

ORACLE®



ORACLE®

Building Innovative Applications with Oracle Database 11g


Mark Drake, Manager, Product Management



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 remains at the sole discretion of Oracle.

Agenda

- Database
 - SQL, PL/SQL, Secure Files
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - In-Memory Database Cache
 - Partitioning
- High Availability Architecture
- Enterprise Manager Packs
- SQL Developer Data Modeler
- SQL Developer





Enhanced SQL and PL/SQL Capabilities Oracle Database 11g

- Pivot
- New analytical window functions
- Real Time SQL Monitoring
- Sequences
 - No Longer need a select
- PL/SQL Hierarchical Profiler
- PL/Scope



Pivot SQL Command

- Simpler and more maintainable
- Better performance


```
SELECT * FROM (SELECT dname ,JOB
FROM EMP e,dept d
where e.deptno = d.deptno)
PIVOT
(COUNT(*) FOR JOB IN (
  'PRESIDENT' as President,
  'MANAGER' as Manager,
  'ANALYST' as Analyst,
  'SALESMAN' as Salesman,
  'CLERK' as Clerk)
)
SELECT  dname,
        sum(decode(job,'PRESIDENT',1,0)) president,
        sum(decode(job,'MANAGER',1,0))  MANAGER,
        sum(decode(job,'ANALYST',1,0))  ANALYST,
        sum(decode(job,'SALESMAN',1,0))  SALESMAN,
        sum(decode(job,'CLERK',1,0))    CLERK
FROM EMP e,dept d
where e.deptno = d.deptno
GROUP BY dname
```



PL/SQL New Features

- Direct Sequence Assignments
 - Old
 - `select my_seq.nextval into l_var from dual;`
 - New
 - `l_var := my_seq.nextval;`
- PL/SQL Hierarchical Profiler
- PL/Scope

Statistics and SQL Analytics



- **Ranking functions**
 - rank, dense_rank, cume_dist, percent_rank, ntile
- **Window Aggregate functions**
(moving and cumulative)
 - Avg, sum, min, max, count, variance, stddev, first_value, last_value
- **LAG/LEAD functions**
 - Direct inter-row reference using offsets
- **Reporting Aggregate functions**
 - Sum, avg, min, max, variance, stddev, count, ratio_to_report
- **Statistical Aggregates**
 - Correlation, linear regression family, covariance
- **Linear regression**
 - Fitting of an ordinary-least-squares regression line to a set of number pairs.
 - Frequently combined with the COVAR_POP, COVAR_SAMP, and CORR functions.
- **Descriptive Statistics**
 - average, standard deviation, variance, min, max, median (via percentile_count), mode, group-by & roll-up
 - DBMS_STAT_FUNCS: summarizes numerical columns of a table and returns count, min, max, range, mean, stats_mode, variance, standard deviation, median, quantile values, +/- n sigma values, top/bottom 5 values
- **Correlations**
 - Pearson's correlation coefficients, Spearman's and Kendall's (both nonparametric).
- **Cross Tabs**
 - Enhanced with % statistics: chi squared, phi coefficient, Cramer's V, contingency coefficient, Cohen's kappa
- **Hypothesis Testing**
 - Student t-test, F-test, Binomial test, Wilcoxon Signed Ranks test, Chi-square, Mann Whitney test, Kolmogorov-Smirnov test, One-way ANOVA
- **Distribution Fitting**
 - Kolmogorov-Smirnov Test, Anderson-Darling Test, Chi-Squared Test, Normal, Uniform, Weibull, Exponential

Note: Statistics and SQL Analytics are included in Oracle Database Standard Edition



Demo



Files in the Database Reinvented





- Best of Both Worlds
- File Capabilities
 - File System Interface
 - High Performance
 - Compression
 - Encryption
 - Deduplication
 - HSM
- Database Capabilities
 - Transactions
 - Query Consistency
 - Advanced Backup and Recovery
 - Powerful Security
 - Flashback
 - Scale up SMPs
 - Scale out Clusters

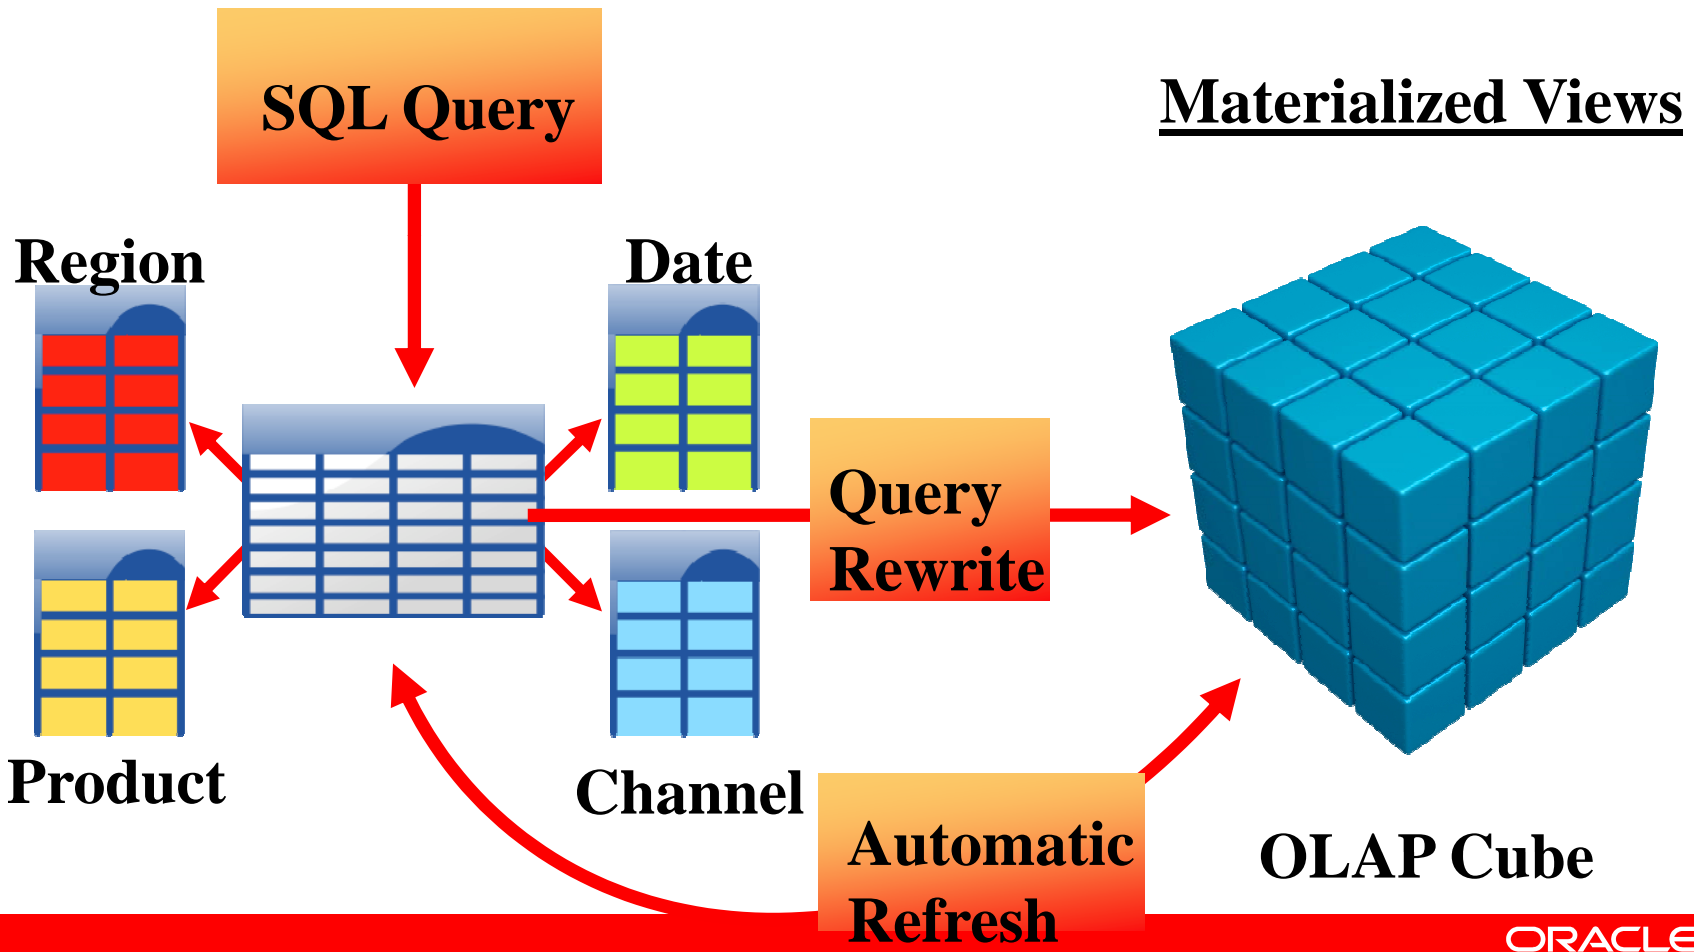
- Files are an integral part of modern database applications
 - Product images, contracts, XML, ETL files, manuals, etc.
- Applications developers want to store business data files in the database to benefit from transactional consistency, and unify HA and Security
 - Poor performance, limited functionality, and lack of access by existing file based tools have held them back
- Oracle Database 11g reinvents files in the database
- SecureFiles provides super fast and powerful file storage
 - Removes performance barrier to storing files in the database
- DBFS provides a file system interface to files in the DB
 - Enables existing file based tools to easily access DB files

SecureFiles Performance

- Performance compared to Linux FS
 - Tests run using both SecureFiles and ext3 in metadata journaling only, no network

