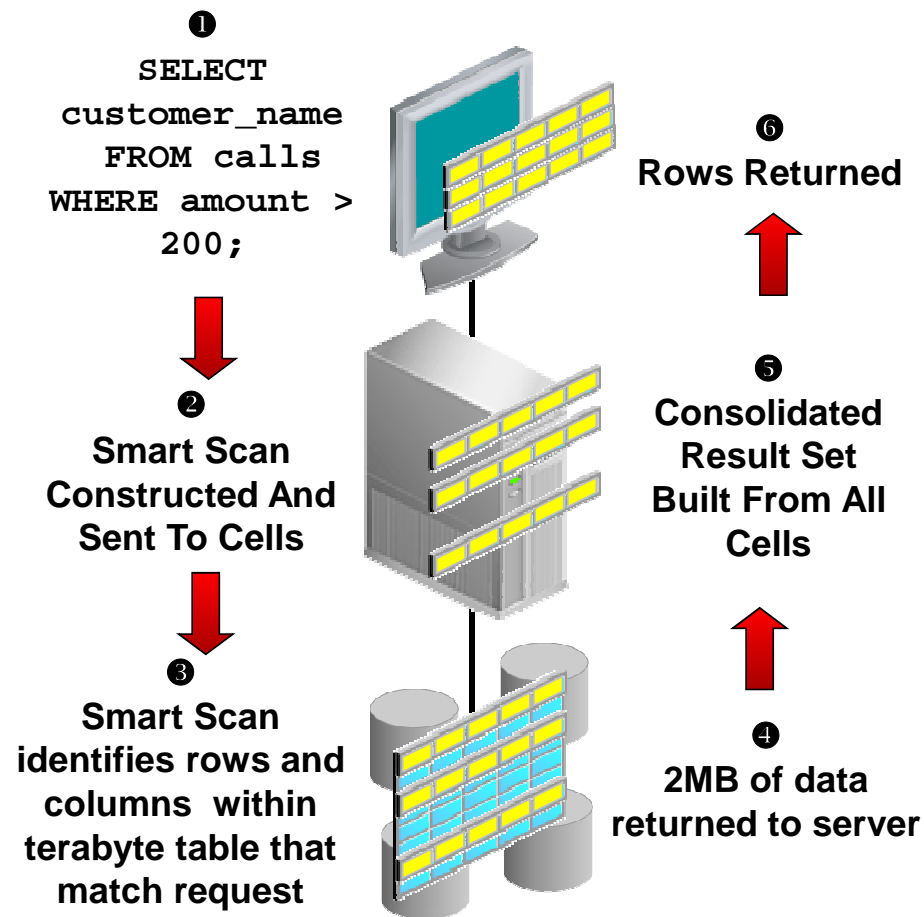
- Add more pipes – Massively parallel architecture
- Make the pipes wider – 10X faster than conventional storage
- Ship less data through the pipes – Process data in storage

Traditional Scan Processing



- With traditional storage, all database intelligence resides in the database hosts
- Very large percentage of data returned from storage is discarded by database servers
- Discarded data consumes valuable resources, and impacts the performance of other workloads

Exadata Smart Scan Processing



- Only the relevant columns
– customer_name
and required rows
– where amount>200
are are returned to hosts
- CPU consumed by predicate
evaluation is offloaded
- Moving scan processing off
the database host frees host
CPU cycles and eliminates
massive amounts of
unproductive messaging
– Returns the needle, not the
entire hay stack



Additional Smart Scan Functionality

- Join filtering
 - Star join filtering is performed within Exadata storage cells
 - Dimension table predicates are transformed into filters that are applied to scan of fact table
- Backups
 - I/O for incremental backups is much more efficient since only changed blocks are returned
- Create Tablespace (file creation)
 - Formatting of tablespace extents eliminates the I/O associated with the creation and writing of tablespace blocks

