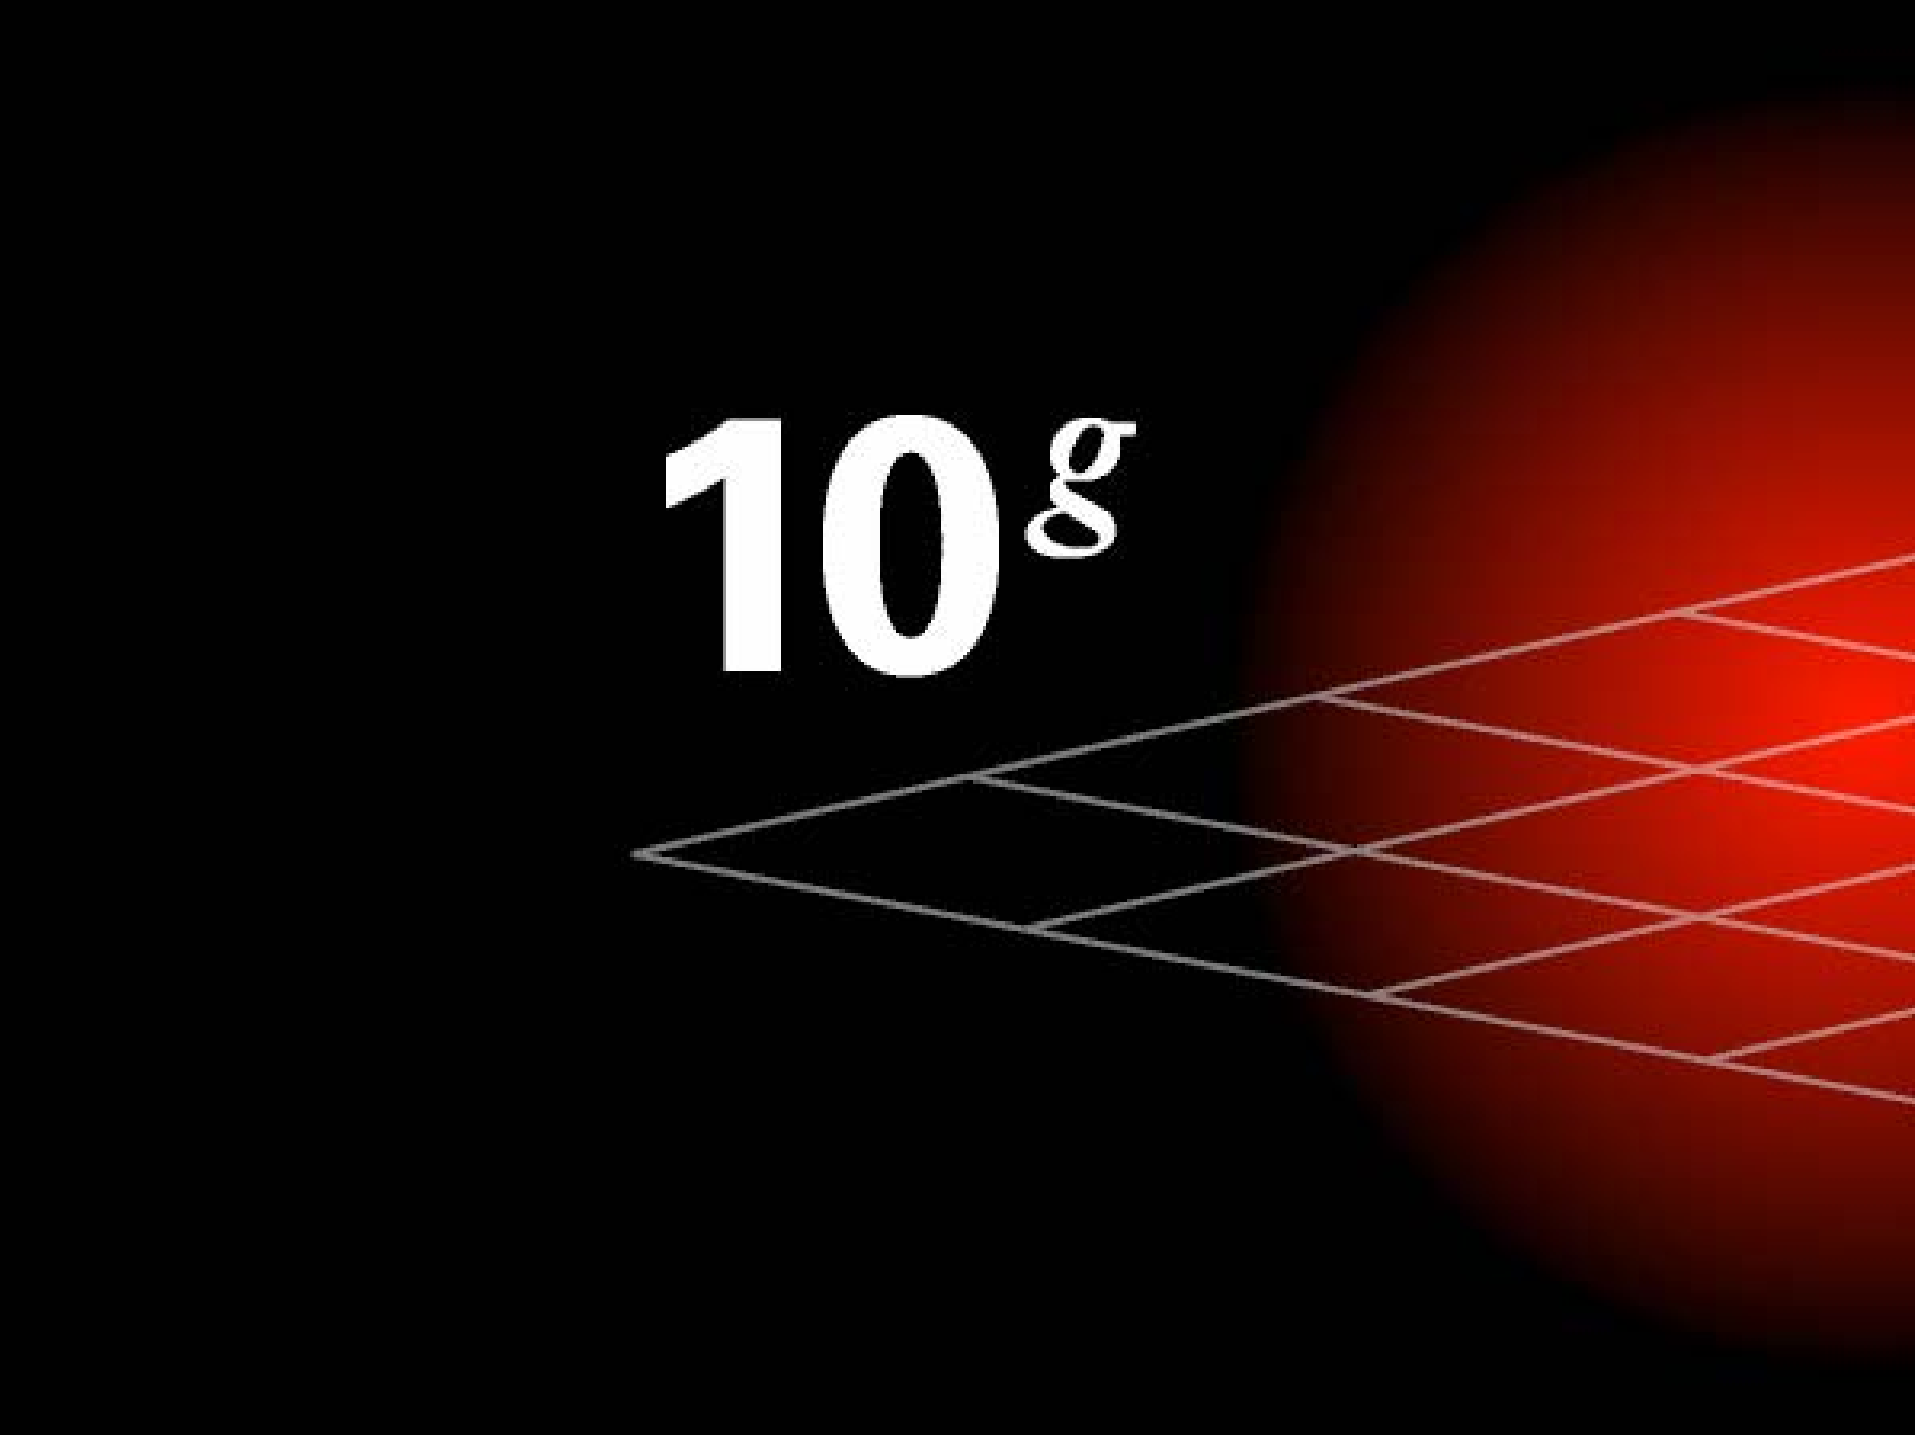
10g

Thomas Kyte

thomas.kyte@oracle.com
Oracle Corporation
http://asktom.oracle.com

Resources

- metalink.oracle.com
- asktom.oracle.com
- otn.oracle.com
  - Discussion forums
  - Software download
  - *Documentation*
  - Technology Corners
- www.oracle.com/oramag
- groups.google.com
  - comp.databases.oracle.*
- IOUG

Oracle10i High Availability Features

- **Backup and Recovery**
  - Automated Disk-Based Backup and Recovery
  - Drop Database
  - RMAN Database Deregistration
  - Backup of Standby Control File
  - Automated TSPITR Instantiation
  - Simplified Recovery Manager Cataloging of Backup Files
  - Automated Channel Failover for Backup and Restore
  - Automated File Creation During Recovery
  - Simplified Backups to Disk
  - Proxy Copy Backup of Archivelogs
- **Backup and Recovery (Cont)**
  - Incrementally Updated Backups
  - Simplified Recovery Through Resetlogs
  - Restore Tolerates Corrupt or Missing Backups
  - Full Database Begin Backup Command
  - Change-Aware Incremental Backups
- **Data Guard**
  - Data Guard Broker Support for RAC
  - Automated LogMiner Configuration
  - Log Miner Support for Index-Organized Tables
  - Secured Redo Transmission
- **Data Guard (Cont)**
  - Fine-Grained Supplemental Logging
  - Named Data Guard Configurations
  - Simplified Zero Data Loss for Data Guard SQL Apply
  - Zero Downtime Instantiation for SQL Apply
  - Real Time Apply
  - Redo Data and Archivelog Compression
- **Flashback**
  - Flashback Database
  - Flashback Reinstantiation
  - Flashback Standby Database
  - Flashback Table
  - Flashback Row History
  - Flashback Drop
  - Flashback Transaction History
- **Online Operations**
  - Enhanced Online Redefinition
  - Improved Handling of DDL Locks on Busy Tables
  - Signature-Based Dependency Tracking Using Synonyms
  - Rolling Upgrades with Data Guard SQL Apply

Oracle10g Manageability Features

- **Infrastructure**
