

ORACLE
OPEN
WORLD

experience

OPENWORLD

November 11–15, 2007

ORACLE®



ORACLE®



SecureFiles – Next Generation Unstructured Data Management

Oracle Open World - November 2007

Ravi Rajamani
Sr. Product Manager

Scott Lynn
PMTS


Agenda

- Enterprise Data
- Current Approaches to managing data
- SecureFiles: Changing the paradigm
- High Performance
- Advanced Features
- Out-of-the-Box benefits with SecureFiles
- Demonstration
- National Ignition Facility – Real world experience

Enterprise Data



Enterprise Data



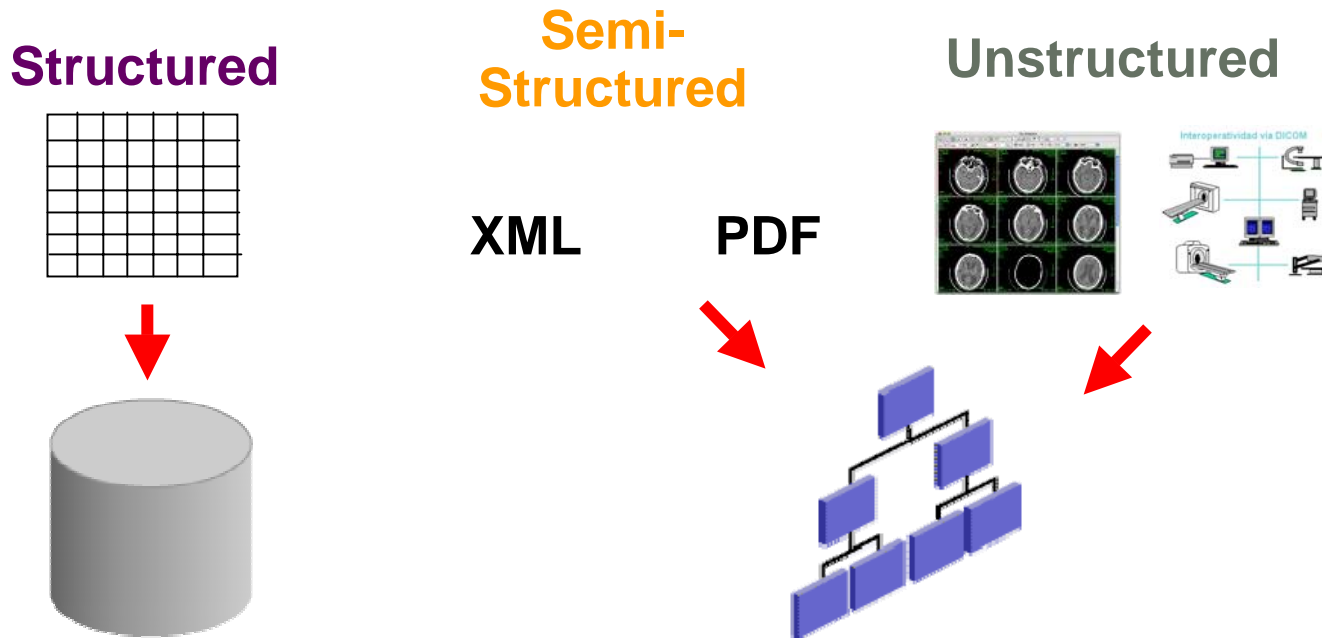
Yearly Data Growth (estimates)

- Structured – 15 – 20%
- Unstructured – 50- 200 %

- Drivers for unstructured data growth
 - Increased digitization of content
 - Healthcare, Finance, Insurance, Banking
 - Compliance
 - Web 2.0
 - Scientific/Research Community

Managing Information

- Organizations need to efficiently and securely manage



- Simplicity and performance of file systems makes it attractive to store file data in file systems, while keeping relational data in DB
- Enterprise applications manipulate both files and relational data
 - Rich user experience, compliance, business integration