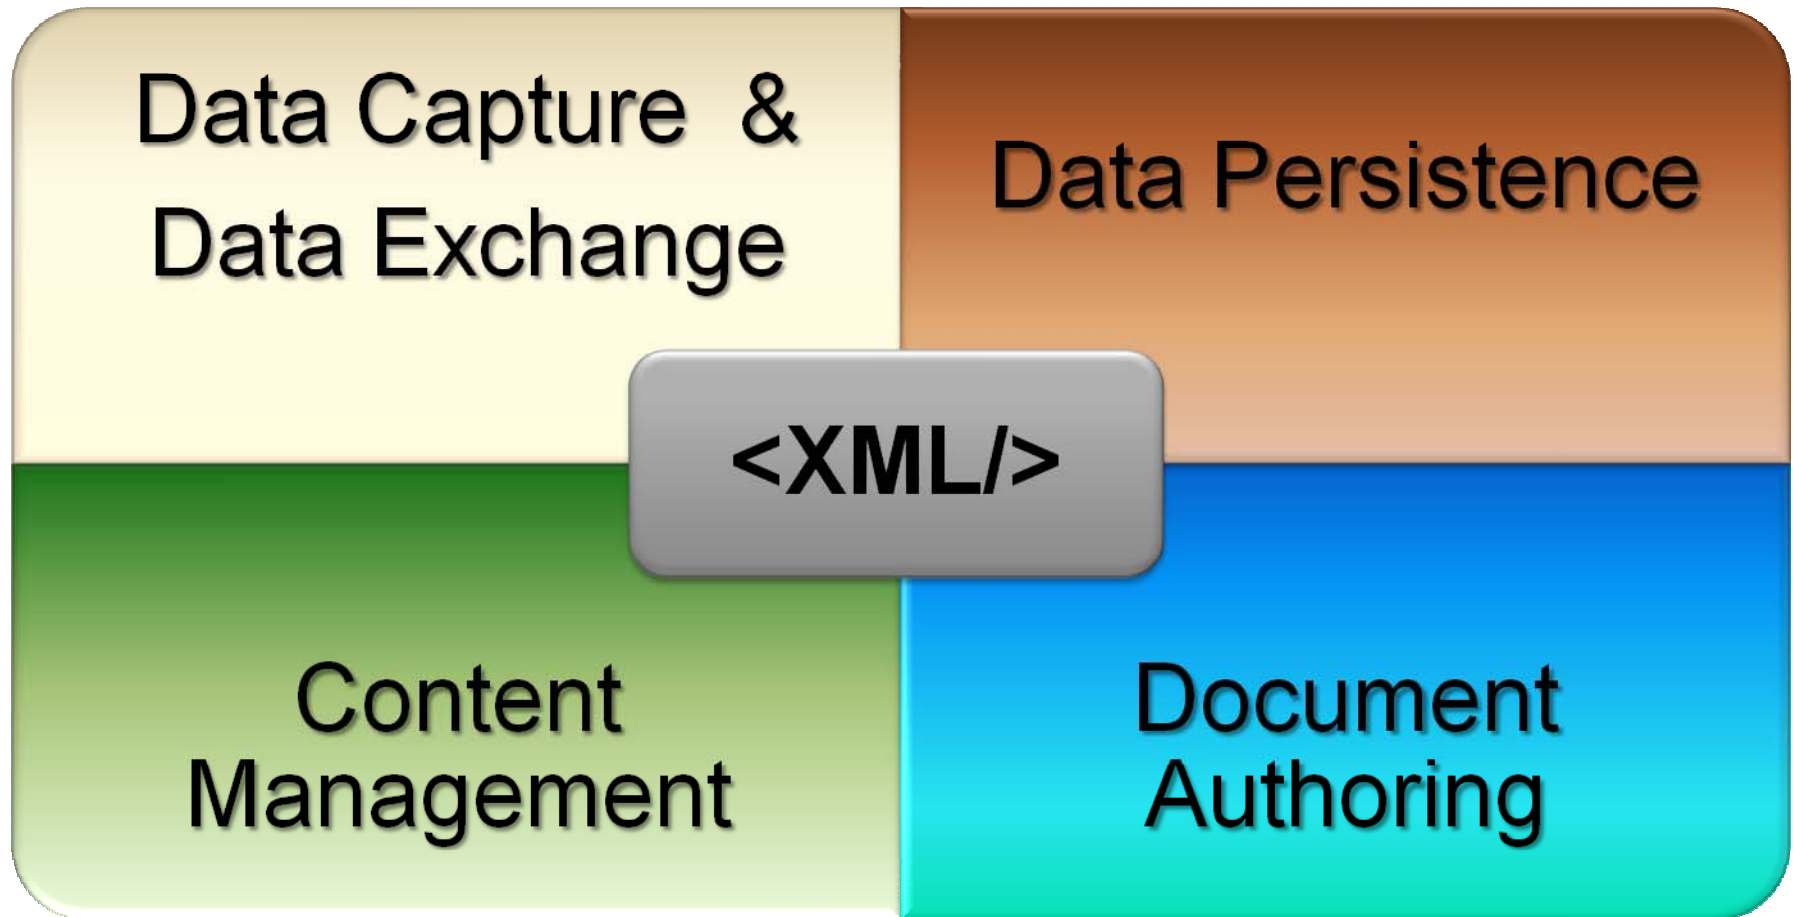


Cube Organized Materialized Views Transparent to SQL Queries





XML Use cases

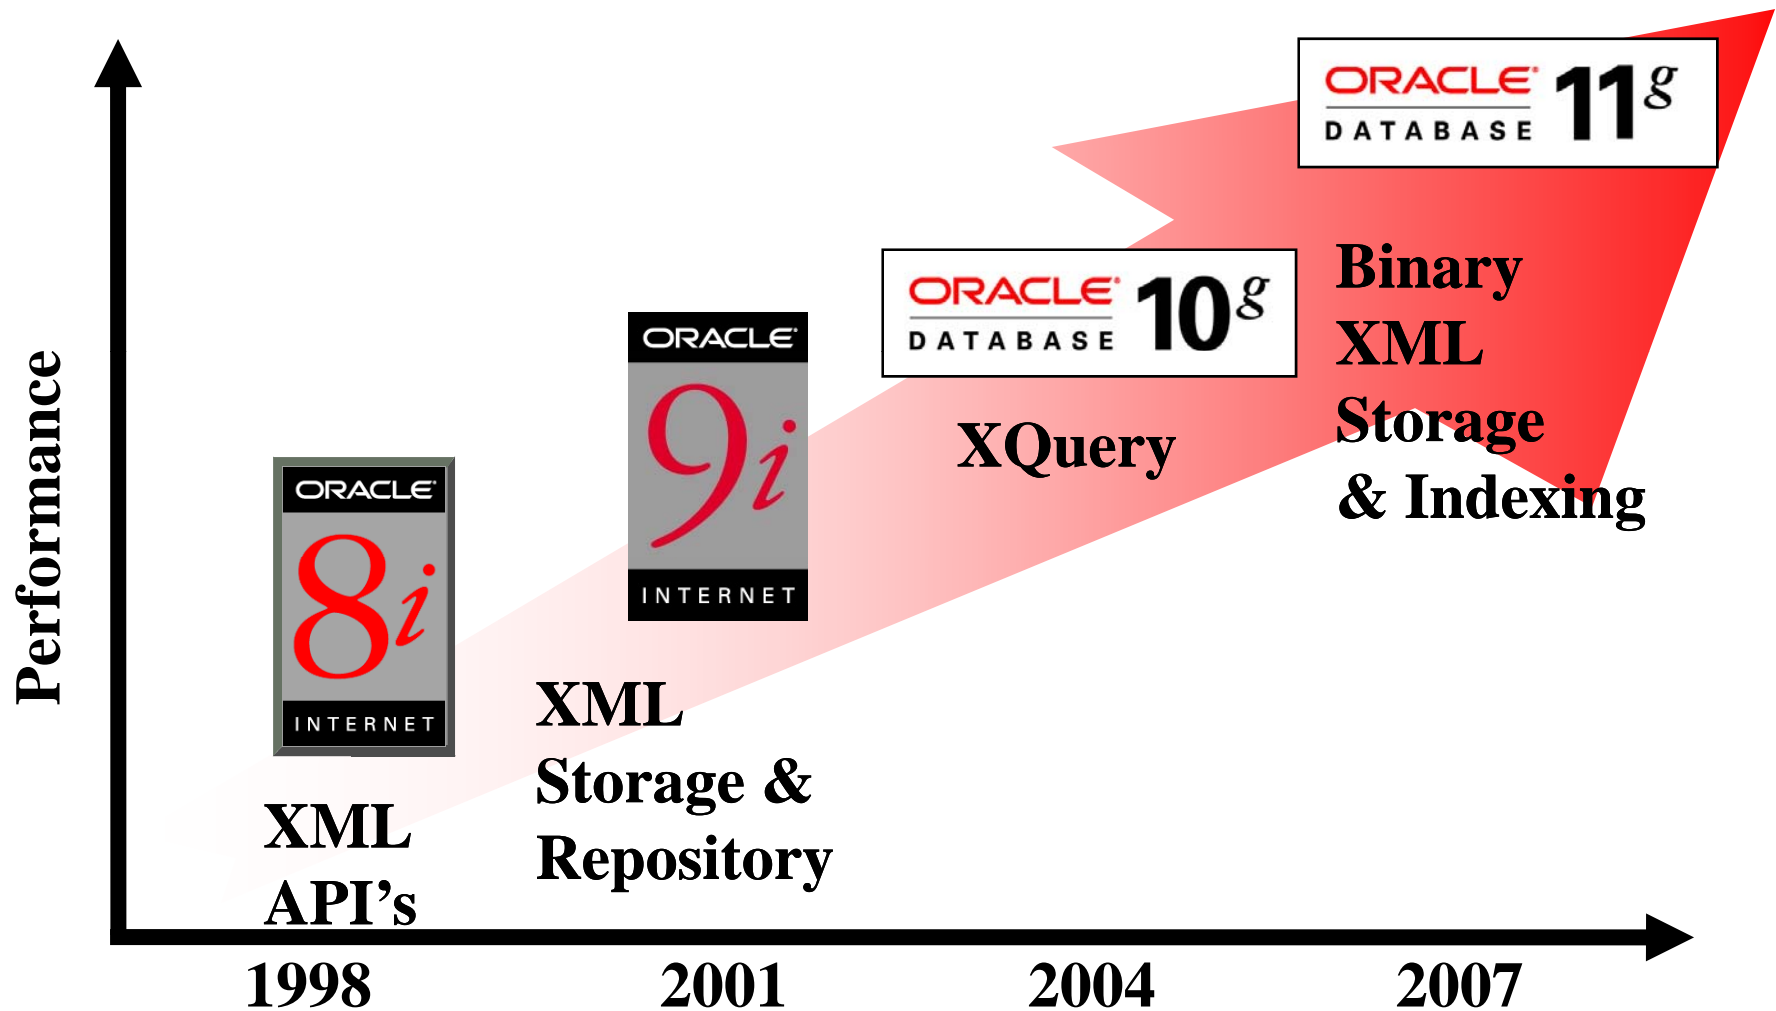




Oracle's XML Vision

- Enable a single source of truth for XML
- Provide the best platform for managing all your XML
 - Flexibility to allow optimal processing of data-centric and content-centric XML
 - Deliver Oracle's commitment to Reliability, Security, Availability and Scalability
- Drive and implement key XML Standards
- Support SQL-centric, XML-centric and document-centric development paradigms

Oracle and XML: Sustained Innovation





Why XML in the Database?

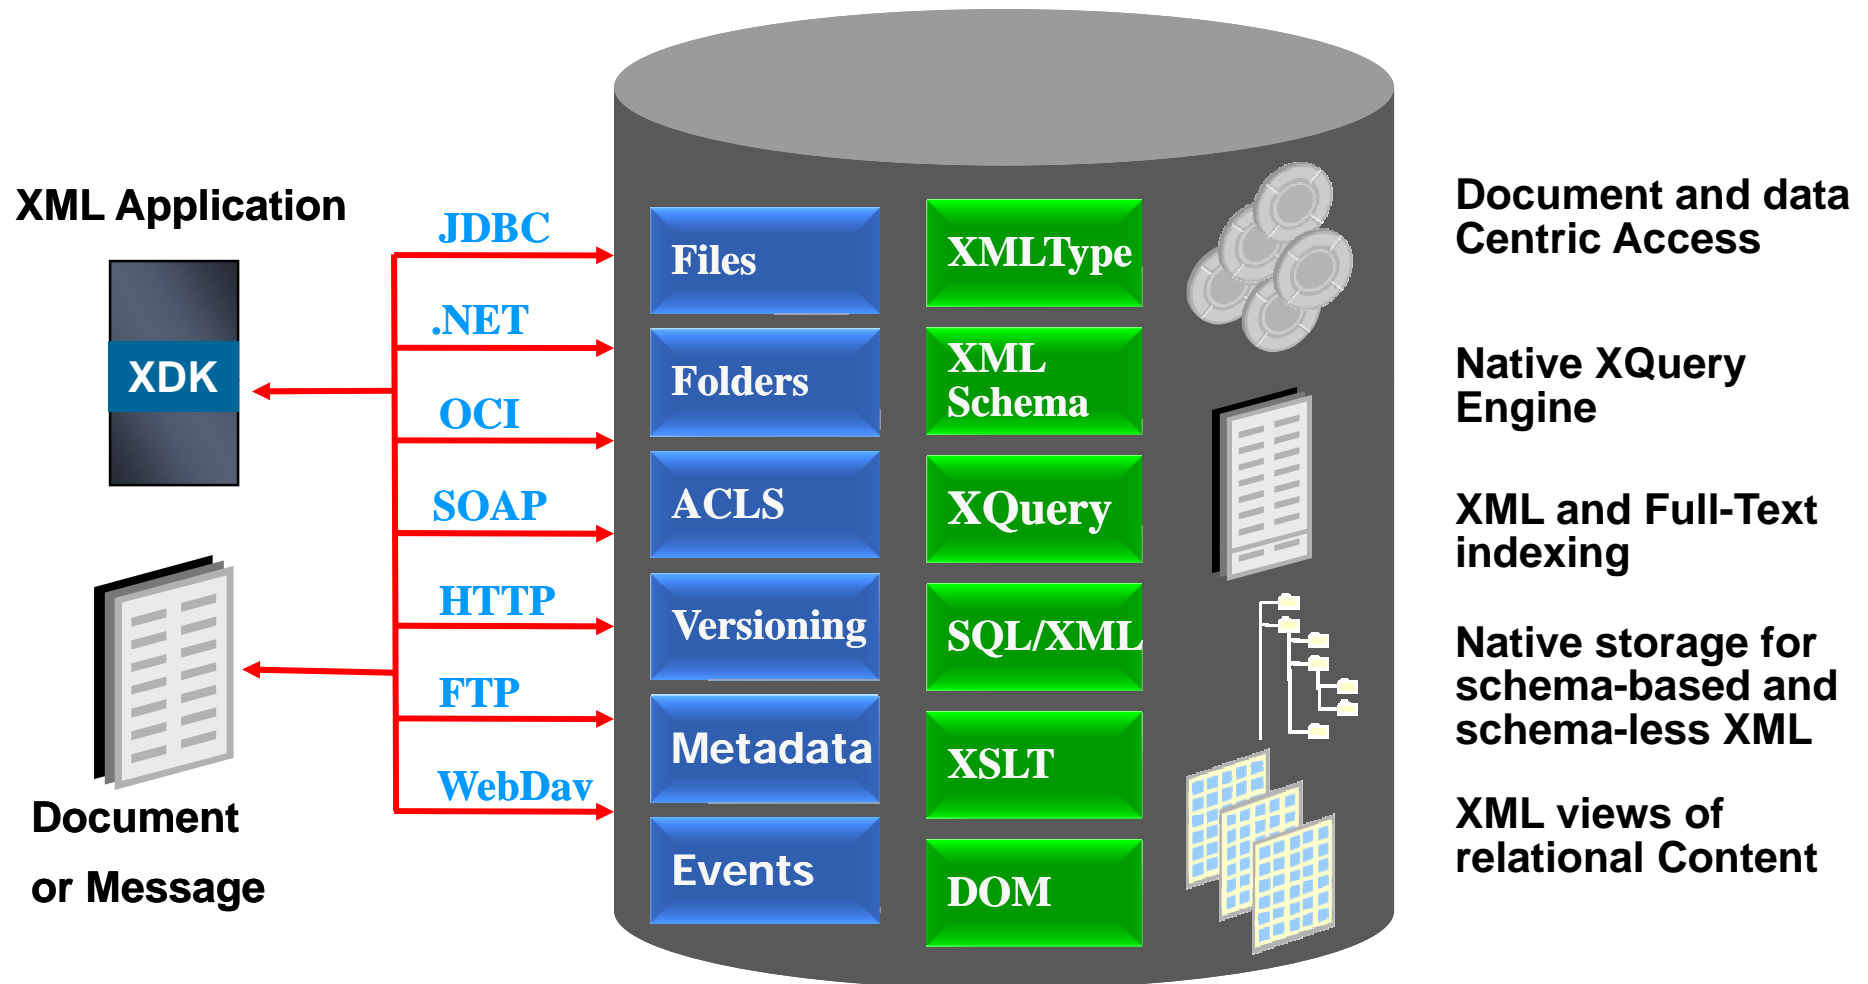
- XML contains mission critical information
 - Interchange with external organizations
 - Web Services
- Need to manage XML effectively and efficiently
 - Number and size of documents increasing
 - Reliability, Scalability, Availability
 - Security
 - Compliance
- Accurate and fast information location and retrieval

What is XML DB?

- XMLType: Native flexible, scalable, XML storage, processing and indexing
- XQuery, XML Schema, XSLT: enable XML centric development
- XMLDB repository: XML specific content management
 - File / Folder access metaphors
 - Xlink for integrity management
 - Xinclude for compound documents
- SQL/XML for SQL-Centric XML publish and SQL/Xquery interaction
- SQL/XML and XML/SQL interoperability
- Standards compliant



Advanced XML Capabilities





Demo





Why Java in the Database?

