

OGh

ORACLE®

Oracle gebruikersclub Holland
onderdeel van  **ORACLE BENELUX**
USER GROUP

Welkom bij de themasessie Oracle Spatial 11g

OGh SIM Commissie



Agenda

- 19:00 - 19:05 Welkom
- 19:05 - 19:35 Managing ALL your Information in 11g
(Han Wammes)
- 19:35 - 20:15 Oracle Spatial 11g (Han Wammes)
- 20:15 - 20:30 Pauze
- 20:30 - 21:00 OBI EE en Oracle MapViewer (Bas
Roelands)
- 21:00 - 21:30 OGH SIM Commissie (Bram
Schiltmans)



ORACLE®

Managing ALL your Information

Han Wammes
Solution Architect
Sales Consulting Technology

Trends

- Enterprise applications incorporating content with data
- Compliance driving content into database
- Documents becoming semi-structured information, driving value of database
- XML becomes the standard to represent / describe Information
- Information Retrieval (Search) is essential to Information Access
- Information Integration critical to locate and correlate data from multiple sources, structured & unstructured



Challenges Of management of Unstructured Data


- Increases Cost of Management
- Impacts Productivity
- Complicates Compliance
- Increases Risk



Strategy

- Evolve Database to manage all Enterprise Information
- Meld Database and File-system metaphors
- Enable Integration of all Enterprise Information Sources
- Enable Rich Information Retrieval capabilities
- Provide solutions built on top of the proven data management platform - database

Evolving the Database....



Complete Information Management

Secure Enterprise Search
Enterprise Wide Search

Content DB and Records DB
Content Management in the Database

XML
Integrated Native XML Database

Text
Text Indexing and Classification

Location & Spatial
Location Enabled Databases

Multimedia
Audio, Image and Video

Relational
Characters, Numbers, Dates, LOBs

ORACLE
D A T A B A S E

10^g


ORACLE

Complete Information Management



- Leverage the Oracle database to manage all Enterprise Content
- Secure structured, semi-structured & unstructured data in the database

Secure Enterprise Search <i>Complete Set of Connectors</i>
Content DB and Records DB <i>Content Management in the Database</i>
XML <i>Binary XML Support</i>
Text <i>RDF and OWL Support</i>
Location & Spatial <i>3D Support</i>
Multimedia <i>Medical Image Support (DICOM)</i>
Relational <i>New SecureFiles LOB Storage</i>



Complete Information Management Relational--LOBs





Files Belong with Relational Data

- Many applications have both files and relational data
 - e.g. Document Management, Medical, CAD, Imaging
- This split compromises security, robustness, and management
 - Disjoint security and auditing models
 - Changes cannot be made atomically
 - Backup and recovery are fragmented
 - Search across relational data and files is difficult
 - Space management is complicated
 - Separate interfaces and protocols
- Two data managers for one application is one too many



SecureFiles

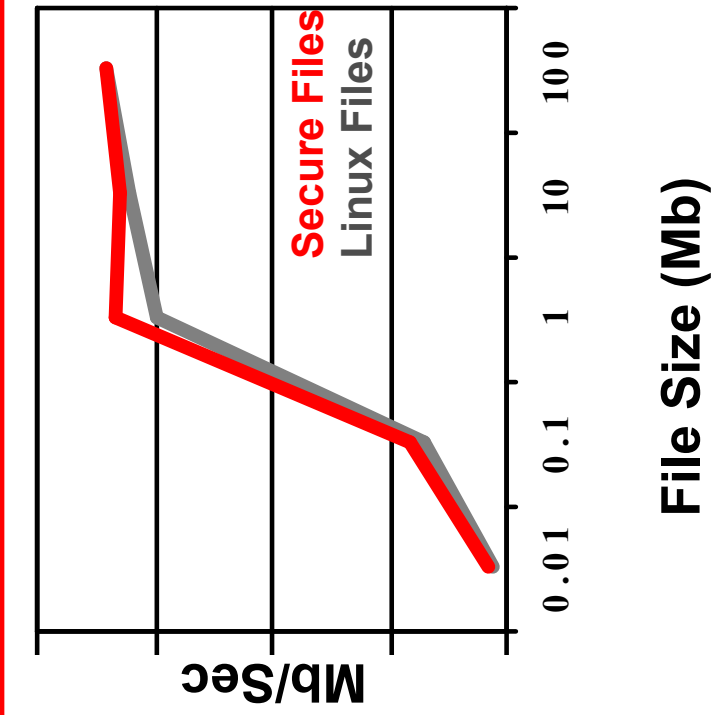
Consolidated Secure Management of Data

- Many applications have both flat files & relational data
 - E.g. Document Management, Medical, CAD, Imaging
- **SecureFiles** is new database feature designed to break the performance barrier that has kept file data out of database
- Similar to LOBs but much faster, and with more capabilities
 - Transparent encryption, compression, deduplication, etc
 - Preserves security, reliability, and scalability of database
 - Superset of LOB interfaces for easy migration from LOBs

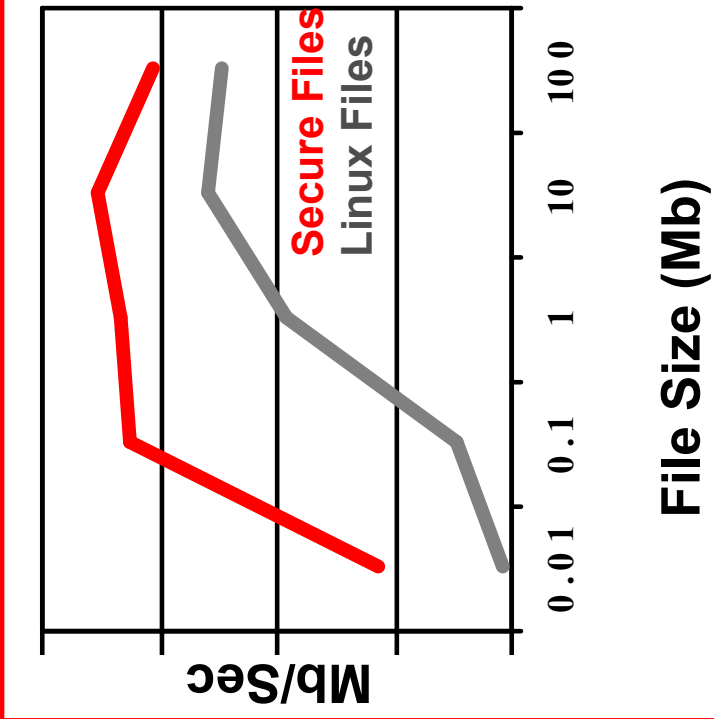
Oracle Secure Files

Breaking the Performance Barrier...

Read Performance



Write Performance





“Piction requires the ability to store very large individual digital images for use in a digital asset management system.