  - Easy OracleJVM Monitoring
  - SYSAUX Tablespace
  - Enhanced Wait Model
  - Database Features Usage Tracking
  - Enhanced Database Time Model
  - Top Consumers Performance Monitoring & Diagnostic
  - JDBC: End-to-end Java/J2EE Tracing
  - SQLJ: End-to-end Java/J2EE Tracing
  - Automated Routine Administration Tasks
  - Proactive Server-based Alerts and Notifications
  - Database Statistics and Workload Respository
  - Threshold-Based Alerts
- **Server Configuration**
  - Automated Storage Management Configuration
  - Automated RAC Services Configuration
  - Simplified Upgrade for RAC and OPS Databases
  - Automated Enterprise Manager Configuration
  - Automated Portable Clusterware Installation
  - Automated Configuration of Recovery Area
  - Out-of-the-box LDAP Configuration
  - Simplified Initialization Parameters
  - Easy Upgrade
  - Upgrade Information Tool
  - Simplified Database Install
  - User Default Tablespace
- **Instance Tuning**
  - User-Initiated Buffer Cache Flushing
  - Database Resource Manager - Adaptive Consumer Group Mapping
  - Database Resource Manager - Fixed CPU Quota
  - New Performance Overview Charts in Oracle Enterprise Manager
  - Improved SQL Reporting Using Oracle Enterprise Manager
  - Integration of Resource Manager and Profiles
  - Automated Checkpoint Tuning
  - Transaction Rollback & Recovery Monitoring
  - Database Performance Analyzer
  - Redo Logfile Sizing Advisor
  - Automated SGA Memory Tuning
- **Application Tuning**
  - SQLAccess Advisor
  - Materialized View Tuning API
  - SQL Tuning Advisor
  - Enhanced SQLAccess Advisor
  - Automatic Optimizer Statistics Collection
- **Backup and Recovery**
  - Enhanced RMAN Reporting
  - Backup Compression
  - Bounded Backup Window
  - Manage RMAN Scripts
  - Recovery Area Full Alert
- **Storage Management**
  - Multiple Default Temporary Tablespace Support for SQL Operations
  - Automated Storage Management
  - Rename Tablespace