- Implement things you cannot do in PL/SQL
- Reuse Java classes and libraries to extend database
 - Reduce development time and costs
 - Extend database functionality
 - e.g., custom image transformation and data format conversion (GIF, PNG, JPEG)
- Partition Java applications between middle-tier and database
 - Compute bound applications in client/middle-tier
 - Data, data-and-compute intensive applications in the database
- Reuse Java skills for database development



Java in the Database: Use Cases

- Trigger-based Notification System using RMI
- Secure Credit-Card Processing using JSSE
- Custom Alert applications that monitor business data
- Sending emails with attachment from within the database
- Produce PDF files from Result Set
- Execute external OS commands and external procedures
- Implement Md5 CRC
- Publish Repository Content to Portal
- Portable Logistic Applications
- Implement Parsers for various File Formats (txt, zip, xml, binary)
- Implement Image Transformation and Format Conversion (GIF, PNG, JPEG, etc)
- Implement database-resident Content Management System
- HTTP Call-Out
- JDBC Call-Out
- RMI Call-Out to SAP
- Web Services Call-Out
- Messaging across Tiers
- RESTful Database Web Services*
- Lucene Domain Index*




What's New for Java in Oracle Database 11g


- Standard compliance:
 - IPv6 Support via System properties `java.net.preferIPv4Stack`, `java.net.preferIPv6Addresses`
- Standards Compliance: JDK 1.5 Support
- Performance: Innovative JIT Compiler
 - Transparent, On-the-fly compilation, no C compiler required
 - Compiled Java codes persisted for future invocations
- Ease of Deployment: JMX Instrumentation for Monitoring
 - Management MBean, standard JMX tools (e.g., JConsole)
- Ease of Development: JDK Command-line like Interface

Agenda

- Database
 - SQL, PL/SQL, Secure Files
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - In-Memory Database Cache
 - Partitioning
- High Availability Architecture
- Enterprise Manager Packs
- SQL Developer Data Modeler
- SQL Developer




Server SQL Results Cache



- Database caches results of queries, sub-queries, or PL/SQL function calls
 - Cache is shared across statements and sessions on server
 - Full consistency and proper semantics
- 2x speedup on hit for worst case of trivial query
- 100x speedup on hit for complex queries
- Statement hints specify caching - `/*+ result_cache +*/`
- Only for very read intensive tables

OCI Consistent Client Cache



Consistent Caching

Simplest queries can speedup:

- 50x in elapsed time
- 20x in CPU time

- Caches query results on client
- Primarily for caching small (10s or 100s of KB) read-intensive tables
 - Queries where network overhead dominates
 - e.g. lookup tables
- Cache is fully consistent
 - Coherence messages bundled into responses to DB calls ensure cache remain consistent
 - Like Cache Fusion extended out to clients

Client Query Result Cache

- Database Configuration (init.ora)

```
client_result_cache_size=200M
```

```
client_result_cache_lag=5000
```

- Client Configuration (sqlnet.ora)

```
OCI_QUERY_CACHE_SIZE=200M
```

```
OCI_QUERY_CACHE_MAXROWS=20
```

- Oracle Database 11g Release 2


- Table Annotation: transparent, no code change

```
alter table emp result_cache;
```

- Oracle Database 11g Release 1

- Query Annotation: hint for Caching the Result Set

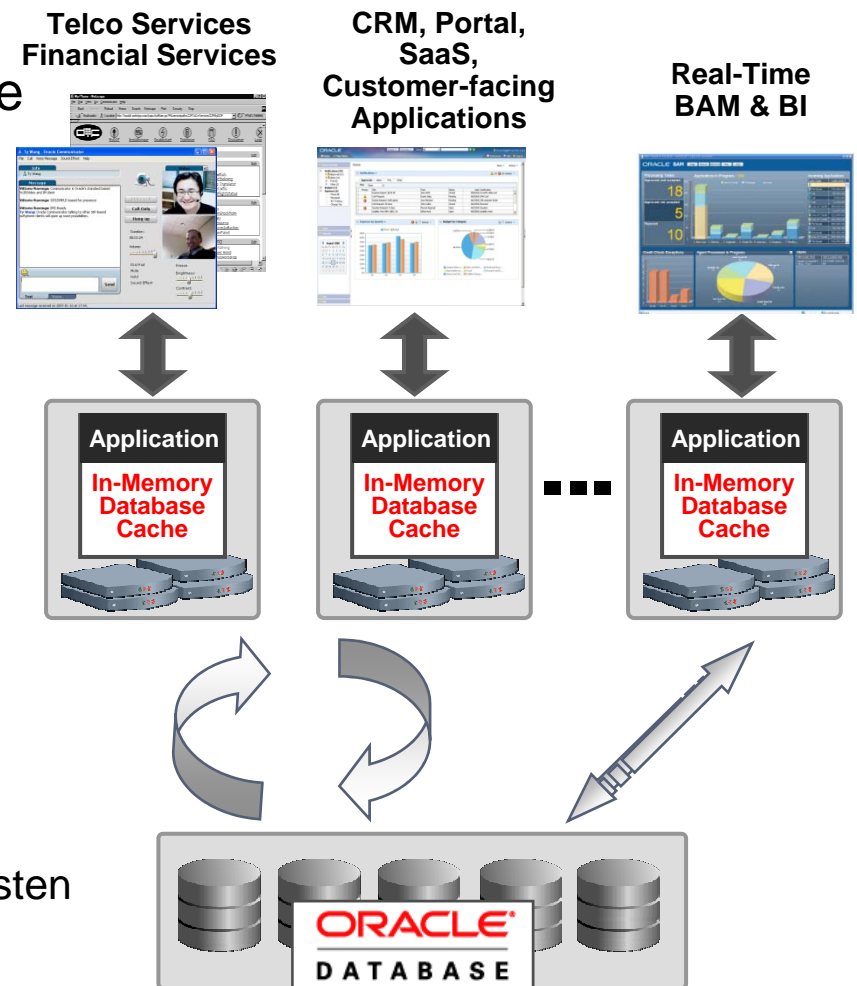
```
select /*+ result_cache */ * from employees
```



Oracle In-Memory Database Cache

Accelerator for Oracle Database Applications

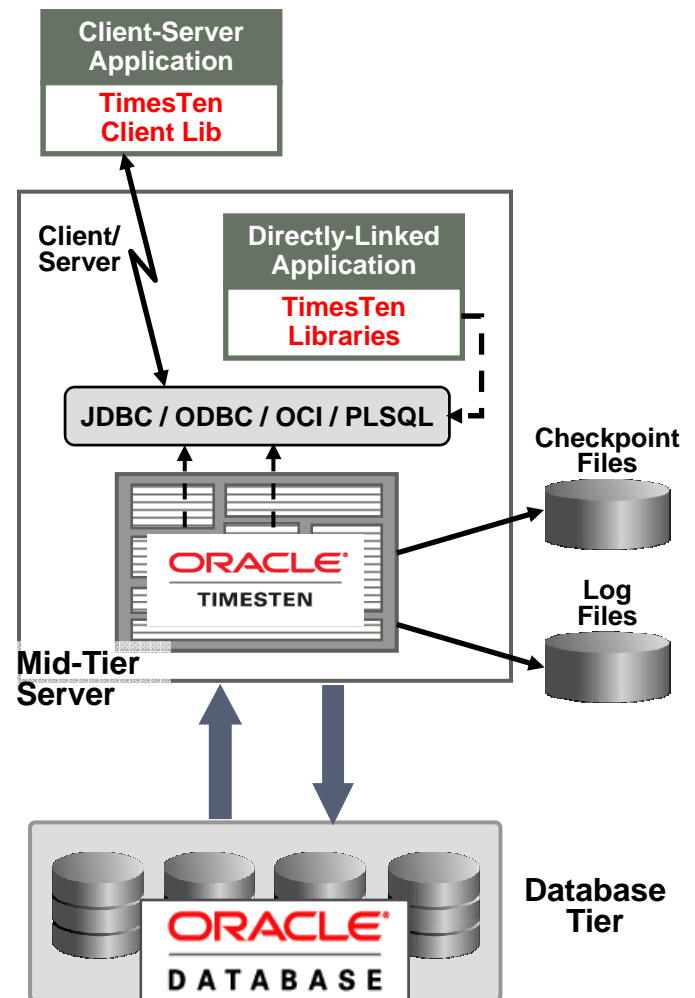
- IMDB Cache is an Option for the Oracle Database Enterprise Edition
- Deployed in the application tier, as a real-time, updatable cache for the Oracle Database
- Reduced response time and increased throughput for Oracle Database applications
- Available for download on OTN
<http://www.oracle.com/technology/products/timesten>



Oracle In-Memory Database Cache

Optimize In-Memory Database Technology

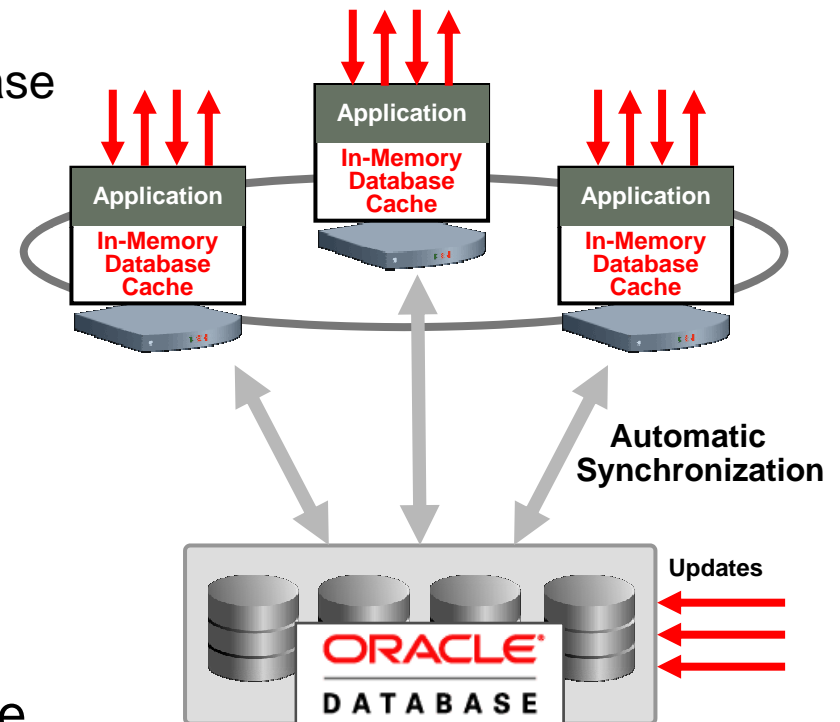
- Built using Oracle TimesTen as a relational cache database
 - Entire database in memory
 - Standard SQL with JDBC, ODBC, OCI, Pro*C, PL/SQL
 - Persistent and durable
 - Fast, consistent response times
 - Cross-tier High Availability
- Cache Oracle Database tables in the application-tier
 - Groups of related tables
 - All or subset of rows and columns