Using Oracle Database 11g and SecureFiles we are able to load gigabyte size blobs faster than in previous versions of Oracle. We saw improvements of over 70% in performance when loading and unloading photos and video files.

All our Oracle Database 11g Piction customers will be using the new SecureFiles feature.”

**-Marcel Kratochvil,
Piction**

ORACLE®

Storage Optimized File Store



- Deduplication
 - Multiple copies of identical SecureFiles are stored only once
 - Great for document management, testing, product “homes”
 - Significantly better performance for copy operations
- Compression
 - Intelligent – Automatically determines if data is compressible
 - 2x to 3x compression for typical files
 - Improves performance by reducing IO, buffer cache requirements, redo generation and encryption overhead
- Transparent to applications



SecureFiles value propositions

- Single security model
- Single view of data
- Single management of data
- High performance




Complete Information Management **XML DB**




XML in Oracle Database

- Native, flexible, high-performance, scaleable XML storage and processing; pioneered in 9i, becoming state-of-the-art in 11g, put the nearest competition in catch-up mode
- Repository integration enables document centric access, security, integrity, and apps development
- First product to support key XML standards including the W3C XQuery recommendation and SQL/SML
- First product to deliver SQL/XML duality
 - SQL:2003 and SQL:2007 compliant XML publishing
 - SQL/XML and XML/SQL interoperability.





Evolution of Oracle's XML Support





11g Performance Improvements for Schema-Optimized XML

- Significant performance improvements
 - Up to 10x performance improvement when ingesting XML
 - Up to 10x performance improvement when generating XML
 - XMLAgg optimization
 - General improved optimization of XQuery expression on Schema-Optimized storage
- Significant optimization for fragment level DML operations on collections



Extending the lead in 11g

- Comprehensive XML storage and indexing
 - Binary XML – Compact and efficient storage representation
- Efficient end-to-end XML Processing
 - Single, compact XML representation shared by client, mid-tier and database
 - Plan to place our Binary XML format into open source
- Enabling XML Application Development
 - Support XQuery 1.0 Standard
 - Support for Web Services in the database
 - Support Content Repository API for JAVA(JSR 170)



Complete Information Management Text






Text in Oracle Database

- Full-text Search
- Natural Language Processing
- Clustering and Classification
- Relational and Hierarchical Data Models
- Semantic Search



What is classification?





Recap of classification and clustering

- Classification
 - Supervised classification of content
 - Two ways: rules or training sets
 - You can group a number of categories into a taxonomy
 - Very useful for defining a common vocabulary in an enterprise
- Clustering
 - Unsupervised classification of patterns into groups
 - The engine analyzes the document collection and outputs a set of clusters with documents on it
 - Very useful for *discovering* patterns or nuggets in collections
 - Could be used as a starting point when there is no taxonomy present



New in 11g

- Composite Index support
- User-defined Relevance Ranking
- Segmentation and Stemming for 32 Languages
- Name Search