Oracle10g Manageability Features (continued)

- **Network Management**
  - Back Up Directory Naming Entries to Local Naming File
  - Dynamic Connection Manager Configuration
  - Easy Connect Naming Method
  - Easy OCI/JDBC Install
  - Improved Network Outage Detection
  - Automated LDAP Discovery for Clients
  - Improved Connection Manager Access Rules
  - Automated Shared Server Configuration
  - Simplified Shared Server Configuration Parameters
- **Space, Object, Transaction Mgmt**
  - Automatic Undo Retention Tuning
  - Segment Shrink Advisor
  - Segment Size Estimation
  - Online Segment Shrink
  - Proactive Tablespace Management
  - Undo Advisor
- **Database Utilities**
  - SQL*Plus - SPOOL CREATE, REPLACE, and APPEND Options
  - SQL*Plus - Show Recycle Bin
  - SQL*Plus - DBMS_OUTPUT After SELECT
  - SQL*Plus - SET SQLPROMPT Runtime Variable Substitution
  - SQL*Plus - COMPATIBILITY Command Line Option
  - SQL*Plus - Run glogin.sql and login.sql After Each CONNECT
  - iSQL*Plus Input Prompting
- **Scheduler**
  - Core Scheduler Features
  - Usability and Manageability
  - Monitoring
  - Recovery
  - Support for Jobs in a Clustered Environment
  - Job Processing
  - ILMS Support
  - Increased Throughput
  - Redo Minimization
  - Dynamic Slave Pool

Oracle10g Business Intelligence Features

- **Real Time Data Warehousing**
  - Asynchronous Change Data Capture
  - Cross-Platform Transportable Tablespaces
  - Enhanced Table Functions
  - External Tables Unload
  - Enhanced MERGE Functionality
  - Data Pump Export and Import Utilities
  - Parallel Data Pump Export and Import
  - SQL*Loader Direct Path Load Support for ROWID Datatype
  - SQL*Loader Direct Path Load Support for VARRAY Datatype
  - SQL*Loader Direct Path Load Support for XMLType Tables
- **VLDB**
  - Bigfile Tablespace
  - Enhanced Partition Management in Enterprise Manager
  - Global Partitioned Indexes - Hash Partitioning
  - Partitioned Index-Organized Tables - List Partitioning
  - Partitioned Index-Organized Tables - LOB Support
  - Partitioned Index-Organized Tables - Global Index Maintenance
  - Enhanced Bitmap Index Performance and Space Management
  - Local Partitioned Indexes Manageability Improvements
- **Summary Management**
  - Summary Management OLAP Query Performance Improvements
  - Summary Management - Enhanced Partition-aware Materialized View Refresh
  - Summary Management - EXPLAIN PLAN Shows Materialized View Access
  - Summary Management - Enhanced Dimensions
  - Summary Management - Nested Materialized View Refresh
  - Summary Management - Enhanced PCT Refresh
  - Summary Management - Query Rewrite Uses Multiple Materialized
- **OLAP**
  - Parallel AGGREGATE Command
  - Parallel SQL IMPORT Command
  - PL/SQL Creation of Analytic Workspaces
  - Intra-Dimensional (Attribute) Aggregation
  - Enhanced Composite Dimensions Indexing
  - PL/SQL OLAP Measure Calculation Definition
  - Complete Analytical Workspace Management
  - Multi-Writer Support
  - XML Interface for Analytic Workspace

Oracle10g Business Intelligence Features (Continued)

- **Data Mining**
  - Frequent Itemsets
  - Attribute Importance Algorithms
  - Enhanced Adaptive Bayes Network
  - JSR-73 RI & TCK - Java Data Mining
  - Multi-User Access Control
  - Complete Oracle Data Mining Management
  - Enhanced Model Seeker
  - PMML Producers / Consumers
  - Enhanced Data Preprocessing
  - Data Mining Rapid Application Development
- **Bio-informatics**
  - DOUBLE and FLOAT Datatypes
  - BLAST Queries
  - Statistical Functions
  - Document Clustering
  - Document Clustering Non-negative Matrix Factorization
  - Document Clustering Neural Nets
  - Document Named Entity Extraction
  - Support Vector Machines
  - Taxonomy Builder
- **SQL Analytics**
  - Grouped Table Outer Join
  - Increased Number of Aggregates per Query
  - SQL Interrow Calculations
  - Upsert Through SQL Interrow Calculations

Oracle10i Security and Directory Features

- **Administration**
  - Extended SQL Support in FGA
  - Communication over SASL
  - Unified User Model
  - Easy Database Registration
  - Uniform Audit Trail
  - OLS Directory Integration
- **Secure Hosting**
  - VPD Support for Parallel Query
  - VPD Static and Dynamic Policies
  - Column-level VPD
- **Integration**
  - Certificate Validation with Certificate Revocation Lists (CRLs)
  - Centralized CRL Management
  - Centralized User Management for Kerberos Users
  - Kerberos-Enabled Database Links
  - Operating System Credential Cache
  - Access to Single Sign On Wallet
  - Key Store Interface
  - End-to-End Identity Propagation
- **Integration (Continued)**
  - Database Authentication with Standard LDAP Password Verifiers
  - Single Station Administration for Password Authentication to Oracle Database
  - Smart Card Support for X509v3 Certificates
  - Transport Layer Security (TLS) Support
  - SSL Session Renegotiation
  - Performance Improvements
  - 4096-Bit Key Size Support

Cool Thing #1

Automatic Storage Management (ASM)

Automatic Storage Management

New capability in the Oracle database kernel

- Provides a vertical integration of the file system and volume manager for simplified management of database files
- Spreads database files across all available storage for optimal performance
- Enables simple and non-intrusive resource allocation with automatic rebalancing
- Virtualizes the storage resources