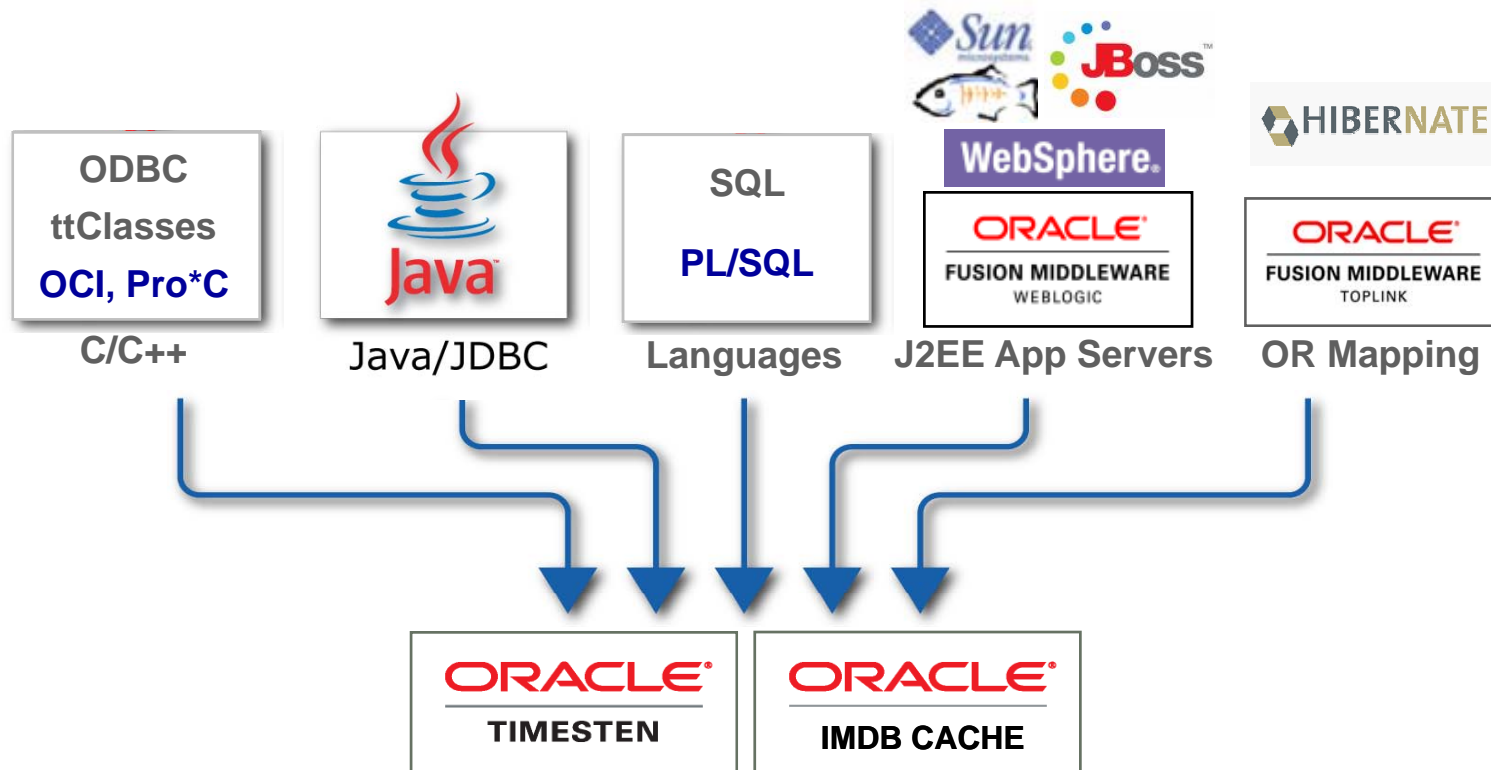
Oracle In-Memory Database Cache

Low-Cost Extreme Performance in the Middle-tier

- Scalable cache grid
- Read-only and updatable caches
 - Access cache tables like regular database tables
 - Joins/search, insert/update/delete
 - Flexible caching options
 - Pre-load or dynamic load data
 - Local cache groups for consistent response time
 - Global cache groups for location transparency and data sharing across all cache nodes
- Automatic data synchronization with the Oracle database



Oracle In Memory Database Cache Application Development




New in 11g release

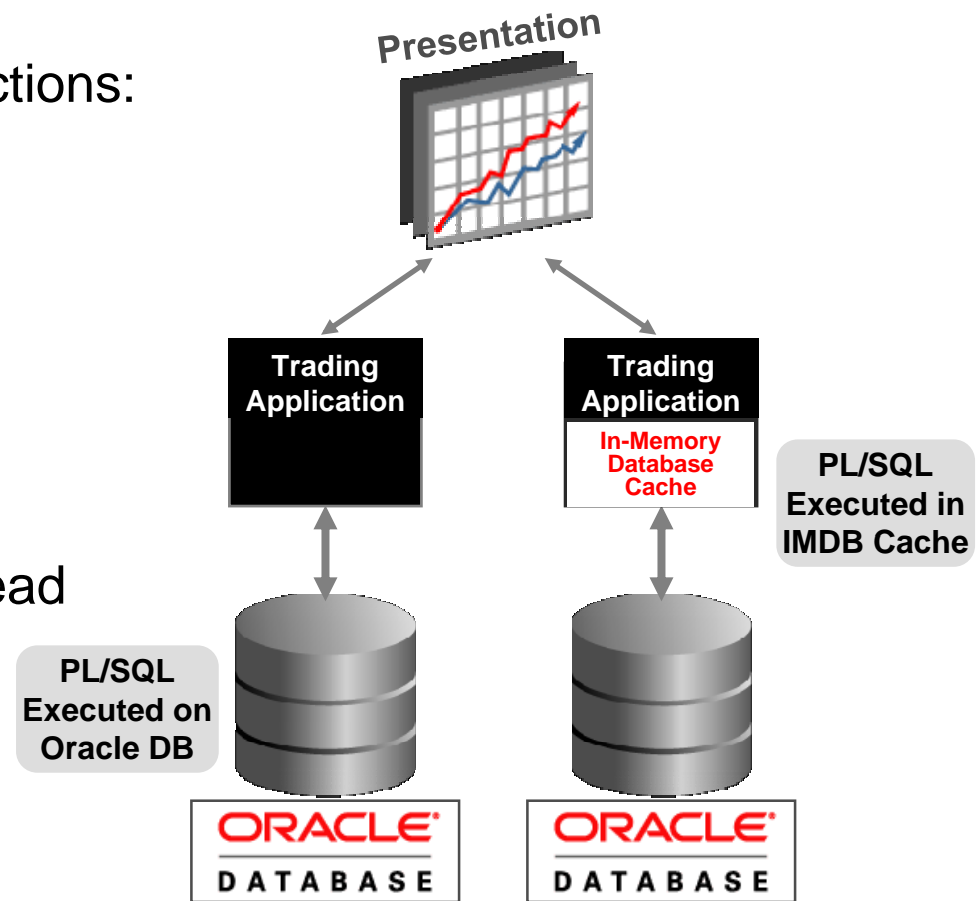
- PL/SQL, Oracle Call Interface (OCI) and Pro*C Support
Planned for CY2010
- ODP.NET data provider, PHP

Order Matching Application Demo

- Three different types of transactions:
 - Place market order
 - Place limit order
 - Process quote
- Business logic implemented in PL/SQL stored procedures
- Application written in Java
- Execute 1000 times in one thread
 - Place an order
 - Process a quote

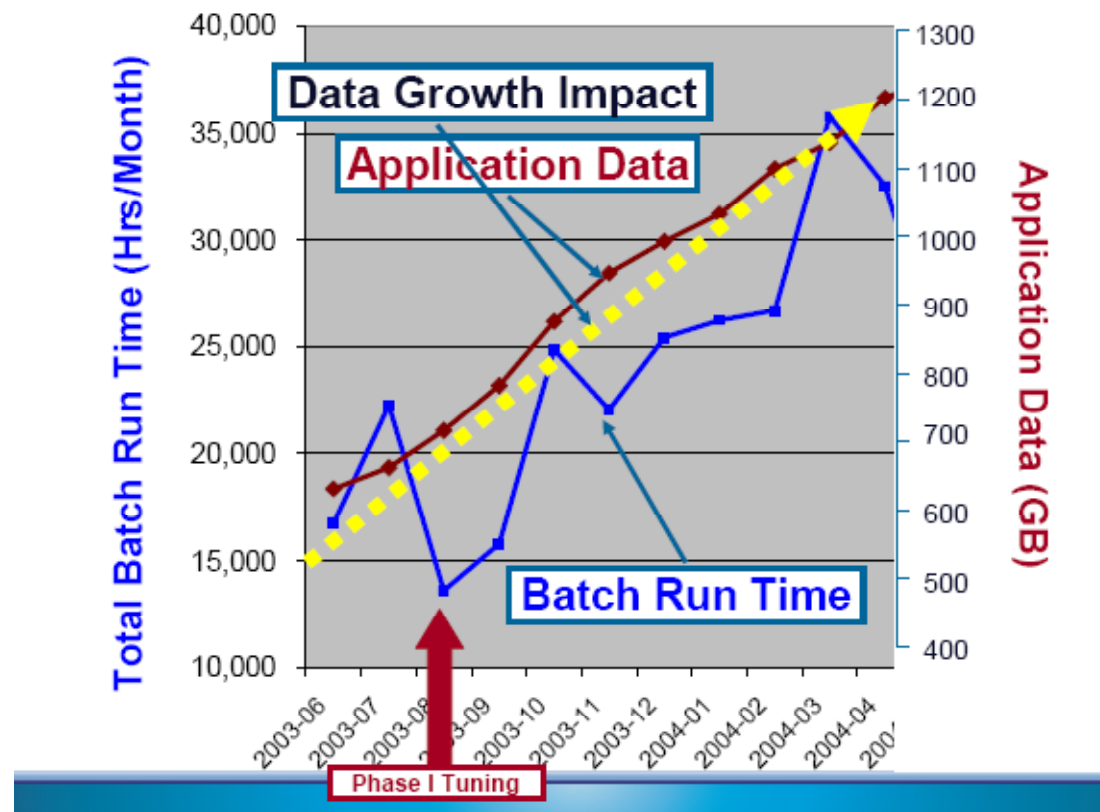


Start Demo



Data Growth Impacts Performance

Real World: Data Growth vs. Performance




ORACLE

Advanced OLTP Compression

Significantly Reduce Storage Usage

- Compress large application tables
 - Transaction processing, data warehousing
- Compress all data types
 - Structured and unstructured data types
- Improve query performance
 - Cascade storage savings throughout data center