Oracle Exadata Storage Server

Hybrid Columnar Compression

- Data stored by column and then compressed
- Useful for data that is bulk loaded or moved
- **Query mode** for data warehousing
 - Typical 10X compression ratios
 - Scans improve accordingly
- **Archival mode** for old data
 - Typical 15X up to 50X compression ratios



ORACLE



Oracle Database 11g Release 2

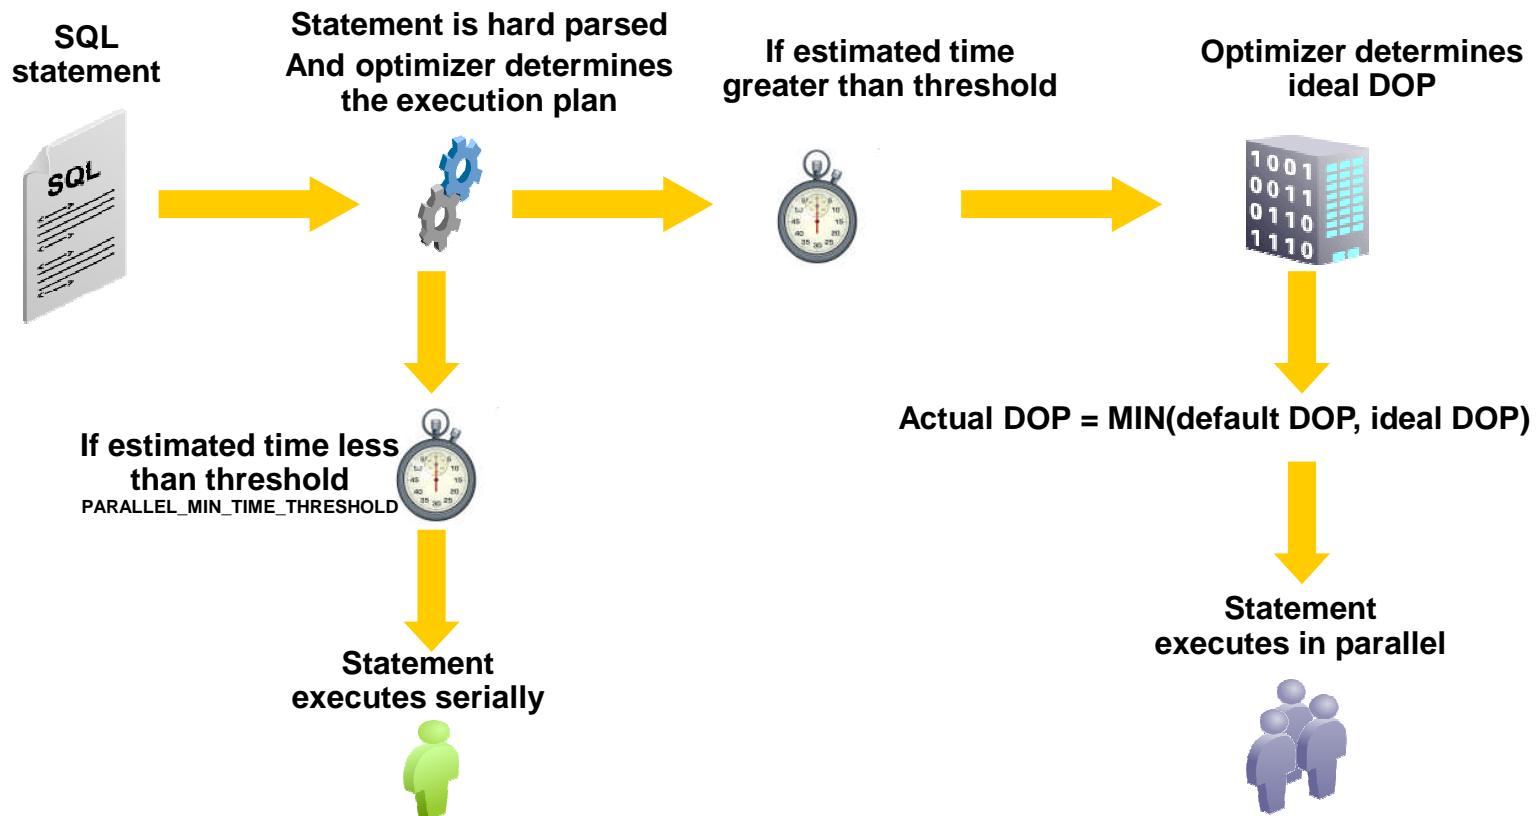
Automated Degree of Parallelism

- Currently tuning parallelism is a manual process
 - one degree of parallelism does not fit all queries
 - too much parallelism can flood system
- Automated Degree of Parallelism automatically decides
 - If a statement will execute in parallel or not (serial execution would take longer than a set threshold – 30 secs)
 - What degree of parallelism the statement will use
- Optimizer derives the DoP from the statement based on resource requirements
 - Uses the cost of all scan operations (full table, index etc...)
 - Balanced against a max limit of parallelism
 - Assigned DOP is shown in Notes section of EXPLAIN PLAN