Complete Information Management Spatial



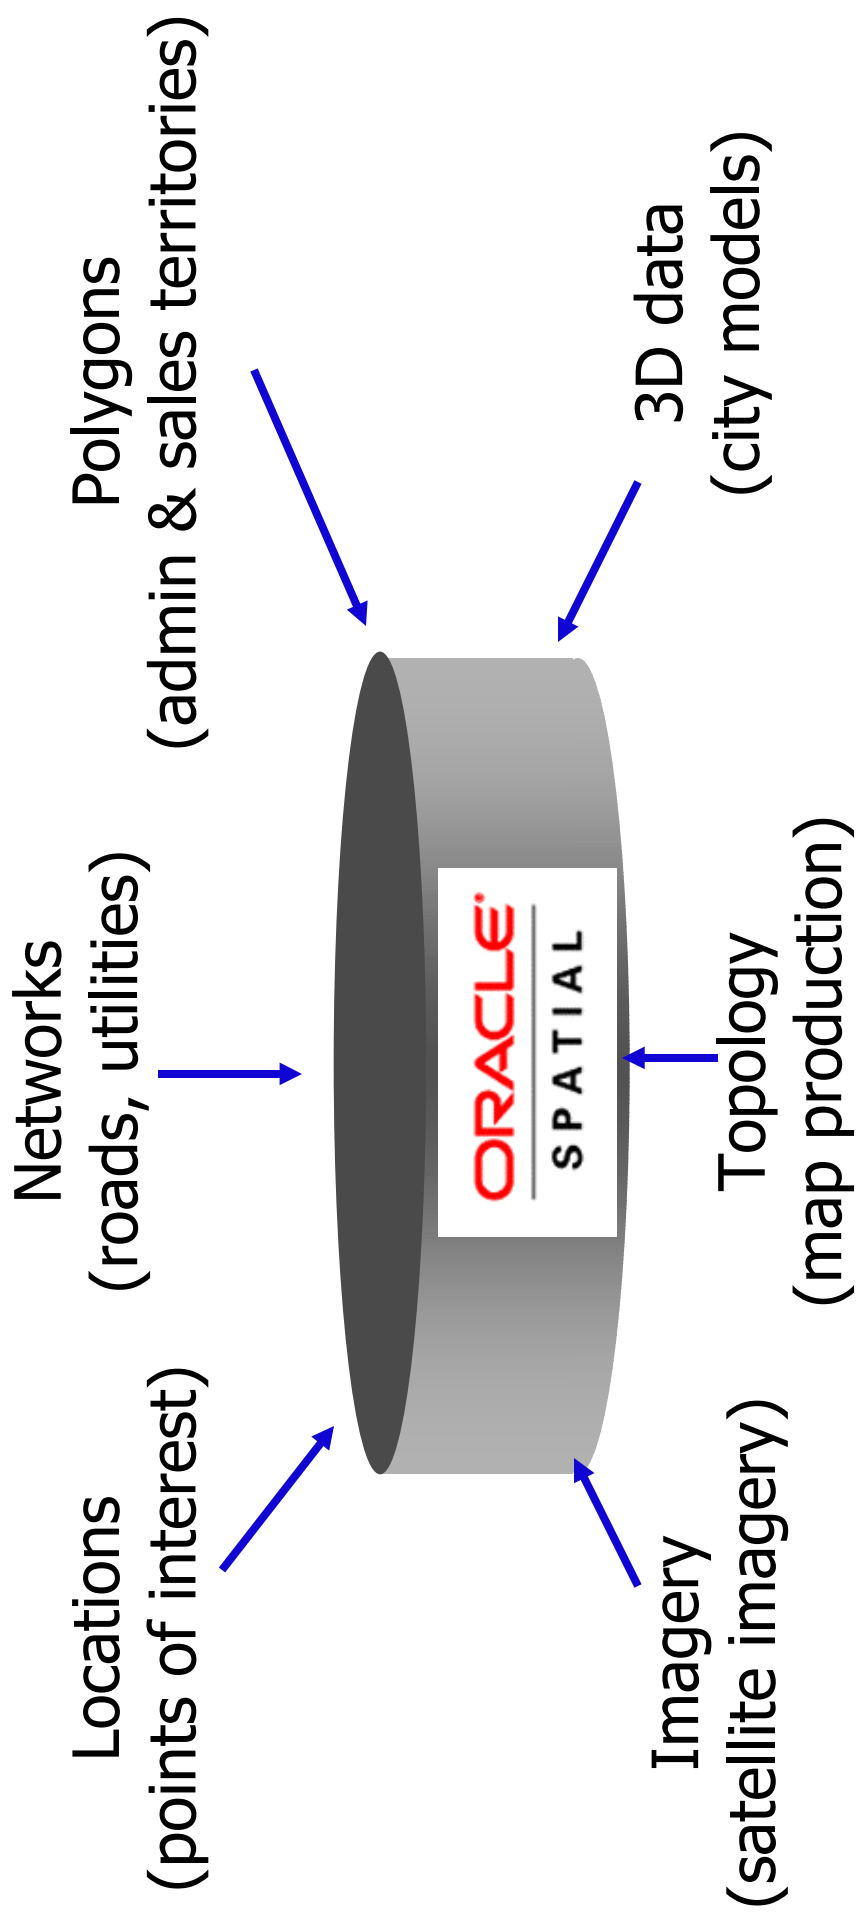


What is Location Data?

- Business data that contains or describes location
 - Street / postal address (customers, stores, factories...)
 - Sales data (sales territory, customer registration...)
 - Assets (cell towers, fire hydrants, electric transformers...)
 - Geographic features (roads, rivers, parks...)
- Anything connected to a physical location
- Every database in the world contains some form of business data that can be leveraged using spatial technologies
- Location is a “universal key”



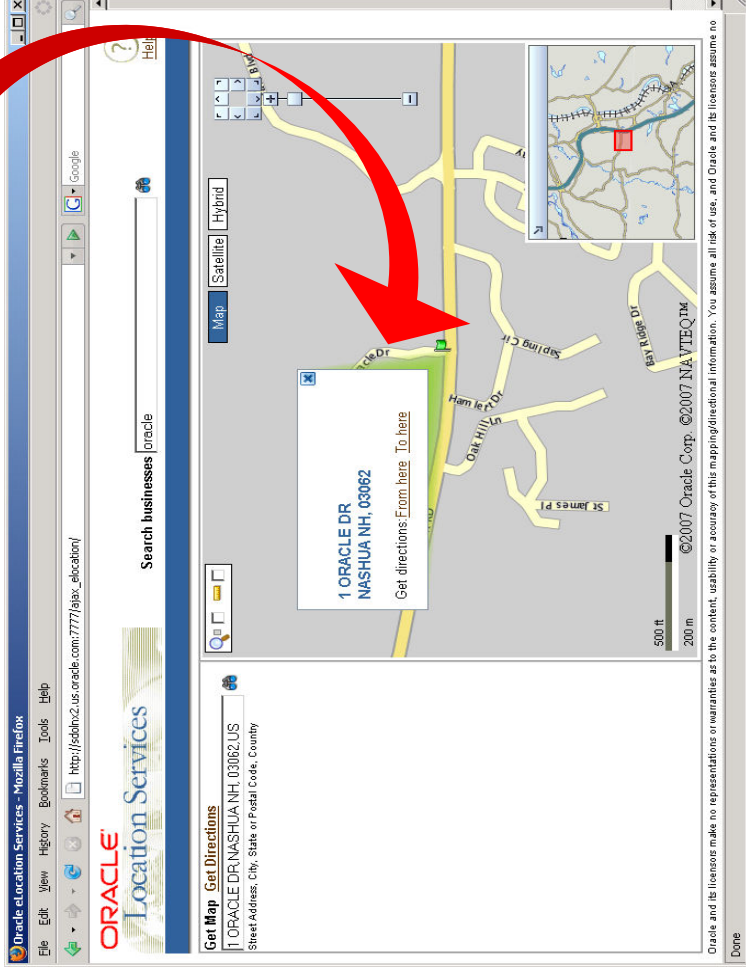
Data: Manage All Geospatial Data Types



Geocode:

- Generates latitude/longitude (points) from address
- International addressing standardization
- Formatted and unformatted addresses
- Tolerance parameters support fuzzy matching
- 100% Java, open and scalable
- Record-level and batch processes
- Data provided by leading data vendors

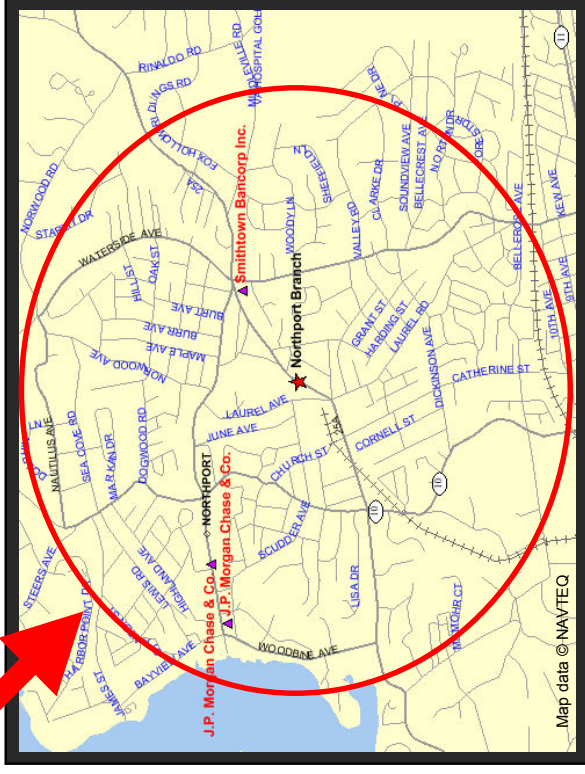
One Oracle Drive, Nashua NH, 03062



Analyze: Geospatial Data

Find all competitors within 2 miles of Northport Branch

```
SELECT c.holding_company, c.location
FROM competitor c,
bank b
WHERE b.site_id = 1604
AND SDO_WITHIN_DISTANCE(c.location,
                        b.location,
                        'distance=2 unit=mile') = 'TRUE'
```



Display: Generate Powerful Maps

Geospatial Operations and Complement Management

Facility Type: School On right click show: School within 1 mi. radius Clear Selections