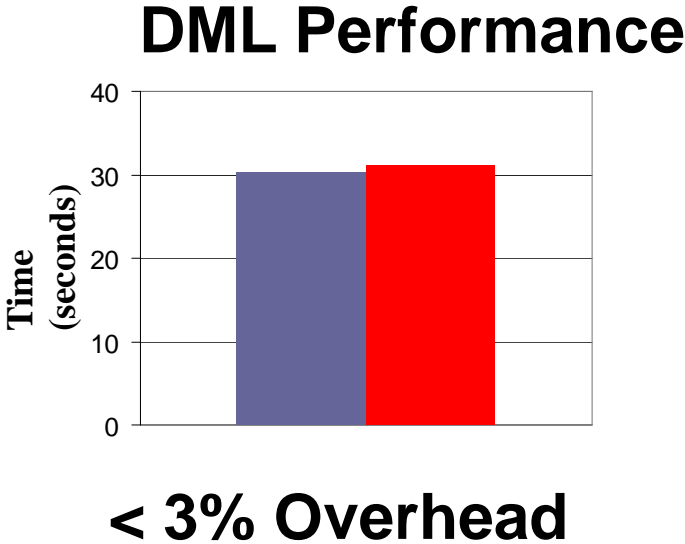
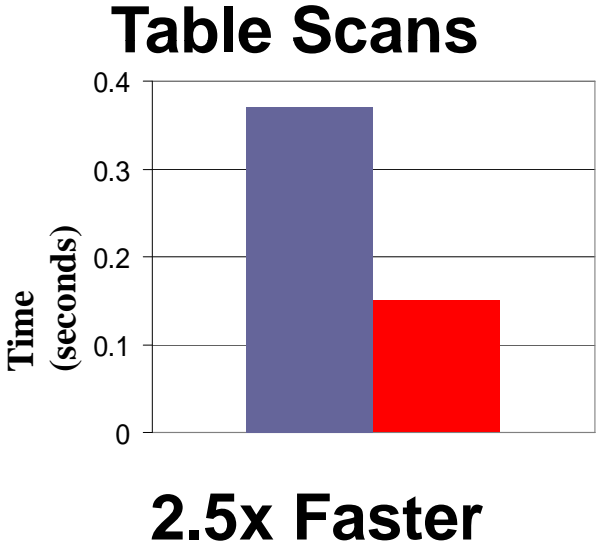



Up To
4X
Compression



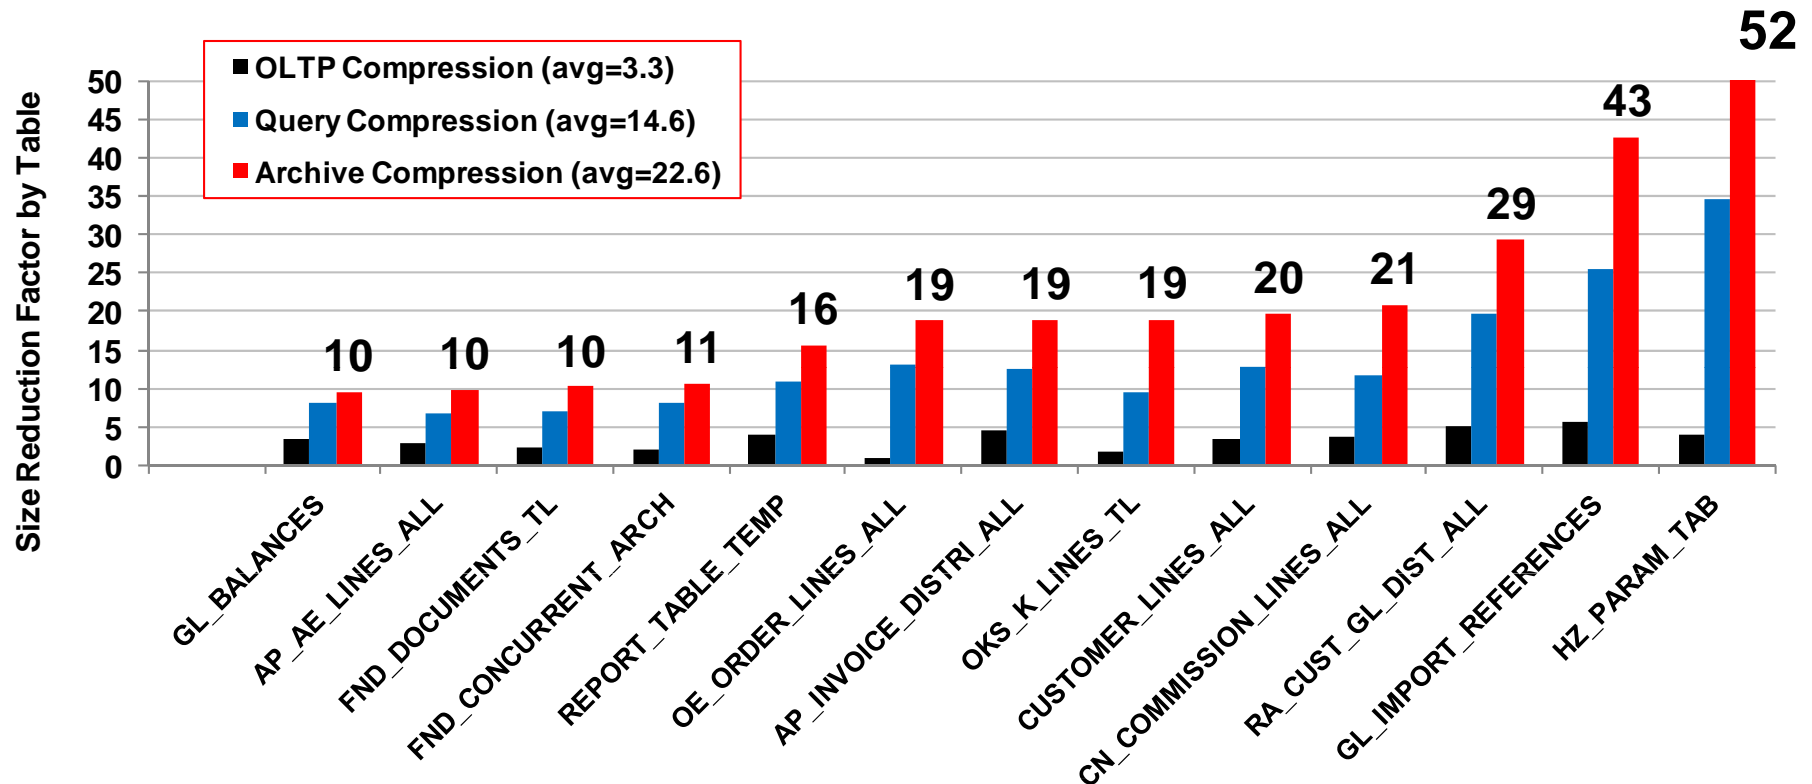
Real World Compression Results

10 Largest ERP Database Tables



Real-World Compression Ratios

Oracle Production E-Business Suite Tables



- Columnar compression ratios
 - Query = 14.6X
 - Archive = 22.6X
 - Vary by application and table