ORACLE

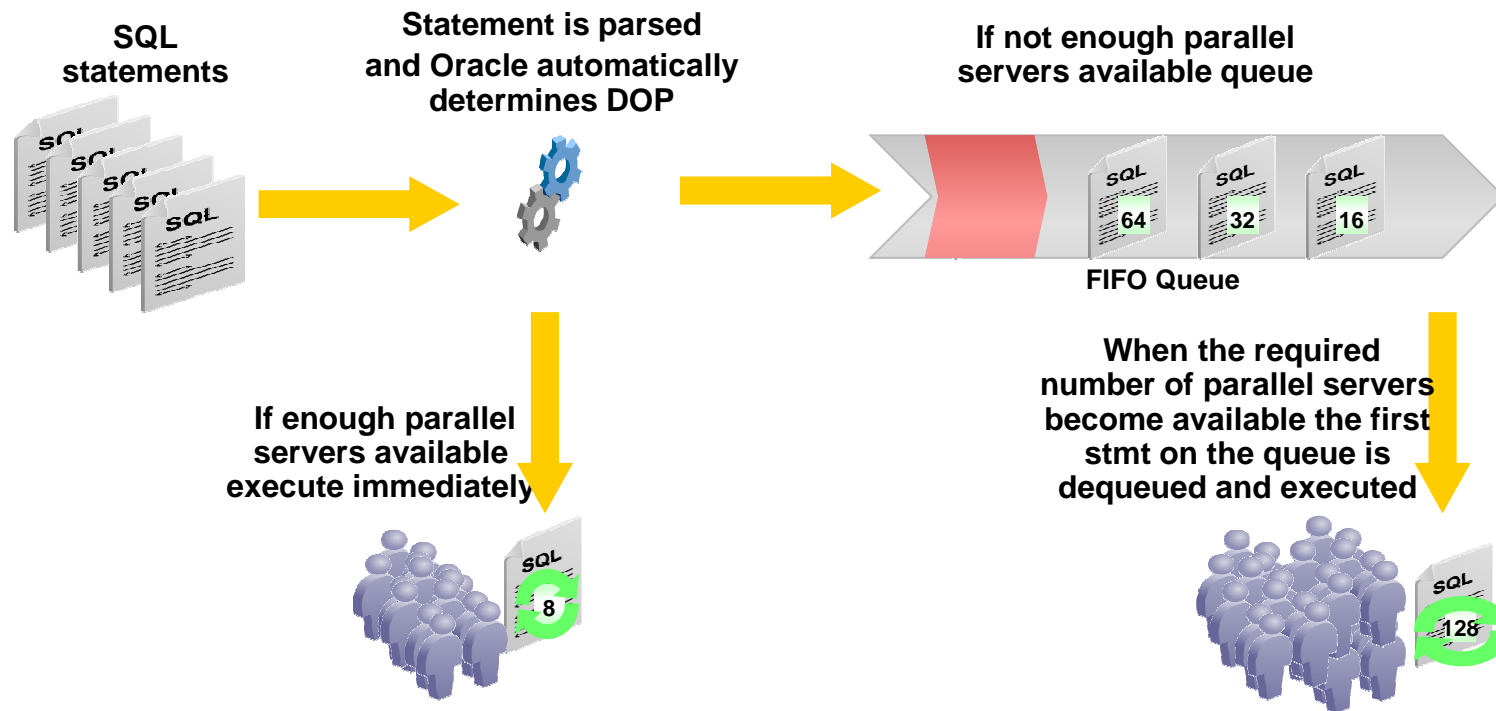
Automated Degree of Parallelism

How it works



Parallel Statement Queuing

How it works




Oracle Database 11g Release 2

In-Memory Parallel Execution

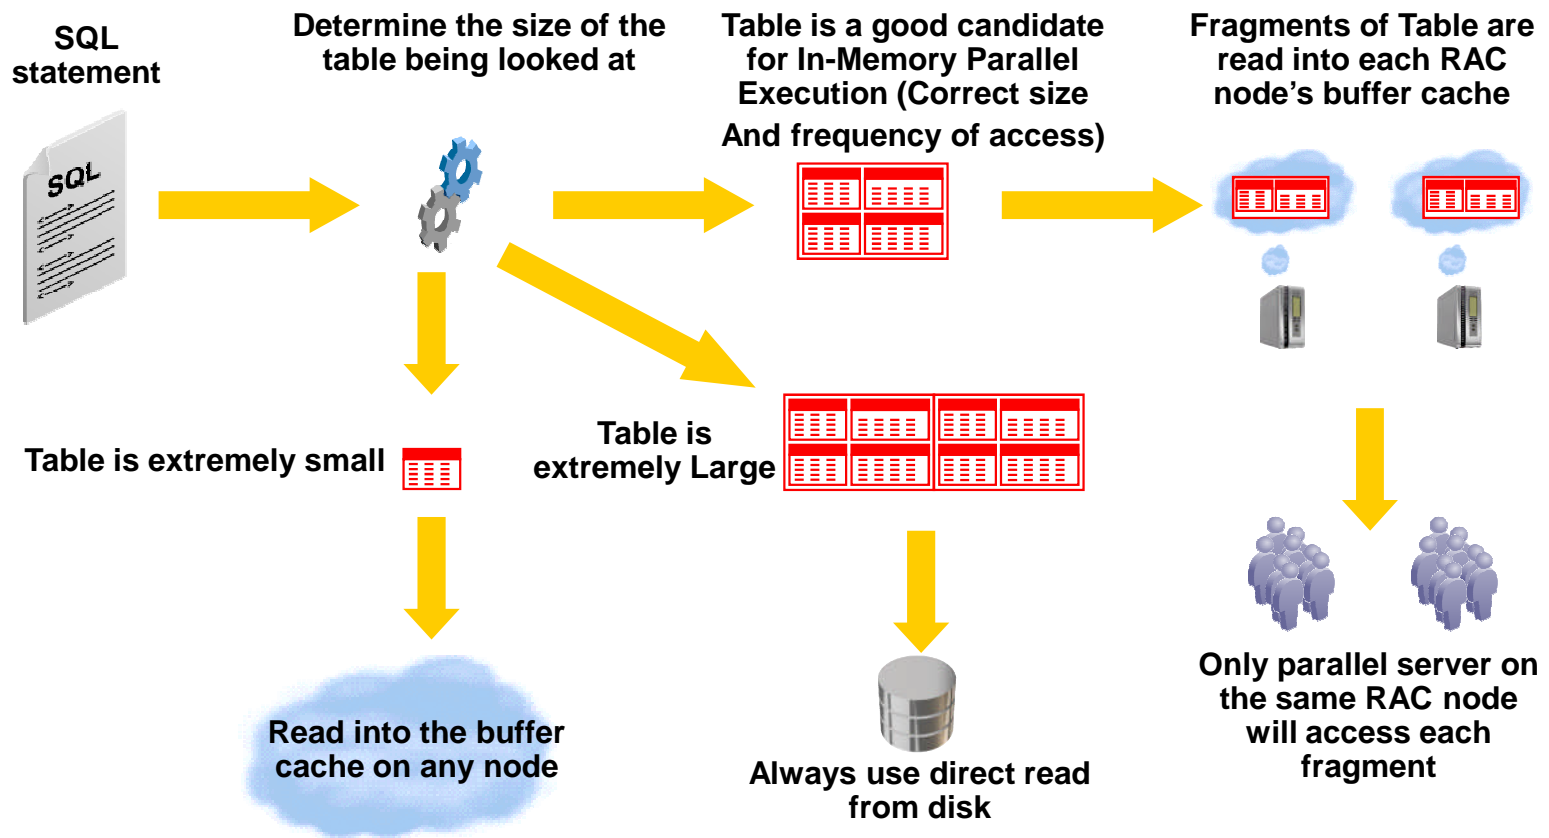
- New commodity servers have have large amounts of memory
- Data Compression also means more data in memory
- Intelligent algorithm places fragments of a table in memory on different nodes
- In Memory Parallel Queries are then executed on the corresponding nodes
- Removes need to perform disk I/O for queries on large tables