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

Oracle SecureFiles

Consolidated Secure Management of Data

- SecureFiles is a new database feature designed to break the performance barrier keeping file data out of databases
- Similar to LOBs but much faster, and with more capabilities
 - Transparent encryption, compression, deduplication, etc.
 - Preserves the security, reliability, and scalability of database
 - Superset of LOB interfaces allows easy migration from LOBs
- Enables consolidation of file data with associated relational data
 - Single security model
 - Single view of data
 - Single management of data




Innovations

- SecureFiles is a major rearchitecture of how the database handles unstructured (file) data
 - Not an incremental improvement to LOBs
- Entirely new:
 - Disk format
 - Network protocol
 - Caching and locking
 - Redo and undo algorithms
 - Space and memory management
 - Cluster consistency algorithms




SecureFiles Vs. NFS

File Read Performance



File Write Performance




— SecureFiles — NFS


- SQL+ File test, single stream, single host
- Using Secure Files is faster across the board
 - Upto 2x faster for Queries, 6x for Inserts
 - Tests run using both SecureFiles and NFS/ext3 in metadata journaling only (default for NFS)
- Filesystem-like performance

SecureFiles Vs. BasicFiles


Writes: New Space



Writes: Reused Space



Read Performance



— SecureFiles — LOBs

SQL+ File Test: Concurrent Reads/Writes, OCI, 4 streams


- Adding Files using New Disk Space: **Up to 2x faster**
- Adding Files Reusing Space: **Up to 22x faster**
- Reads up to **3x faster**

Breaking Performance Records!

– *Made for Speed*

- Our Lab Setup
 - 3 EMC CX700 connected through 2 switches, 12 LUNs spanning the 3 arrays
 - 4 RAC Nodes
 - Xeon 3.4GHz, 2 CPUs, each hyperthreaded, 6GB of memory
 - HBA is 2Gbs - limits I/O bandwidth of a node to 276 MB/s
 - OCI, 8 concurrent streams total.
 - SecureFiles: Filesystem_like_logging, NoCache
- Provides Sustained Throughput of
 - 776 MB/s for File Read
 - 462 MB/s for File Writes
 - **Satisfies 38 TB/day of data ingestion rate** using off-the-shelf hardware & software


High Performance Experiment



Advanced Features - Compression

- SecureFiles introduces Compression for unstructured data
- Huge storage savings
 - Industry standard compression algorithms
 - 2-3x compression for typical files (doc, pdf, xml)
- Automatically detects if SecureFile data is compressible
 - Skips compression for already compressed data
 - Auto-turn off compression when space savings are minimal or zero
- Server-side compression
 - Allows for random reads and writes to SecureFile data
- Two levels of compression provide different compression ratios
 - Compression Levels: MEDIUM (default), HIGH
 - Higher the degree of compression, higher the latency and CPU overhead incurred
- Part of the Advanced Compression Option

Advanced Features - Deduplication



- Enables storage of a single physical image for duplicate data
- Significantly reduces space consumption
- Dramatically improves writes and copy operations
- No adverse impact on read operations
 - May actually improve read performance for cache data
- Specially useful for content management, email applications and data archival applications
- Part of the Advanced Compression Option

Advanced Features - Encryption

- Extends Transparent Data Encryption (TDE) functionality to SecureFile data
 - Data encrypted on disk
- Support for industry-standard encryption algorithms
 - 3DES168
 - AES192 (default)
 - AES256
- Unified security level for both file and relational data
- Part of the Advanced Security Option

SecureFile Interfaces

- SecureFiles can be accessed by both database clients and file system clients
- Database clients use standard LOB interfaces
 - JDBC, ODBC, OCI, OCCI, .NET, PL/SQL
- File system clients use the file system protocols implemented in the XML DB or Content DB repository
 - FTP access
 - WebDav Access
 - Http Access