Partitioning

Divide and conquer large tables and indexes



Large Table


Difficult to Manage



Partition

Easier to Manage

Improve Performance




Composite Partition


Higher Performance

Match business needs

Reduce the impact of Data Growth
Partition for performance, management and cost



95% Less Active **5% Active**




Low End Storage Tier
2-3x less per terabyte



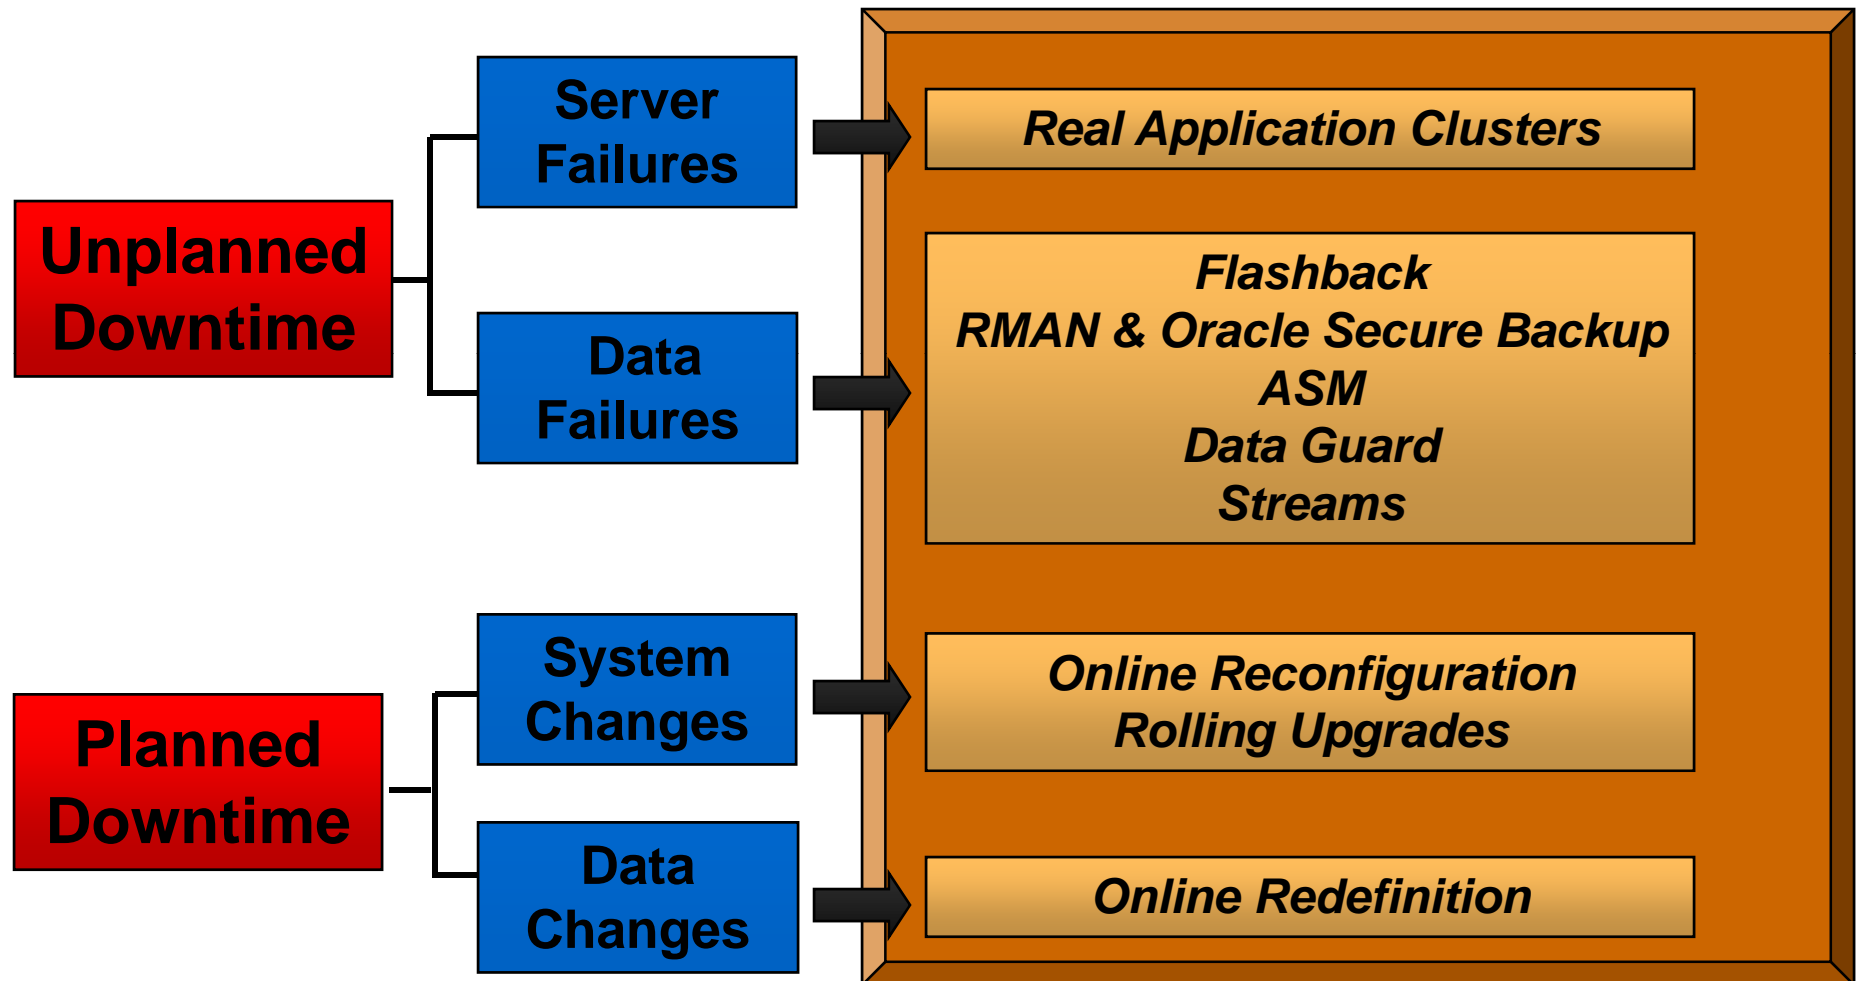
High End Storage Tier

Agenda

- Database
 - SQL, PL/SQL, Secure Files
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - In-Memory Database Cache
 - Partitioning
- **High Availability Architecture**
- Enterprise Manager Packs
- SQL Developer Data Modeler
- SQL Developer




Oracle's Database HA Solution Set



Flashback Technologies

Error Detection & Correction

- Flashback revolutionizes error recovery
 - View 'good' data as of a past point-in-time
 - Simply rewind data changes
 - Time to correct error equals time to make error







- Low impact
- Excellent tool for configuring QA, Dev and Training databases
- Flashback is easy – simple commands, no complex procedure, e.g.

Flashback Query: `select * from Salary AS OF '12:00 P.M.' where ...`

Flashback Database: `FLASHBACK DATABASE TO TIMESTAMP TO_TIMESTAMP`

`('12-10-2008 10:00:00', 'DD-MM-YYYY HH24:MI:SS');`

Recovery at the Level of Business Objects – not Bits & Bytes



Continuous Data Protection (CDP) Built Within the Database

Edition-based Redefinition

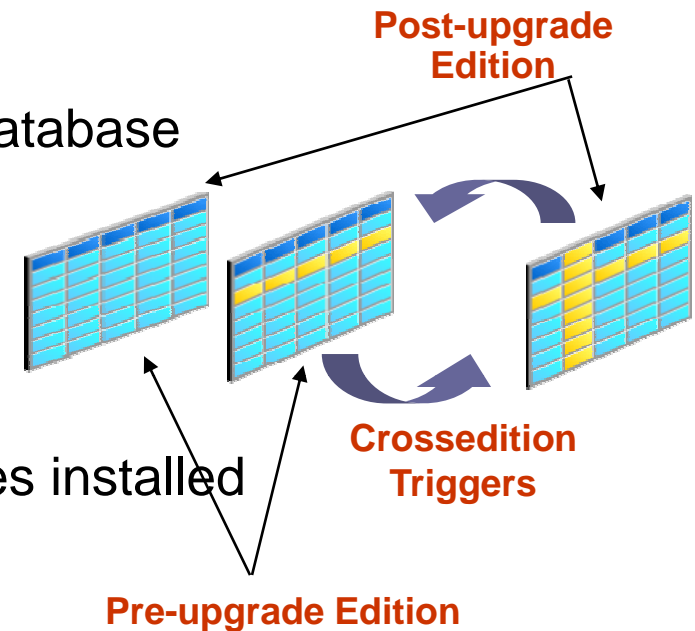
Overview

- Enables online application patches and upgrades
- Allows old and new version of application to co-exist even though schema is changed by new version
- Capabilities primarily used by application developers

Edition-based Redefinition


How Does it Work?

- Maintains logical versions of changed database objects, through:
 - Edition
 - Editioning View
 - Crossedition Trigger
- PL/SQL code changes and view changes installed in the privacy of a new edition
- New data changes made to new columns/tables not seen by old edition
- Editioning view exposes a private projection of a table into each edition
- Crossedition trigger propagates changes made by old edition into new edition's columns, or vice-versa



Agenda

- Database
 - SQL, PL/SQL, Secure Files.
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - In-Memory Database Cache
 - Partitioning
- High Availability Architecture
- **Enterprise Manager Packs**
- SQL Developer Data Modeler
- SQL Developer