**Real Application
Clusters**



In-Memory Parallel Execution

How it works: **PARALLEL_DEGREE_POLICY = AUTO**



ORACLE




Controlling Auto DOP, Queuing and In-Memory Execution

- `PARALLEL_DEGREE_POLICY`
 - `MANUAL` – disables Auto DOP, statement queuing and in-memory execution and defaults to behavior prior to 11gR2
 - `LIMITED` – Enables Auto DOP for some statements
 - Those with hints or that access tables and indexes created with `PARALLEL` clause
 - Disables queuing and in-memory execution
 - `AUTO` – enables all

World's Fastest DW Machine

- Clustered commodity servers
 - large amounts of memory
- Compress more data in memory
- Intelligent algorithm
 - Places table fragments in memory on different nodes
- Reduces need for disk I/O
 - Speeds query execution



ORACLE




World's Fastest OLTP Machine

With Sun FlashFire Technology on Exadata

- Huge semiconductor memory hierarchy
 - 400 Gigabytes DRAM
 - 5 TB Smart Flash Cache – Not Flash Disk !!!
- 1 Million random I/Os per second
 - Eliminates most physical disk I/O
- 3x database compression for OLTP
 - Compressed 1.2 TB Database in DRAM
 - Compressed 15 TB Database in Flash Cache

Start Small and Grow



Oracle Database Machine with Exadata Storage

- Simplicity of Shared-Everything (i.e. traditional SMP)
 - No physical design tricks required (e.g. tables replicated to each node)
 - Adaptive to different workloads (not just read-only)
- Scalability of Shared-Nothing
 - Each additional Exadata Storage Server adds CPU and I/O bandwidth
 - Linear scalability
- Compression of a Column-wise database
 - Typical 15x to 25X up to 50x compression
 - Improved query performance for queries against subset of columns
- Simplified Information Infrastructure!