Search Results Driving Directions

- BANCROFT ELEMENTARY SCHOOL**
1755 NEWTON ST NW
WASHINGTON, DISTRICT OF COLUMBIA
20010
- PREPARATORY SCHOOL OF THE DIS**
3220 17TH ST NW
WASHINGTON, DISTRICT OF COLUMBIA
20010
- SACRED HEART SCHOOL**
1625 PARK RD NW
WASHINGTON, DISTRICT OF COLUMBIA
20010
- SAN MIGUEL MIDDLE SCHOOL**
1525 NEWTON ST NW
WASHINGTON, DISTRICT OF COLUMBIA
20010

Employee counts by class/category for radius

Casual	66	Admin	0
FT	2447	Distribution	236
PTF	143	Executive	161
PTR	119	Facilities	1
Relief	100	Maintenance	63
TE	0	Nurse	30
TMP	0	Teacher	1256
Total	2875	Total	1747

Police
Hospital
University
School
Museum
Recreation
Library
Judicial

©2007 Johnson Intel.amb Enterprise - Powered by Oracle Spatial™. Map Data ©2007 NAVTEQ™

Longitude, Latitude : -77.035, 38.936

Integrate Into Your Business Application

The screenshot displays the Oracle BI Interactive Dashboards interface. The browser address bar shows the URL: <http://obee-demo:9704/analytcs/saw.d17Dashboard>. The dashboard is titled "Public Safety Dashboard" and includes navigation tabs for "Filtered Map", "Pin Map", "Lasso Map", "Crime Analysis Dashboard", "Location Analysis", and "Person Search".

The main content area is divided into three sections:

- Region - Offence**: A dropdown menu showing "Central" selected for the region and "GRAND LARCENY" selected for the offence.
- Complaints by Region - Offence - Year**: A table showing complaint counts for the Central region across the years 2002, 2003, and 2004. The total count is 266.
- Region-Offence Map**: A map of the Central region showing various service areas. A legend indicates the number of complaints per area: 0 (white), 3 (green), 7 (yellow), 13 (orange), and 50 (red).

The table data is as follows:

Region PO Desc	2002	2003	2004	Total
LARCENY,GRAND FROM AUTO	2	15	9	26
LARCENY,GRAND FROM BUILDING,IN	11	43	35	89
LARCENY,GRAND FROM OPEN AREAS,	1	6	4	11
LARCENY,GRAND BY ACQUIRING LOS	1	4	4	8
LARCENY,GRAND BY CHECK USE	1	2	3	6
LARCENY,GRAND BY CREDIT CARD U	3	3	3	9
LARCENY,GRAND FROM PERSON,PICK	8	7	15	30
LARCENY,GRAND FROM PERSON,PURS	1	6	7	14
LARCENY,GRAND FROM PERSON,INCL	32	15	47	94
LARCENY,GRAND PERSON,NECK CHAI	3	1	4	8
Total	67	114	85	266

The line chart below the table shows the trend of complaint counts over the years 2002, 2003, and 2004. The Y-axis represents the "Complaint Count" ranging from 20 to 50. The legend indicates the data series for "LARCENY,GRAND FROM AUTO", "LARCENY,GRAND FROM BUILDING,IN", "LARCENY,GRAND FROM OPEN AREAS,", "LARCENY,GRAND BY ACQUIRING LOS", "LARCENY,GRAND BY CHECK USE", and "LARCENY,GRAND BY CREDIT CARD U".

The bottom right corner of the dashboard shows the Oracle BI Interactive logo, the user "Start OC4J", and the time "9:47 PM".



Oracle Value Proposition

Enhance Oracle Database and Oracle Fusion Middleware Application Server with location analysis and mapping

- Oracle Locator and Spatial Option, Oracle MapViewer
- Mature technology – 12 years and 4 major releases

Enable enterprise business applications

- CRM, business intelligence, Enterprise Resource Planning, Asset Management, Field Service, Supply Chain, etc.
- Adopt GIS capabilities already in DBMS and App Server

Commitment To Standards

- Open Geospatial Consortium (OGC)


Partner With Leading Vendors

- System Integrators: Solutions implementation
- Data suppliers: World wide map and imagery content
- Geospatial software vendors: Specialized tools




New in 11g

- Spatial Web Services
- 3D spatial data management
- Inclusion of worldwide map data
- Others Enhancement
 - Routing engine enhancement
 - Network Data Model enhancement
 - GeoRaster enhancements
 - Support for SQL/MM spatial types and operators



Oracle Spatial 11g Enables 3D Applications




AEC City Modeling

Enterprise Mash-ups
with 3D models



Petroleum Exploration




**Complete Information
Management**
Semantic Database
RDF and Ontology support




Resource Description Framework (RDF)

- **Originally conceived as W3C's metadata model**
 - Document metadata for digital libraries, content rating, site maps, etc.
- **Simple graph data model**
 - Leverages syntactic extensibility and modularity of XML namespaces
 - Provides global extensibility through a common data model
 - Directed labeled graph: "subject/property/object"
 - Nodes are called "resources" and links "properties"



RDF Triples:

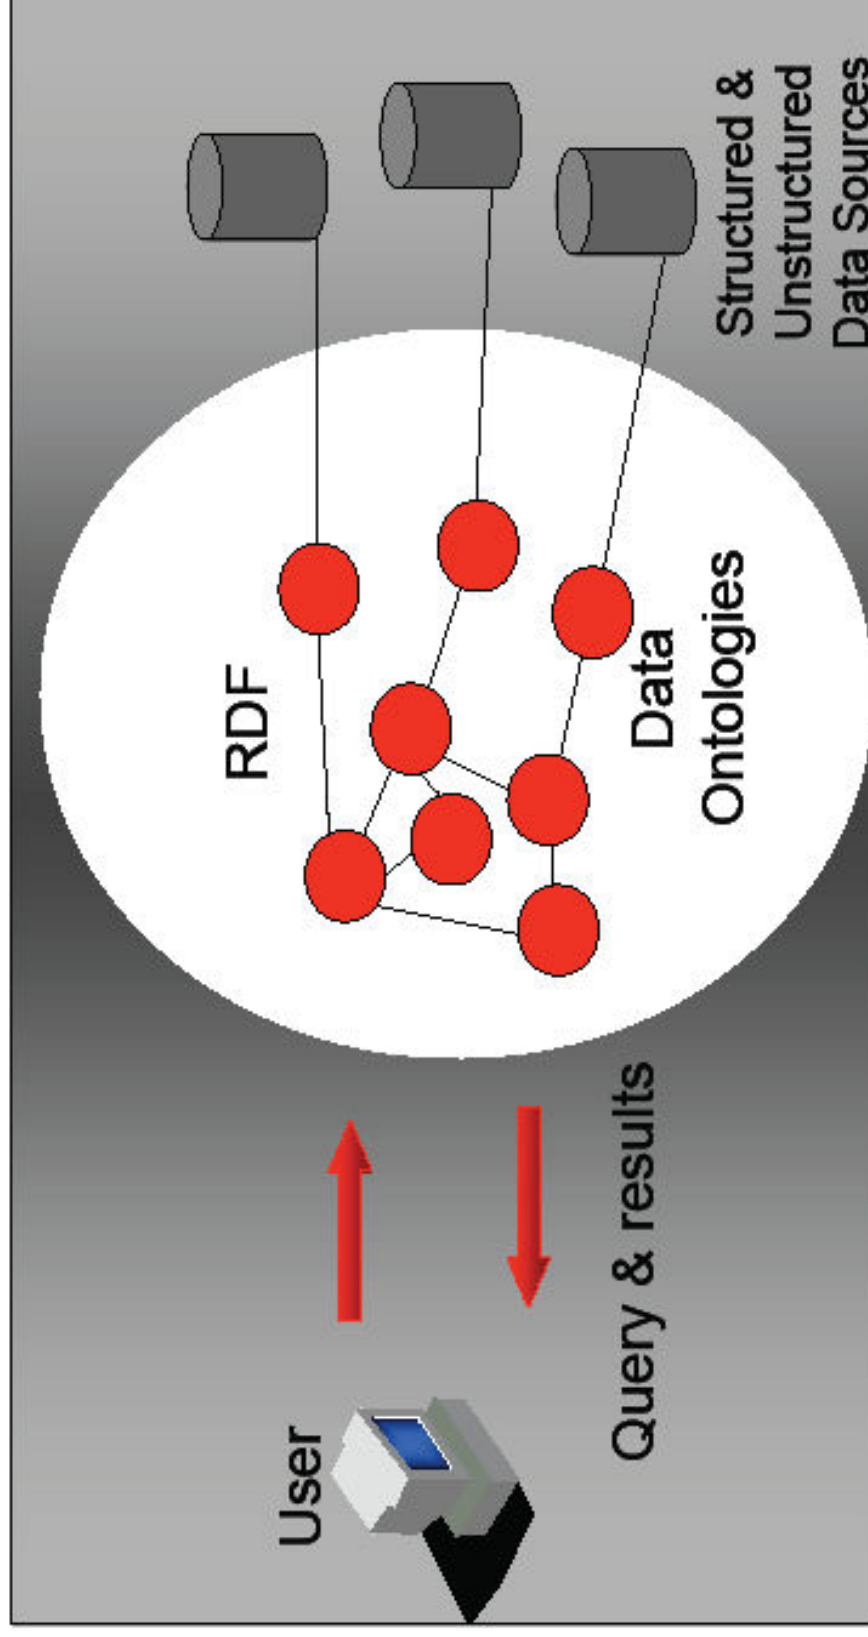
- {S1, P1, O1}
- {S1, P2, O2}
- {S2, P2, O2}



Why is this Useful?

- Designed to represent knowledge in a distributed world
- A method to decompose knowledge into small pieces, with rules about the semantics of those pieces
- RDF data is self-describing; it “means” something
- Allows you to model and integrate DBMS schemas
- Allows you to integrate data from different sources without custom programming
- Allows data re-use from multiple sources
- Supports decentralized data management
- Infer implicit relationships across data

Application Integration



Customers Managing Semantic Relationships in the Database

- Enterprise Information Integration
 - Metadata mapping of application data models
 - Use of structured data
- **Early adopters:** Life Science & Pharma, Utilities
- Knowledge Management
 - Analyze social networks “who / what relates to X ?”
 - Create ontology to facilitate ETL – what to extract?
 - Use of unstructured and structured data
- **Early adopters:** Intelligence, Public Health, Life Science





New in 11g

- 10gR2 introduced RDF support for graph data model
- 11g adds new data model for scalable, secure, standard-based storage and query of OWL and RDF Ontologies
 - SQL-based query of these Ontologies
 - Scales up to a billion “triples”
- “Semantic Match” SQL Operator to query any relational data using referenced Ontology
 - Expands results to include related & semantically relevant content
- Inferencing Support
 - Native, practical subset of OWL inferencing capability
 - Built-in RDFS inferencing
 - User-defined rules

Example: Query with Semantic Operators

Find <id, diagnosis> info for all patients who have been diagnosed as afflicted with diseases of type **Immunodeficiency_Syndrome** that are within a specified distance from it.

```
SELECT id, diagnosis
FROM Patients_Data
WHERE SEM_RELATED ( diagnosis,
                    'rdfs:subClassOf',
                    'Immunodeficiency_Syndrome',
                    'Cancer_ontology', 1) = 1
AND SEM_DISTANCE (1) <= 2;
```



Complete Information Management Multimedia





Multimedia in Oracle Database

- Multimedia content subject to database enforced referential integrity and transaction control
- Metadata extraction, editing and indexing
- Media-specific features such as metadata editing, thumb nailing and format conversion
- Database enforced private and public metadata
- Simplifies secure delivery of content to streaming media servers



New in 11g

- Multimedia
 - 3X performance improvement for common image processing operations
 - Large media handling (up to 128 TB)
 - Java Advanced Imaging (JAI) support

New in Oracle Database 11g


Critical New Data Types



RFID
Data Types



DICOM
Medical Images



3D Spatial
Images

DICOM is...

- Digital Imaging and Communications in Medicine
 - A communication protocol and a file format
 - An evolving standard with approximate yearly revisions
- Huge metadata content ~ 2,000 standard attributes
- Content covers many aspects of radiology imaging, waveforms, Structured Reports:
 - Image (encoding, display, management), Patient/visit management, Service/command semantics
- Click for [DICOM Standard web site](#)



Digital Imaging and Communications in Medicine

HEMMA, Suite 4047
1506 North 17th Street
Rosslyn, VA
22209
Ph: (703) 844-3205


DICOM Medical Imaging Support

- Used by virtually all medical devices (CT, MRI, PET, ...)
- Database support the data management needs of clinical medicine and life sciences
- Built-in support for metadata extraction and searching
- Support for privacy regulations (HIPAA), annotation and format conversion