Automatic Storage Management

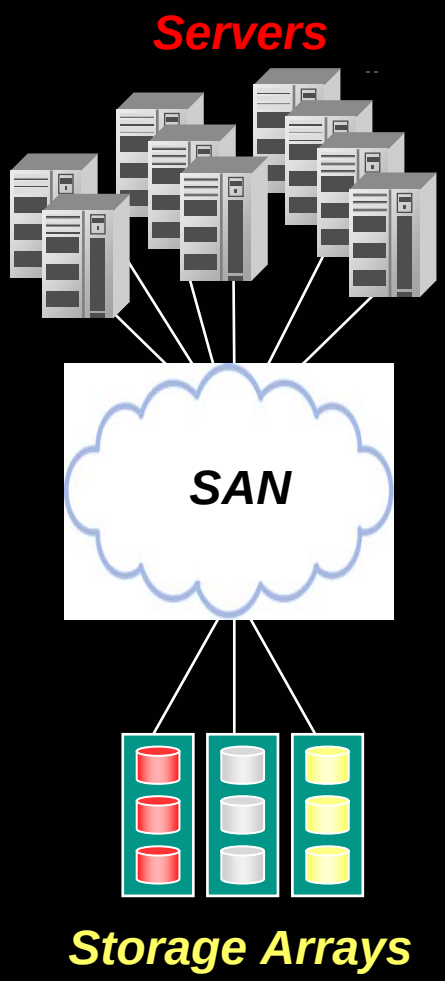

Application & DB Mgmt

File System Mgmt

Volume Mgmt

Zoning

LUN Masking

Assign Volumes Ports

Configure LUNs

This is where ASM fits into the stack

Operational Stack

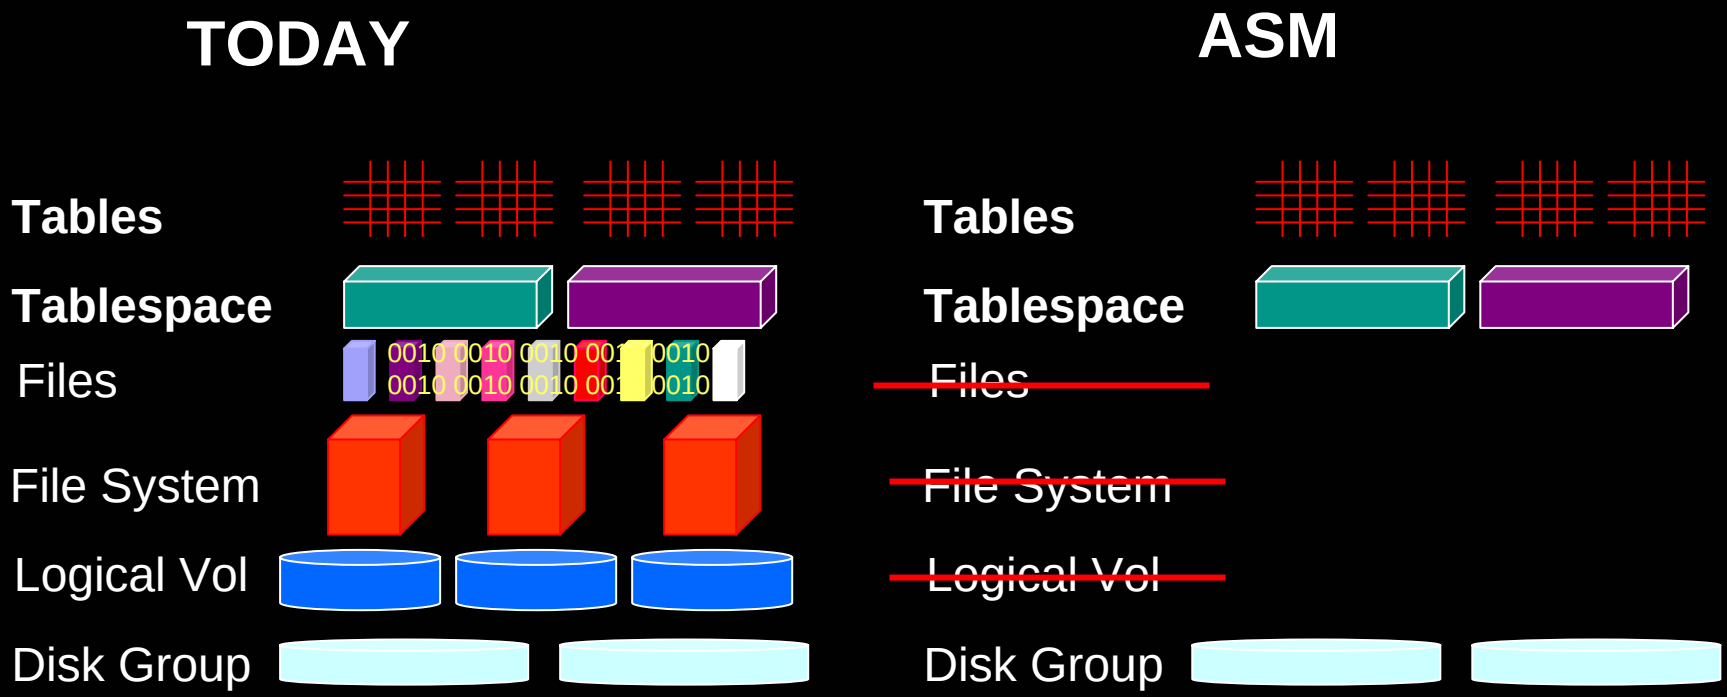

- NO mapping required
- Simpler to set up

ASM Extends SAME Concept

SAME = Stripe and Mirror Everything

- Spread I/Os across as many LUNs as possible
- Eliminates manual I/O tuning in static storage configurations

ASM = Automatic Storage Management

- Allows dynamic and online reconfiguration of storage resources
- Efficient relocation of data during rebalance
- Eliminate manual I/O tuning in all storage configurations