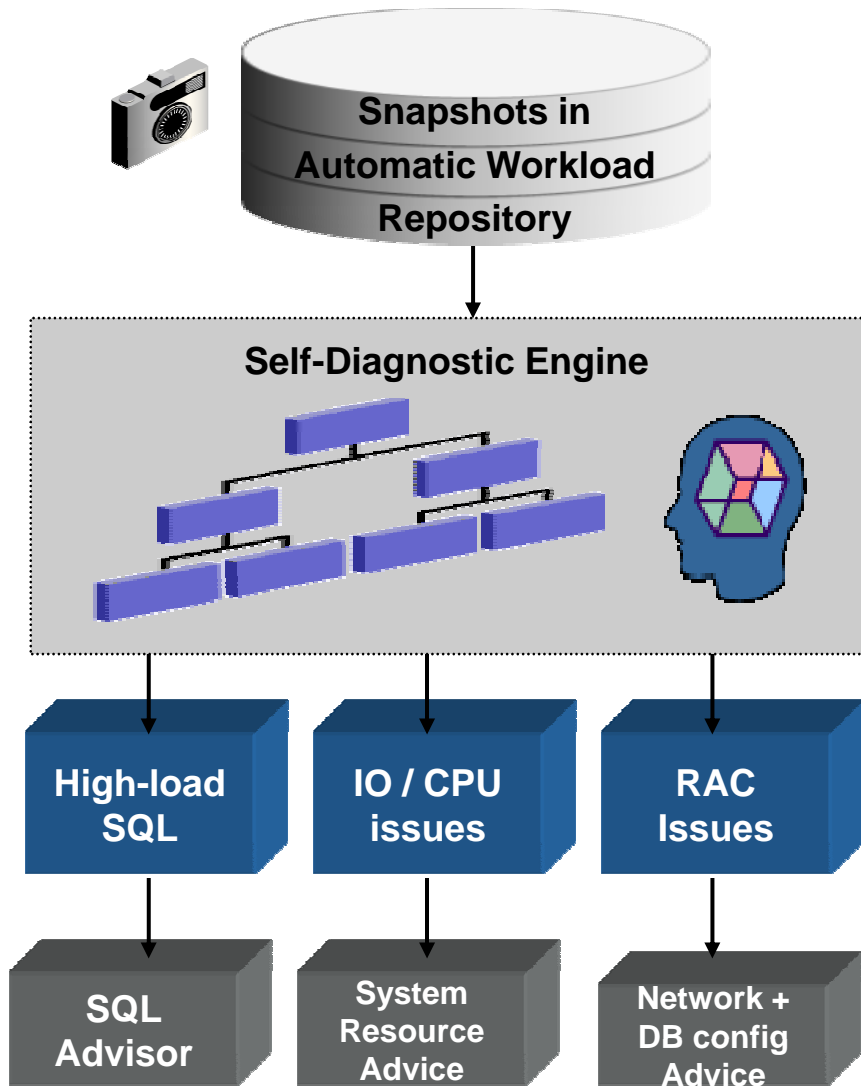


Enterprise Manager Packs

- **Diagnostics Pack**
 - Automatic Database Diagnostic Monitor (ADDM)
 - Automatic Workload Repository (AWR)
 - Performance Monitoring
 - Active Session History (ASH)
- **Tuning Pack**
 - SQL Tuning Advisor
 - SQL Profiles
 - SQL Access Advisor
 - SQL Tuning Sets
 - Real Time SQL Monitoring

Automatic Performance Diagnostics

AWR and ADDM



- Top Down Analysis Using AWR Snapshots
- Classification Tree - based on decades of Oracle performance tuning expertise
- Performance expert, now a RAC Specialist too in 11g
- Real-time results
 - Don't need to wait hours to see the results
- Pinpoints root cause
 - Distinguishes symptoms from the root cause
- Reports non-problem areas, e.g., I/O is not a problem

Real-Time SQL Monitoring


Looking inside SQL Execution

- Automatically monitors long running SQL
- Enabled out-of-the-box with no performance impact
- Monitors each SQL execution
- Exposes monitoring statistics
 - Global execution level
 - Plan operation level
 - Parallel Execution level
- Guides tuning efforts

Status	Duration (s)	SQL ID	SQL Text	Start	Session	Parallel	Database Time (s)	IO
✓	362	8fu14h4ndh7gb	select count(*) fro...	10:16:02 PM	1720		324	56K
✗	61	7m52x97c6td76	select count(*) fro...	10:09:49 PM	1720		60	11K
✗	70	7m52x97c6td76	select count(*) fro...	10:08:28 PM	1720		68	11K
✓	1466	7m52x97c6td76	select count(*) fro...	05:56:16 PM	1658		1197	115K
✗	21	0k3n92v33vbav	select c.customer_...	Wed Oct 24 2007 10	1645		19	
✗	530	bhyfcbx5wss4	select count(*) fro...	Wed Oct 24 2007 10	1539		1054	155K
✓	59	9x8aupdv8qa5z	SELECT gl_dis...	Wed Oct 24 2007 10	1670		100	22K
✓	401	9st19c50cf9pc	select nvl(sum(de...	Wed Oct 24 2007 10	1645		401	10K
✗	82	9st19c50cf9pc	select nvl(sum(de...	Wed Oct 24 2007 10	1595		81	3706
✓	344	cwt8t2fvasz4q	select count(*) fro...	Wed Oct 24 2007 10	1647		347	2954
✓	329	9x8aupdv8qa5z	SELECT gl_dis...	Wed Oct 24 2007 10	1670		971	25K
✗	36	4gqvwff3uzpnq	select count(*) fro...	Wed Oct 24 2007 10	1689		71	17K
✓	283	1fr1jw7gwr1q7	select count(*) fro...	Wed Oct 24 2007 10	1595		554	90K
✓	219	cgh6u1gp92u7r	select count(*) fro...	Wed Oct 24 2007 10	1689		222	4025
✓	97	fasmjk1d8r8sp	SELECT dtp.tab...	Wed Oct 24 2007 09	1607		96	8892
✓	268	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 09	1673		266	60K
✓	267	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 09	1539		265	59K
✓	266	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 09	1550		263	59K
✓	366	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1539		363	64K
✓	394	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1565		392	65K
✓	312	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1626		309	57K
✓	405	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1689		404	64K
✓	372	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1621		370	64K
✓	364	498g203ms1zj6	SELECT a.customer...	Wed Oct 24 2007 08	1539		363	64K

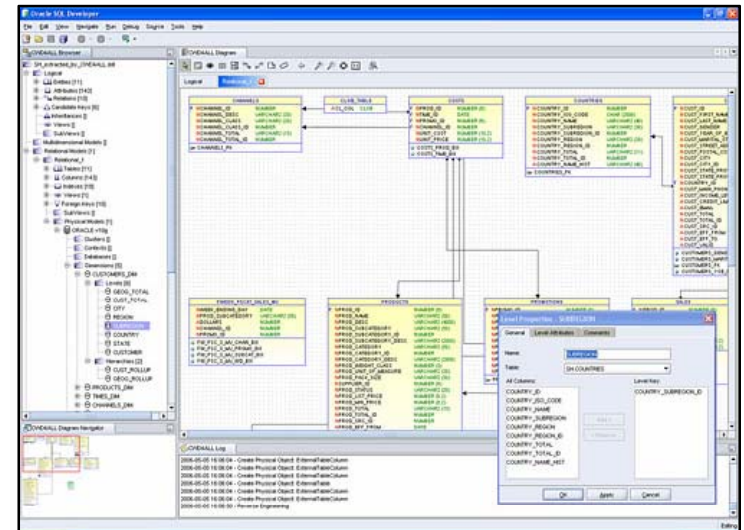
Agenda

- Database
 - SQL, PL/SQL, Secure Files
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - In-Memory Database Cache
 - Partitioning
- High Availability Architecture
 - RAC
 - RAC One-Node
- Enterprise Manager Packs
- **SQL Developer Data Modeler**
- SQL Developer



What is SQL Developer Data Modeler?

- A diagramming and data modeling tool
- A single tool for different users and functionality
 - Data Architect builds logical data models
 - Database Developer models relational models (tables and columns)
 - DBA adds tablespaces, partitions
- Use data models to
 - Verify accuracy and completeness of data requirements and business rules with customers
 - Build standards-driven DDL scripts






Oracle SQL Developer Data Modeler

- Supports a variety of visual models including
 - Entity Relationship (ERD)
 - Relational
 - DataTypes (SQL99)
 - Multi-Dimensional
- Records details for physical implementation
- Imports from various sources
 - DDL and Dictionary import from Oracle and non-Oracle databases
 - Oracle Designer repository
 - Ca Erwin Data Modeler
- Exports to various sources
 - DDL for Oracle and non-Oracle databases


Agenda

- Database
 - SQL, PL/SQL, Secure Files
 - Analytics
 - Secure Files
 - XML
- Performance By Design
 - Query Cache
 - Partitioning
- High Availability Architecture
 - RAC
 - RAC One-Node
- Enterprise Manager Packs
- In-Memory Database Cache
- SQL Developer Data Modeler
- SQL Developer



Oracle SQL Developer

Targeting the Database Developer



- Lightweight, graphical interface
 - Enhanced database development productivity
 - Browsing, creating, editing, debugging and reporting
 - Integrated migration with third-party databases
- Easy installation
 - Download and unzip
 - Uses thin JDBC driver => No Oracle Home required
- Free and fully supported
- Adoption
 - Over 1 million downloads
 - Distributed with Oracle Database 11g
 - 1.5 million users
- Provides framework for extensions



SQL Developer 2.1 – Features

- PL/SQL Unit Testing
- Data Modeler Viewer
- Migration support for IBM DB2 UDB and Teradata
- Updated Data Grids
 - Manage columns, filter on data
- New SQL Worksheet
 - Dockable dbms_output, multiple worksheets
- Increased Connections navigator support for
 - Jobs, Editions (for 11gR2), XML DB Repository
- Updated display editors
 - PL/SQL edit mode, subpartitions
- Version Control Support for Serena Dimensions, Perforce
- Updated filtering mechanism
 - Schema level, generated objects




Demo



Database Recap

Supporting the new normal with a thousand cores and petabytes of data requires new algorithms




• Scalable Execution

- Scale to 1000 Cores systems
 - Parallelism, locks, connections
- Transparent caching in middle tier

• Scalable Availability

- Scale Backup and recovery to Petabyte databases
- Optimized cache fusion algorithms
- Use Standby to isolate faults and improve scaling



• Scalable Storage

- Disks get bigger but not faster
- Compression to speed data access
- Faster LOBs
- More sophisticated data partitioning
- Exadata Storage Server



• Scalable Management

- Manage larger systems
- Realistic testing at scale
- Manage ultra-complex applications
 - e.g. Oracle Fusion Apps





For More Information

<http://search.oracle.com>

or

www.oracle.com/database