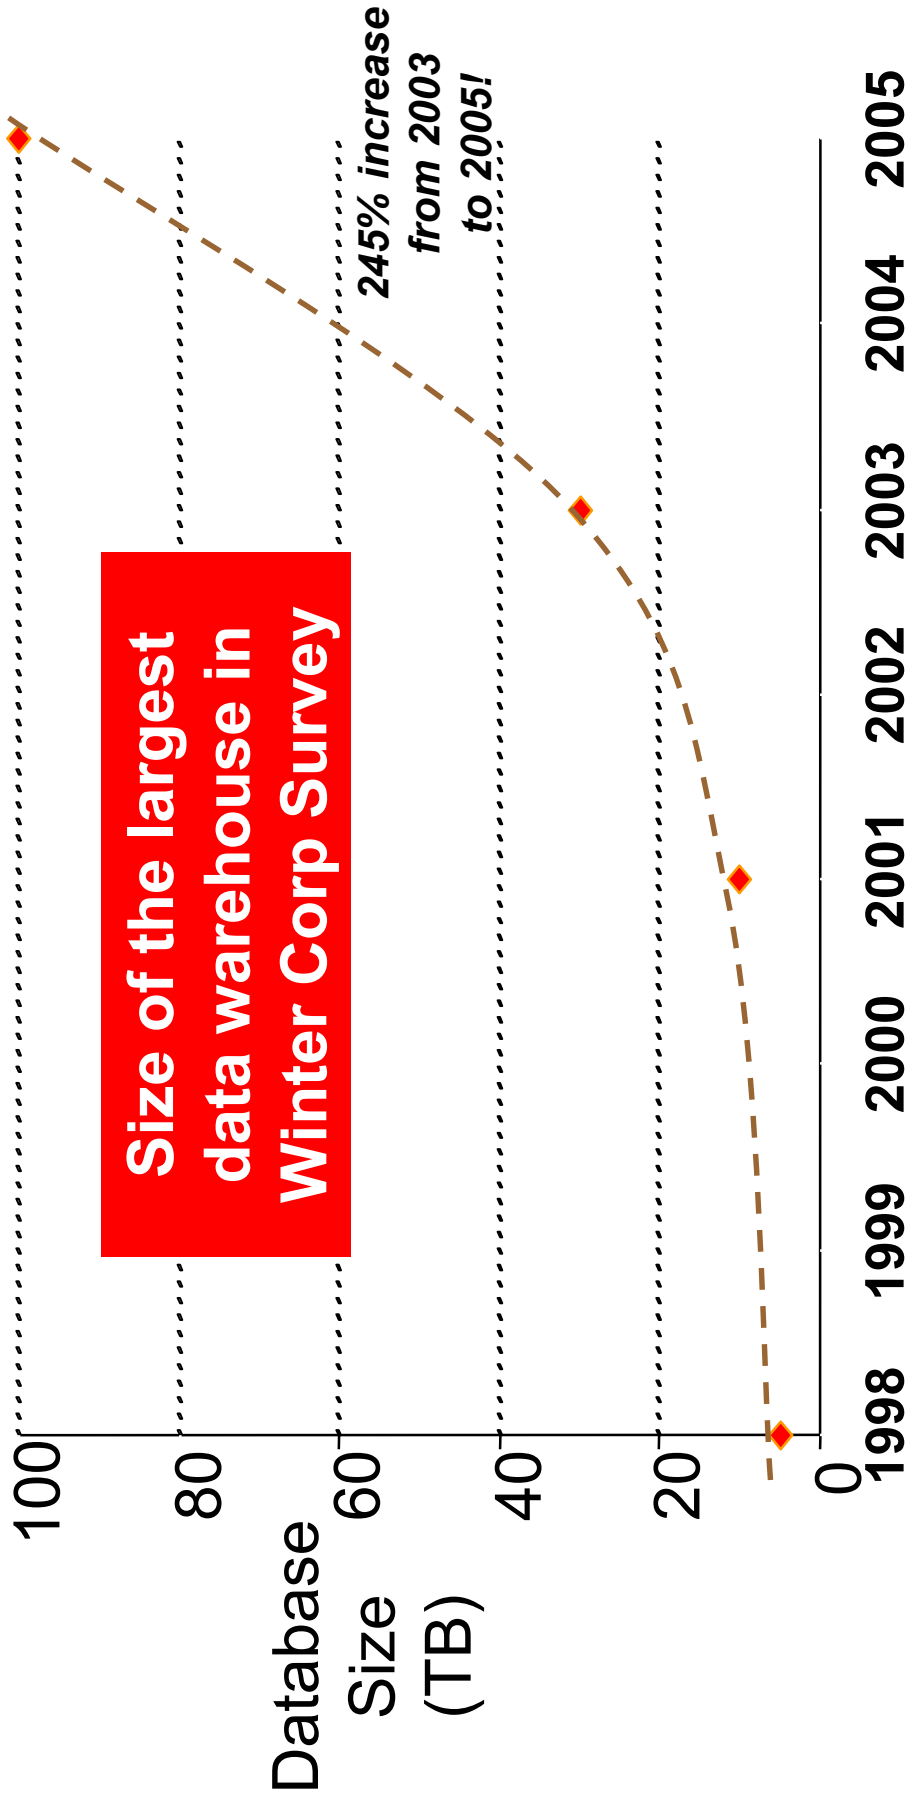




Complete Information Management Information Lifecycle Management



Growing Data Volumes






Data is Tripling every Two Years

And storage costs increase accordingly...

If 1gb storage = \$72






So how do we control storage costs?

Lets take example of 13.5 terabytes...