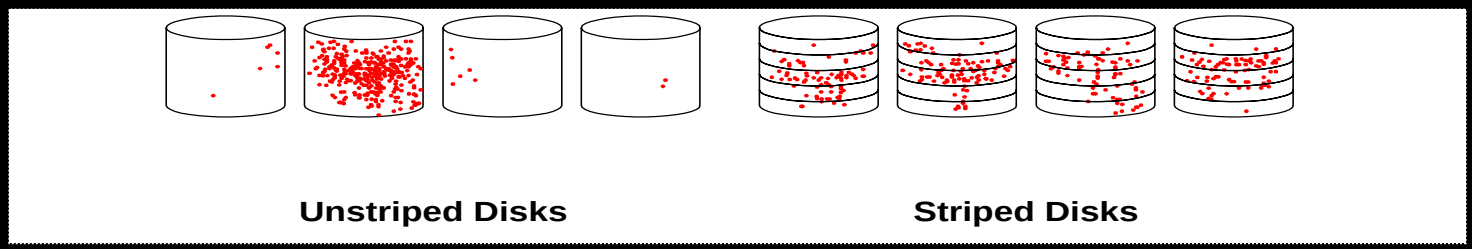

ASM Objectives

- Easier management of storage resources
- Optimize productivity of DBAs
  - Eliminate day to day I/O tuning tasks
  - Built in best practices
  - Eliminate database physical configuration layout
  - Manage disk groups not files
  - Non-intrusive storage resource addition/reduction
- Common storage management interface
  - Works the same on any OS
  - Same interface for any storage

ASM – How it Works

- A separate ASM “server” must be started before DB instances can access ASM files
- A Disk Group is a collection of disks managed as a logical unit in which ASM files are stored
- Every ASM file striped across all disks

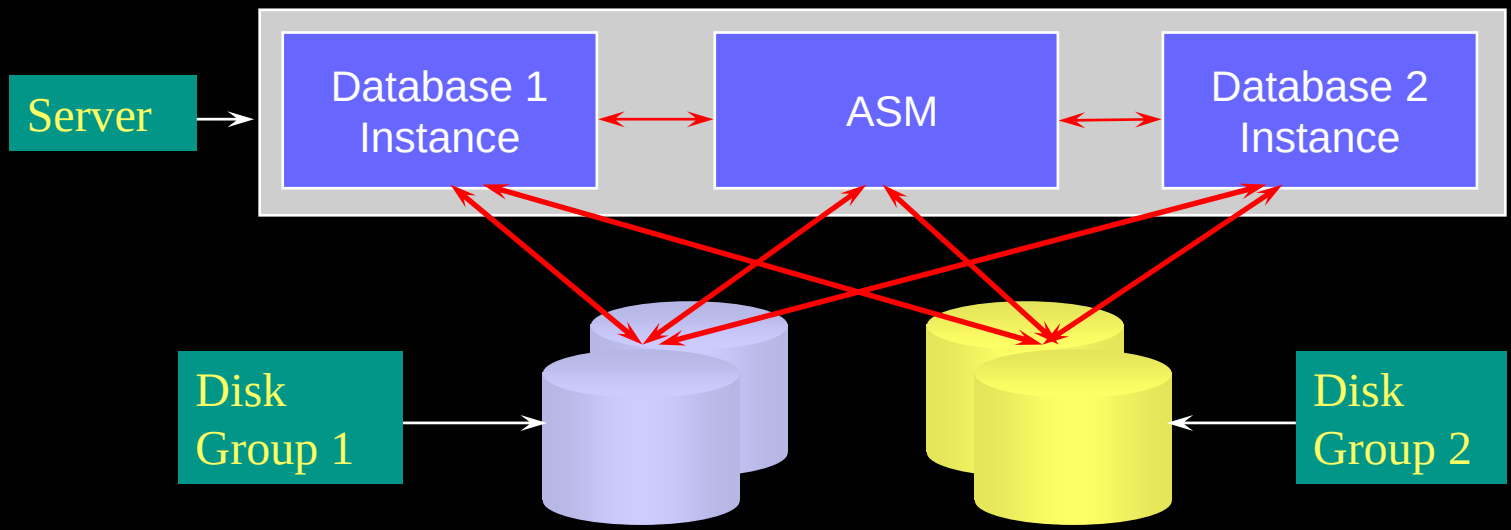

Disk Groups & Failure Groups

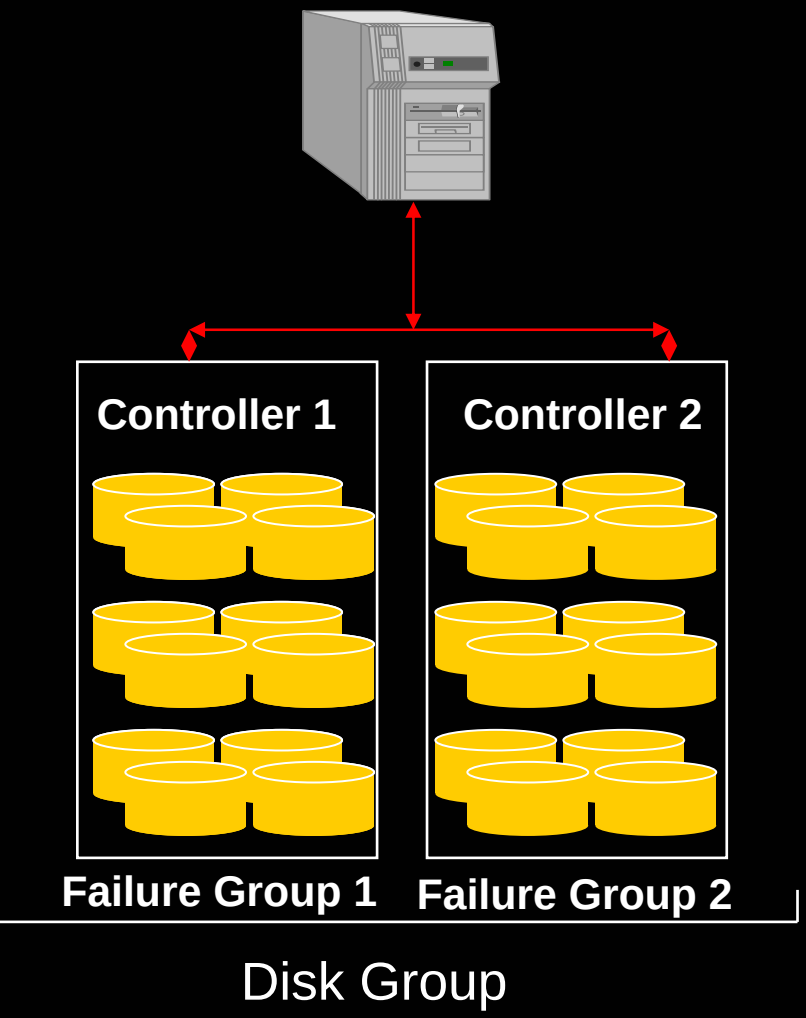

- Storage resources organized into Disk Groups and Failure Groups
- A *Disk Group* is a collection of disks managed as a logical unit in which ASM files are stored
- *Disk Groups* assigned by DBAs
- A *Failure Group* is a set of disks sharing a common resource whose failure needs to be tolerated
  - Redundant copies of an extent are stored in separate Failure Groups
- *Failure Groups* assigned by DBAs or automatically by ASM