ORACLE

Oracle Stack

Complete, Open and Integrated

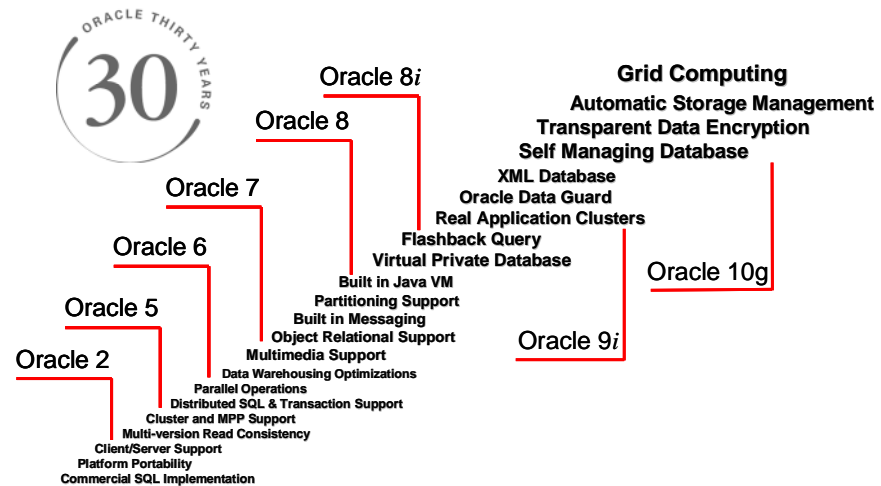
- Standard components
- Certified configurations
- Comprehensive security
- Higher availability
- Easier to manage
- Lower cost of ownership
-



ORACLE®


Design Goals 2000 – 2009 and Beyond

- Take advantage of HW price/performance curve
- Improve utilization rates for servers/processors, storage, memory
- Increasingly automate repeatable, labor intensive tasks
- Reduce or eliminate the risk of change or user error
- Simplify the Information Architecture



Oracle Database 11g Release 2

What are my upgrade paths?



ORACLE

Customer Resources

OTN Upgrade Page

<http://www.oracle.com/technology/products/database/oracle11g/upgrade/index.html>

shortcuts | **GETTING STARTED** | **DOWNLOADS** | **DOCUMENTATION** | **FORUMS** | **ARTICLES** | **SAMPLE CODE** | **TUTORIALS**

Printer View | E-mail this page | Bookmark

ORACLE
DATABASE **11g** Oracle Database Upgrade

Upgrading to Oracle Database 11g provides the latest in efficient, reliable, secure data management for mission-critical on-line transaction processing applications, query-intensive data warehouses, and content management and Web2.0 applications. The right planning, preparation, and upgrade steps will make the upgrade process simpler, faster, and more predictable from start to finish.

Latest (May 2009)

- [Fujitsu SPARC Enterprise: Migrating from Oracle9i Database to Oracle Database 11g](#) (Posted May 2009)
 - [11g R1 Upgrade Workshop Presentation](#)
 - [May 2009 Update](#)
- [Part 1 of 2](#) (May 2009 - ~5.6M)
- [Part 2 of 2](#) (May 2009 - ~6.1M)
- [Upgrade Oracle 9i->11g: A Customer Experience](#) (November 2008)

Technical Information

- [Different Upgrade Methods for Upgrading your Database](#) (Metalink)
- [Complete Checklist for Manual Upgrades - 11gR1](#) (Metalink)
- [Complete Checklist for Upgrades - 11gR1 Using DBUA](#) (Metalink)
- [Database Server Upgrade/Downgrade Compatibility Matrix](#) (Metalink)

Database Upgrade Services

The Database Upgrade Partner Program includes Oracle partners who offer Oracle Database 11g Upgrade Services to

Documentation

- [11g Upgrade Companion](#) (Metalink)
- [Oracle 11g Release 1 Documentation](#)
- [Oracle 11g Release 1 Upgrade Guide](#)

Upgrade Resources/Papers

- [Discussion Forum](#)
- [Benefits of Upgrading to Oracle Database 11g](#) (July 2008)
- [Database Rolling Upgrade Using Transient Logical Standby: Oracle Data Guard 11g](#) (June 2008)
- [Database Rolling Upgrade Using Data Guard SQL Apply - Oracle Database 11g and 10gR2](#) (April 2008)

ORACLE



For More Information

<http://search.oracle.com>

oracle database 11g



or

www.oracle.com/database