If 1gb storage = \$72






Traditional Storage Approach

All data resides on single storage tier

High Performance
Storage Tier
= \$72 per Gb

Active

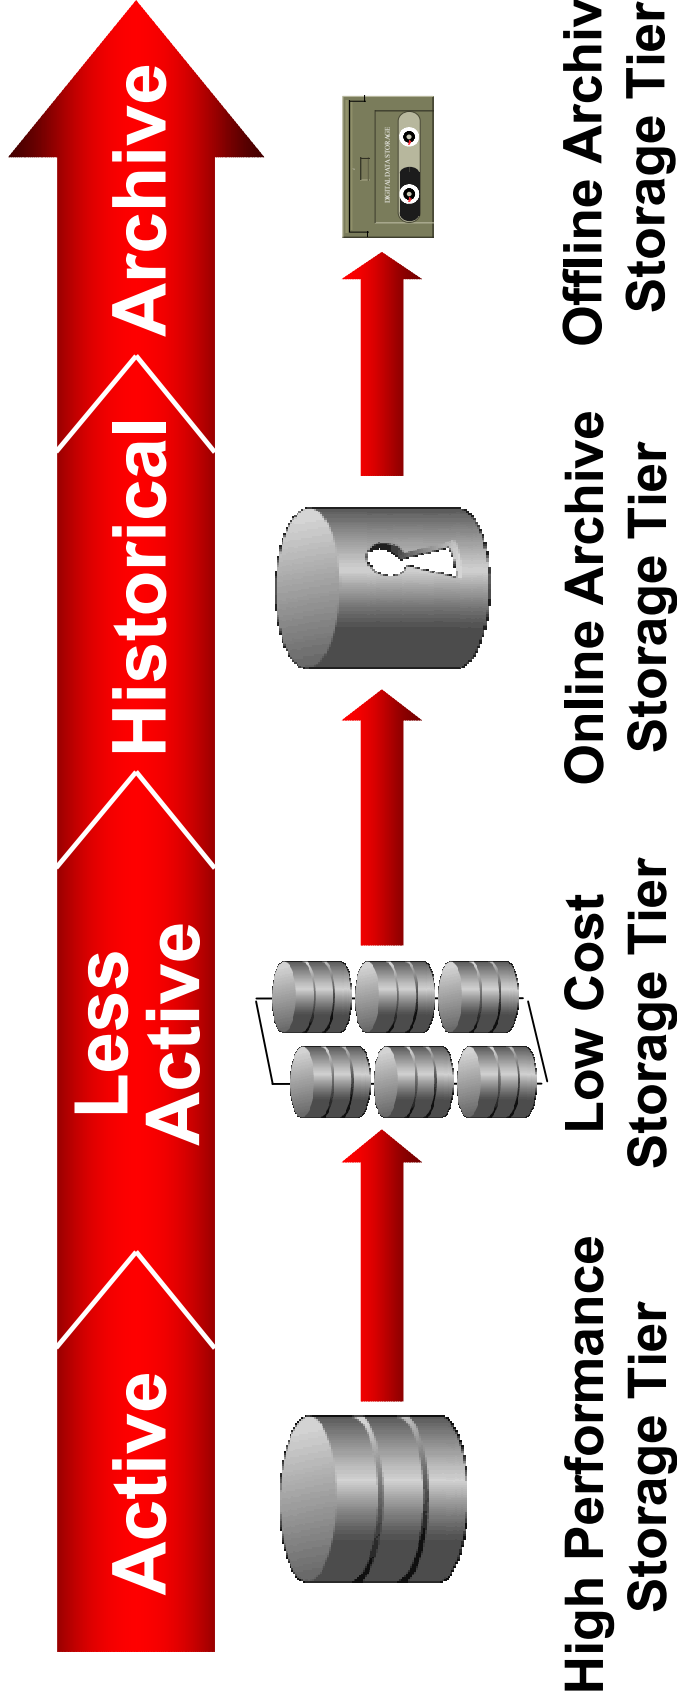


All data on active
= \$972,000!

Information Lifecycle Management

“The policies, processes, practices, and tools used to align the business value of information with the most appropriate and cost effective IT infrastructure from the time information is conceived through its final disposition.”

Storage Networking Industry Association (SNIA) Data Management Forum






Information Lifecycle Management

Partition data onto appropriate storage tier

High Performance
Storage Tier
= \$72 per Gb

Low cost
Storage Tier
= \$14 per Gb

Read only
Storage Tier
= \$7 per Gb








Information Lifecycle Management

Move data onto appropriate storage tier

High Performance
Storage Tier
= \$72 per Gb

Low cost
Storage Tier
= \$14 per Gb


Read only
Storage Tier
= \$7 per Gb





Information Lifecycle Management

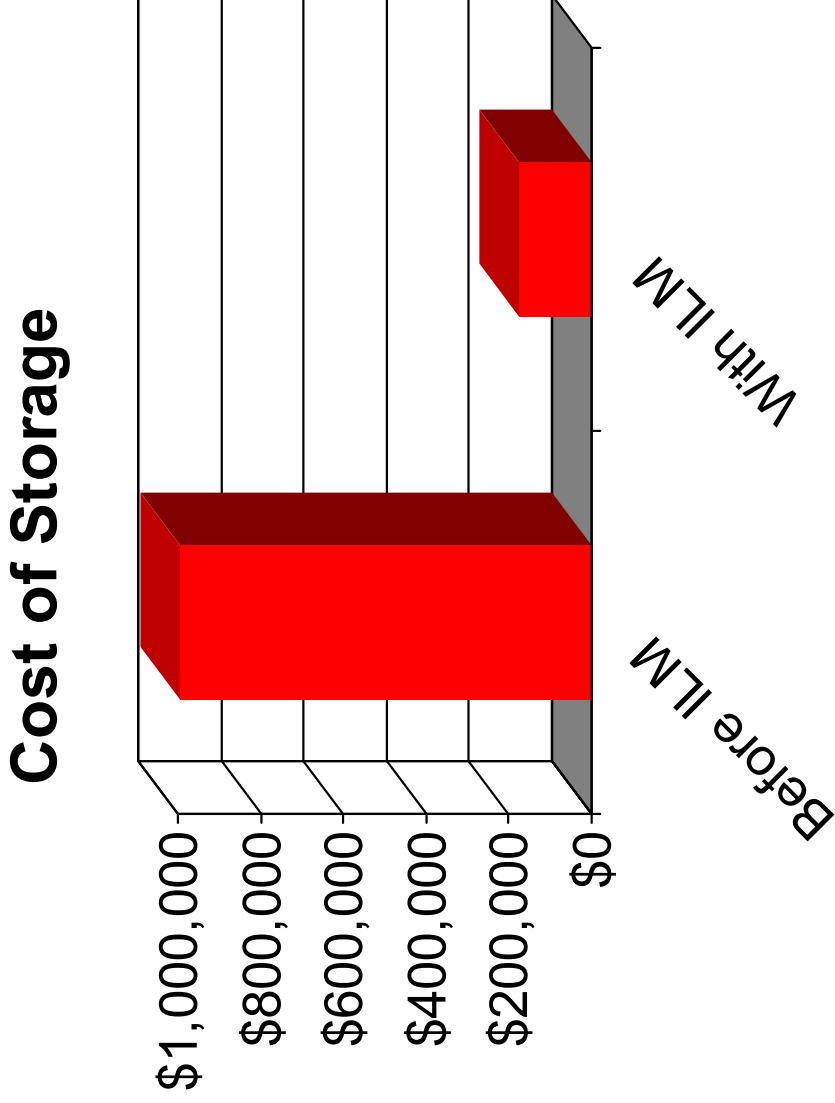
Reduce storage costs accordingly





Impact of Partitioning?


Significantly reduces storage costs...



Advanced Compression Option

Compression for Mainstream

- **Compress Large Application Tables**
 - Transaction processing, data warehousing
- **Compress All Data Types**
 - Structured and unstructured data types
- **Typical Compression of 2-3 X**
 - Cascade storage savings throughout data center



Up To


3X

Compression

ORACLE®

Advanced Compression

Reduces storage requirements across all tiers...