ASM Functions Drill Down

- Getting started
- Defining a DB
- Adding/deleting/migrating storage
- Hot spots

Traditional vs ASM – Setup

1. Determine required storage capacity
2. Install and configure Volume Manager, File System
3. Architect data layout to avoid hot spot
4. Create logical volumes
5. Create file systems
6. Install Oracle software
7. Create database

1. Determine required storage capacity
2. Install Oracle software
3. Create Disk Groups
4. Create database

Traditional vs ASM – Add Disk

1. Add Disk to OS
2. Create volume(s) with Volume Manager
3. Create File System over volume
4. Figure out data to move to new disk
5. Move data to new files
6. Rename files in database
7. Re-tune I/O

1. Add Disk to OS
2. Issue the Add Disk command

Traditional vs ASM – Remove Disk

1. List all data that is on disk
2. Choose existing filesystem to hold data from dropped disk
3. Move data to new files
4. Rename files in database
5. Remove disk from OS
6. Re-tune I/O

1. Issue drop disk command
2. Remove the disk from OS when rebalance completes

Traditional vs ASM – Tune I/O

1. Daily monitor I/O performance
2. Discover hot spots
3. Figure out how to remove hot spots
4. Move data to new files
5. Rename files in database
6. Verify that hot spot is gone

This space intentionally left blank

ASM - Integration with Key Storage Hardware

- Vendor library for disk discovery and I/O
- Efficient I/O interface
- Data description allows end-to-end validation (HARD)
- Participating Storage Vendors
  - EMC
  - Network Appliance
  - Hitachi
  - HP
  - XIOtech

ASM Benefits

- Enables non-disruptive storage deployment
  - Database is always available
  - Online changes to storage resources (Add disks, remove disks and migrate to new storage)
- Eliminates manual performance tuning tasks
  - Automates database layout
  - Automates file naming
  - Obsoletes manual I/O tuning
- Increases DBA productivity
  - Reduces the number of objects to manage
  - Fewer interdepartmental dependencies
  - Manage more DB with same personnel


It’s all about flashing back…

Human Error

- Estimated to be the biggest single cause of downtime
- Need to quickly determine what happened and fix it
  - Localized damage
    - Needs surgical detection and repair
    - Example – removed wrong person named ‘Smith’
  - Widespread damage
    - Requires drastic action to avoid long downtime
    - Example – batch job deletes this month’s orders
- Analysis and correction using traditional recovery is slow and complex
  - Restore database to point in time and extract data
- Oracle10*g* is a breakthrough release for human error correction

Flashback Time Navigation

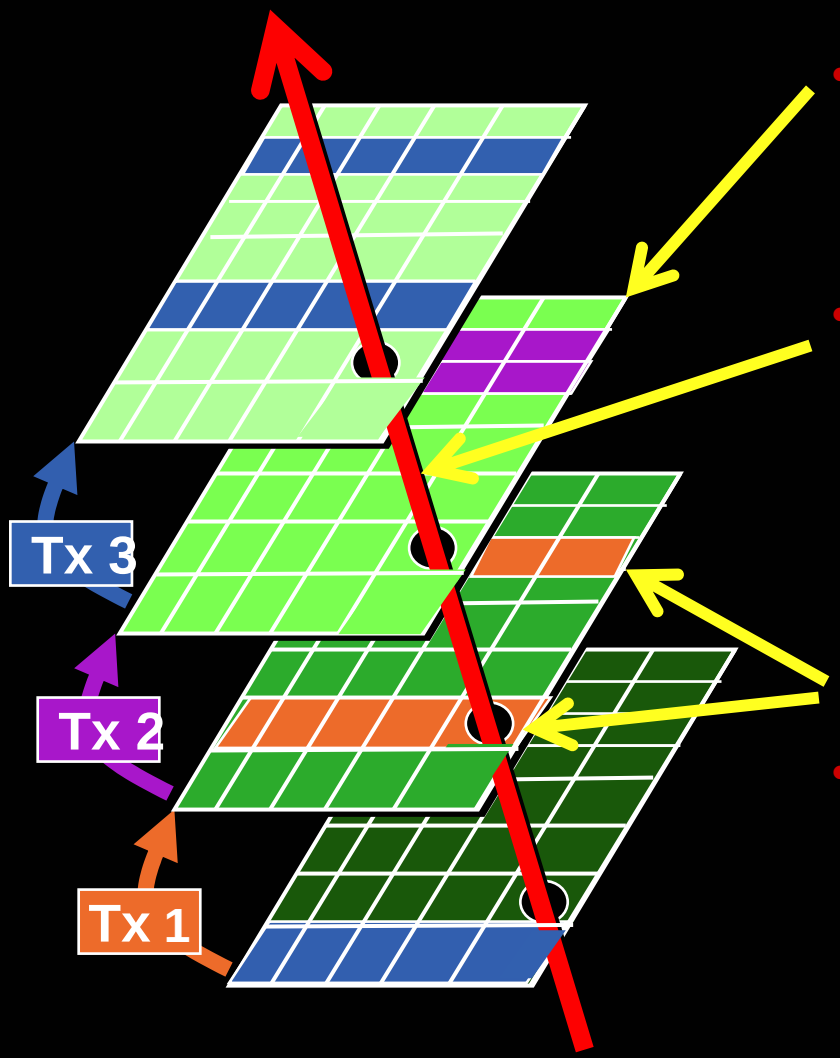

- **Flashback Query**
  - Query all data at point in time

```
Select * from Emp AS OF '2:00 P.M.' where …
```