Out of the Box Benefits

- Easy conversion from LOBs to SecureFiles using Online Redefinition
 - Existing LOB applications benefit with no changes
- New installations
 - By setting `db_securefiles= ALWAYS`
- File system based applications can switch to SecureFiles using Oracle's native implementation of industry standard files protocols
 - FTP, WebDAV, Http
- SecureFiles is fully integrated with
 - XML DB, Oracle Multimedia, Oracle Spatial, Content DB



Demonstration

On-line Redefinition

- Table (Before Redefinition)

```
create table doc_tab (pkey number(10) not null,  
2                      document clob)  
3  lob(document) store as doc_tab_document_lobseg  
4  (nocache logging retention);
```

- Redefinition Table

```
create table doc_tab (pkey number(10) not null,  
2                      document clob)  
3  lob(document) store as SECUREFILE doc_tab_document_lobseg_sf  
4  (nocache logging retention auto COMPRESS DEDUPLICATE);
```

Application Transparency

- Create Table

```
create table doc_tab (pkey number(10) not null,  
                    document clob)  
lob(document) store as SECUREFILE doc_tab_document_lobseg  
  (nocache logging retention auto  
   COMPRESS DEDUPLICATE ENCRYPT);
```

- Create text index on compressed , deduplicated and encrypted SecureFile

```
begin  
ctx_ddl.create_index_set('document_iset');  
ctx_ddl.add_index('document_iset','document');  
end;  
  
/  
CREATE INDEX document_idx ON doc_tab(document) INDEXTYPE IS  
  CTXSYS.CONTEXT;
```



National Ignition Facility – Real World Experience

Philip A. Adams

The Best of Files and Databases

- SecureFiles have all the leading-edge file system capabilities
 - Deduplication, Encryption, Compression, Logging, Versioning
- SecureFiles have advanced database features not in file systems
 - Transactions, Read Consistency, Flashback
 - Readable Standby, Consistent Backup, Point in Time Recovery
 - Fine Grained Auditing, Label Security
 - Sliding Inserts
 - XML indexing, XML Queries, XPath
 - Real Application Clusters
 - Automatic Storage Management
 - Partitioning and ILM
 - Search across meta-data and file content
- Capabilities go far beyond any other database or file system
 - Having the best of both worlds removes the need to compromise

For More Information

- Oracle Technology Network (OTN)
 - <http://www.oracle.com/technology/products/database/securefiles/index.html>
- SecureFiles Tutorial
 - http://www.oracle.com/technology/obe/11gr1_db/datamgmt/securefile/securefile.htm
- Advanced Compression White Paper
 - <http://www.oracle.com/technology/products/database/oracle11g/pdf/advanced-compression-whitepaper.pdf>

Demonstration

Oracle DEMOgrounds (O07)
Moscone West Exhibition Hall

Sessions From Oracle Development

Monday, Nov 12

- S291688 - Oracle Database 11g: Next-Generation Performance and Scalability, 4:45 - 5:45 pm, Moscone South 104

Tuesday, Nov 13

- S292033 - Oracle Database 11g: Next-Generation Data Management Platform, 10:45 - 11:45 pm, Moscone South 304

- S291861 – Advanced Compression in Oracle Database 11g, 12:15 - 1:15 pm, Moscone South 304

- S291857 – Next Generation Unstructured Data Management using SecureFiles, 3:15 - 4:15 pm, Moscone South 304

- S291857 – Total Recall using Flashback Data Archive, 4:45 - 5:45 pm, Moscone South 304

Thursday, Nov 15

- S291862 - Improve the performance and manageability of your NFS environment with Oracle Database 11g Direct NFS Client, 11:30 am - 12:30 pm, Moscone South 304

- S291689 – Using Oracle Database in Amazon Web Services Environment, 1:00 - 2:00 pm, Moscone South 304

Demos From Oracle Development

Monday, Nov 12 – Thursday, Nov 15
Oracle DEMOgrounds, Moscone West

SecureFiles

Advanced Compression

Total Recall

Recovery Manager (RMAN) and Flashback Technologies



QUESTIONS
ANSWERS





ORACLE IS THE INFORMATION COMPANY

ORACLE®