\$49,800

\$67,700

\$58,000

Lets use compression factor of 3

\$16,600


\$22,600

\$19,400



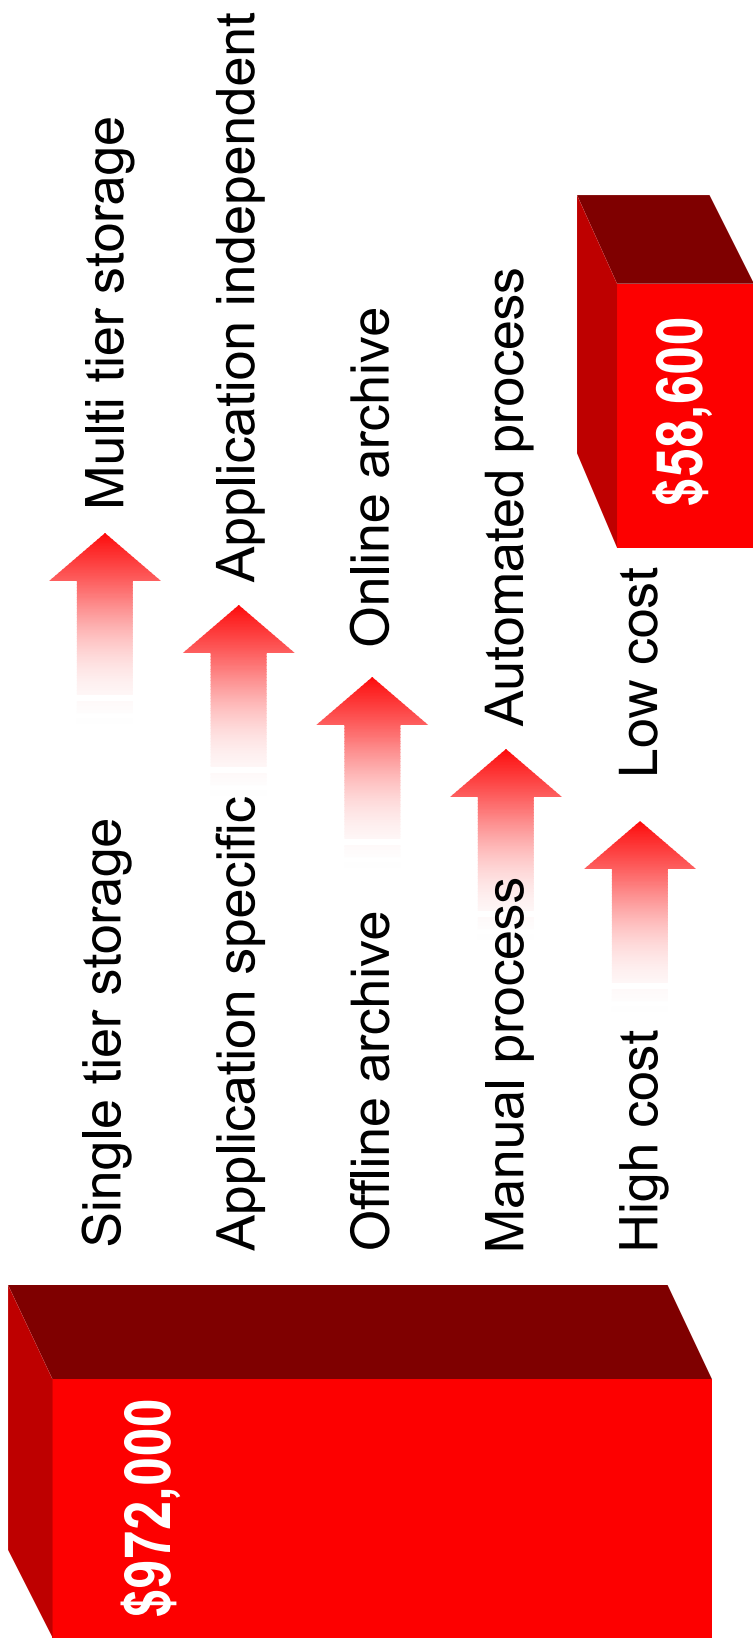
Impact of Data Compression

Compression further reduces storage costs...



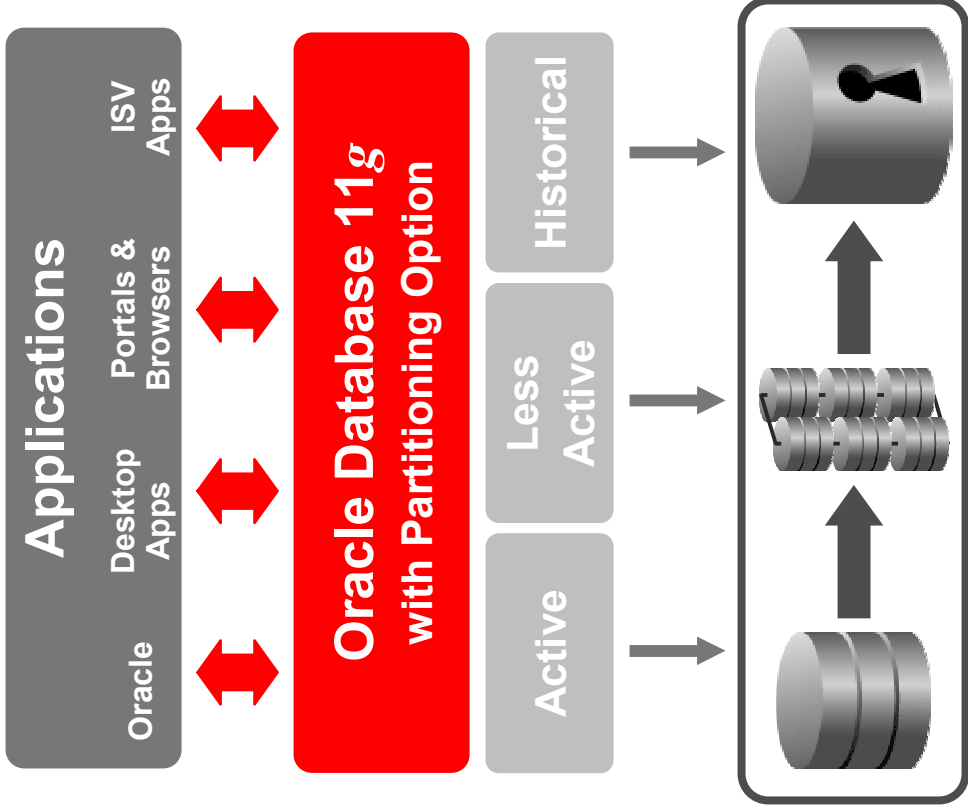
Why Oracle Database 11g?

Managing data growth



Oracle ILM

Optimized storage cost and performance



- **Low Cost**
 - Can use low cost storage to reduce the cost of retaining data
- **Transparent to Applications**
 - Applications are unchanged
- **Enforceable Compliance Policies**
 - Sophisticated techniques to define and enforce data policies
- **Fine Grained Control**
 - Lifecycle of groups of business data is managed down to individual rows
- **Hardware Independent**



Conclusion



Managing All Your Information

- The key trends driving unstructured data into Database
- Oracle is the leading platform for managing all your information
 - Dominant in Location and Spatial
 - Leaders in XML Database
 - Proven in Information Retrieval and Media Management



For More Information

Look out for

[www.oracle.com](http://www.oracle.com/database/11g)

**A dedicated microsite for Oracle Database 11g with more information,
assets and the latest updates**