- **Flashback Versions Query**
  - See all versions of a row between two times
  - See transactions that changed the row

```
Select * from Emp VERSIONS BETWEEN
 '2:00 PM' and '3:00 PM' where …
```

- **Flashback Transaction Query**
  - See all changes made by a transaction

```
Select * from DBA_TRANSACTION_QUERY
 where xid = '000200030000002D';
```

Flashback Database

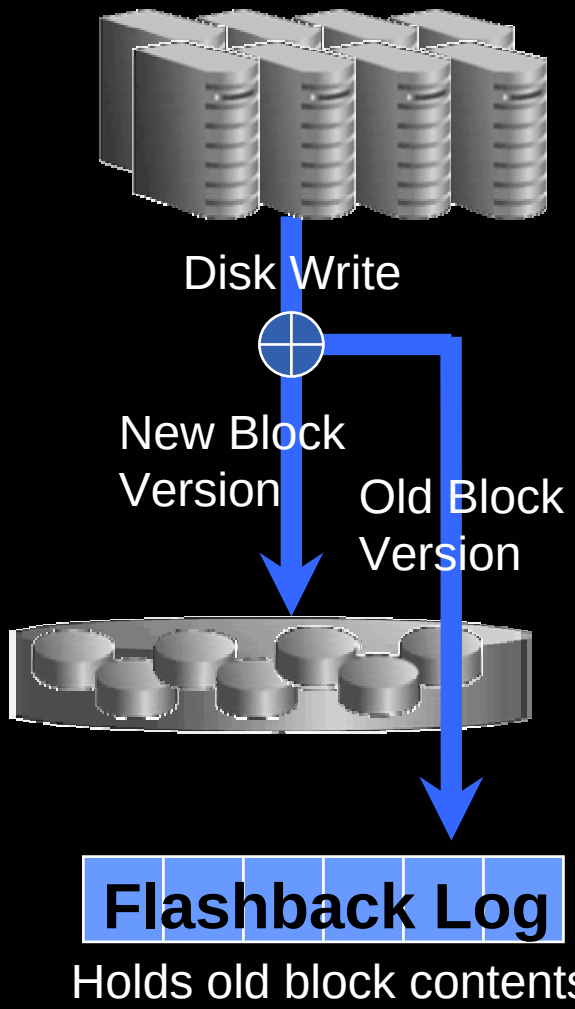

- A new strategy for point in time recovery
- Flashback Log captures old versions of changed blocks
  - Think of it as a continuous backup
  - Replay log to restore DB to time
  - Restores just changed blocks
- It's fast - recover in minutes, not hours
- It's easy - single command restore

  SQL> Flashback Database to '2:05 PM'

"Rewind" button for the Database

Flashback Error Correction

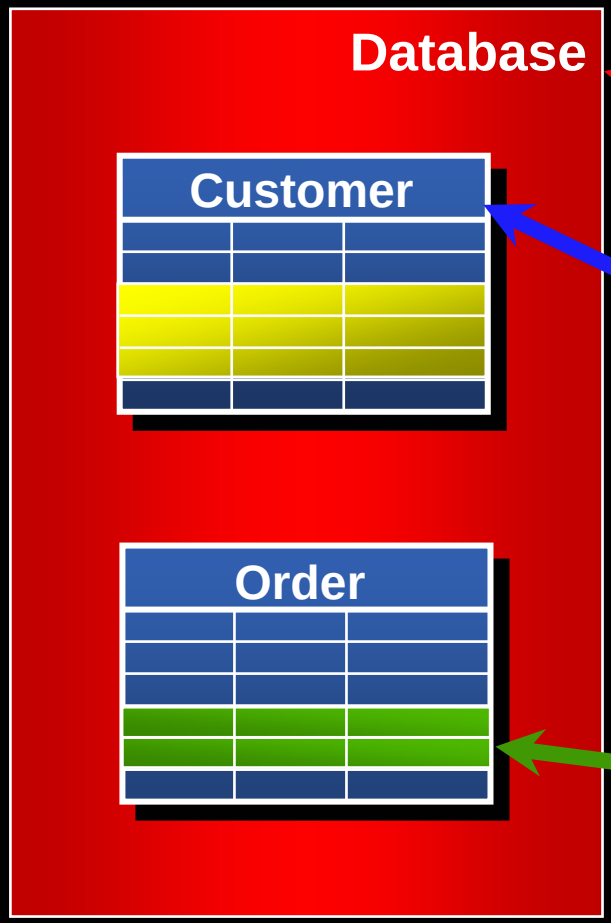

- Recovery at all levels
- Database Level
  - Flashback Database restores the whole database to time
    - Uses Flashback Logs
- Table Level
  - Flashback Table restores rows in a set of tables to time
    - Uses UNDO in database
  - Flashback Drop restores a dropped table or a index
    - Recycle bin for DROPs
- Row Level
  - Flashback Rows restores rows to time
    - Uses Flashback Query

Flashback for All Users

END USER
- Flashback Query
- Flashback Row History

DEVELOPER
- Flashback Row History
- Flashback Transaction History
- Flashback Table

DATABASE ADMIN
- Flashback Database
- Flashback Drop

SYSTEM ADMIN
- Data Guard

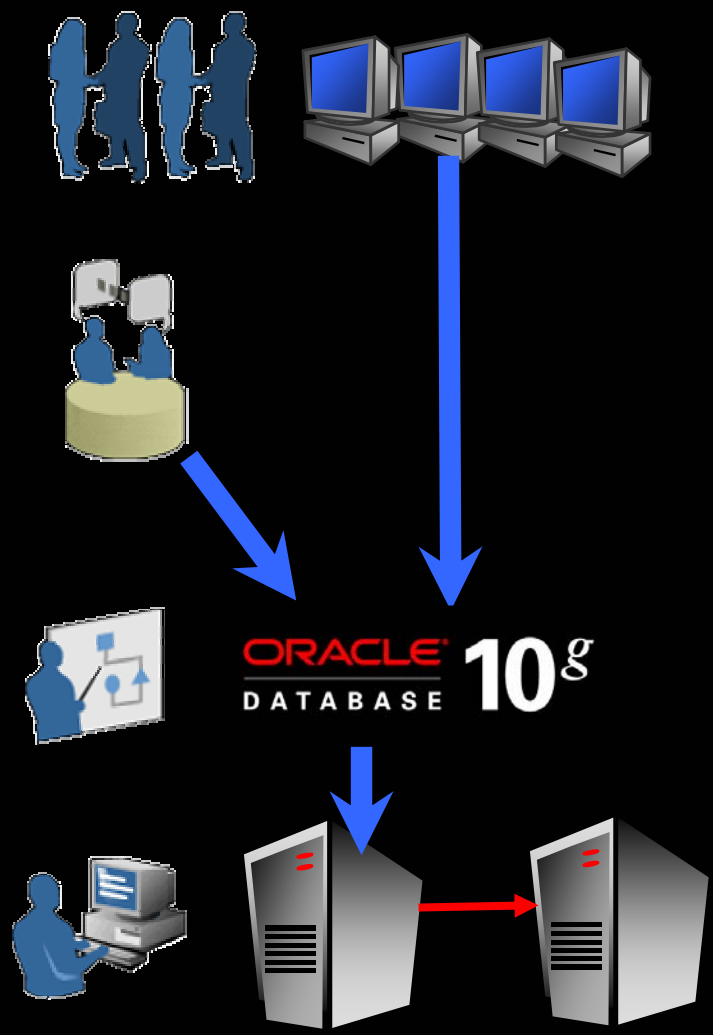

Cool Thing #3

PL/SQL just got faster… and more

What's new in PL/SQL in Oracle10^{g} Database ?

- Transparent and Semi-Transparent performance improvements
- New language features
- New Supplied Packages

Transparent performance improvements

- New optimizing compiler
  - Brand new code generator
  - New PVM instructions
  - Some front-end improvements
  - Native Floats
- NATIVE Compilation
  - Configuration simplified
  - Interoperable with RAC
  - Interoperable with backup

```
procedure P is
  a positiven := 1 /* binary_integer subtypes */;
  b positiven := 1; c positiven := 1; d positiven:=1;
  t0 integer;
begin
  t0 := Dbms_Utility.Get_Time();
  for j in 1..1000000
  loop
    b := j /* pls_integer to binary_integer conversion */;
    d := 42 /* constant assignment within a loop */;
    c := d + b /* can be combined... */;
    a := b + c /* ...c not used except here */;
  end loop;
  Dbms_Output.Put_Line(Dbms_Utility.Get_Time()-t0);
end P;
```

Performance comparison

Environment	Time *(sec)*
9*i*R2 interpreted	11.5
9*i*R2 native	5.2
10*i*R1 interpreted	3.4
10*i*R1 native	1.2

Program to extract unique words
and their frequencies from text field

Tested on the complete Shakespeare plays, line by line

VERSION	CODE_TYPE	OPTIMIZER_LEVEL	AVG	STDDEV	Count
10iR1	native	opt_level_2	41.45	.06	5
10iR1	native	opt_level_1	45.77	.13	5
9iR2	native	n/a	59.51	.32	5
10iR1	interpreted	opt_level_2	62.76	.72	5
10iR1	interpreted	opt_level_1	68.21	.68	5
9iR2	interpreted	n/a	91.56	.56	5

What's new in PL/SQL in Oracle10*i* Database ?

- (Semi)transparent performance improvements
- New language features
- New Supplied Packages

New language features

- “Inherited” new SQL features
  - IEEE arithmetic
  - Regular expressions
  - Multiset operations
  - User defined quote character

```
insert into t values
  ( q'|Don't do that.|' )
```

$\pi / 4 = 1 - 1/3 + 1/5 - 1/7 + 1/9...$

```
create or replace function pi return number
as
    subtype my_number is number;
    last_pi my_number := 0;
    delta   my_number := 0.00000001;
    pi      my_number := 1;
    denom   my_number := 3;
    oper    my_number := -1;
    negone  my_number := -1;
    two     my_number := 2;
begin
    loop
        last_pi := pi;
        pi := pi + oper * 1/denom;
        exit when (abs(last_pi-pi) <= delta );
        denom := denom + two;
        oper := oper * negone;
    end loop;
    return pi * 4;
end;
```

```
create or replace function pi return number
as
    subtype my_number is BINARY_DOUBLE;
    last_pi my_number := 0;
    delta   my_number := 0.00000001;
    pi      my_number := 1;
    denom   my_number := 3;
    oper    my_number := -1;
    negone  my_number := -1;
    two     my_number := 2;
begin
    loop
        last_pi := pi;
        pi := pi + oper * 1/denom;
        exit when (abs(last_pi-pi) <= delta );
        denom := denom + two;
        oper := oper * negone;
    end loop;
    return pi * 4;
end;
```

IEEE arithmetic

- Takes 5 million iterations to converge to the specified limit of change
- The *number* version takes approx 410 sec
- The *binary_double* version takes approx 86 sec *(everything else the same)*
- Opens up a new range of numerical programming for PL/SQL

Data Type	Time *(sec)*
number	410.01
Binary_double	85.69

New language features

- Bulk binding of sparse arrays (FORALL syntax)
  - INDICES OF when the bound array is sparse
  - VALUES OF when a helper array picks out elements from the bound array
- Scenario
  - Read from a table (*eg* submitted but unvalidated orders)
  - Analyze and split into two new tables (*eg* ready-to-dispatch, get-back-to-customer)

New language features

- Compiler warnings
  - Severe, Performance and Informational categories
  - Can turn off/on by category
  - Can turn off/on by individual warning
  - *SP2-0804: Procedure created with compilation warnings*
  - *PLW-07203: parameter 'IO_TBL' may benefit from use of the NOCOPY compiler hint*

What's new in PL/SQL in Oracle10*i* Database ?

- (Semi)transparent performance improvements
- New language features
- New Supplied Packages

New Supplied Packages

- Utl_Mail
  - Its API mirrors the mental model of the end user of eg Netscape Communicator
  - No knowledge of SMTP protocol required (compare with using Utl_Smtp)
  - Some restrictions, to ensure simplicity
- Utl_Compress
  - Provides an API to compress and uncompress either a RAW or a BLOB (server side)

New Supplied Packages

- Utl_DBWS
  - Database Web Services
  - Expose/interact with all PL/SQL as a web service
- Utl_I18N
  - Help developers build multilingual applications
  - Conversion functions for character sets
  - Mappings of Oracle Character sets to IANA (Internet Assigned Numbers Authority)
  - And more
- Just to name a few

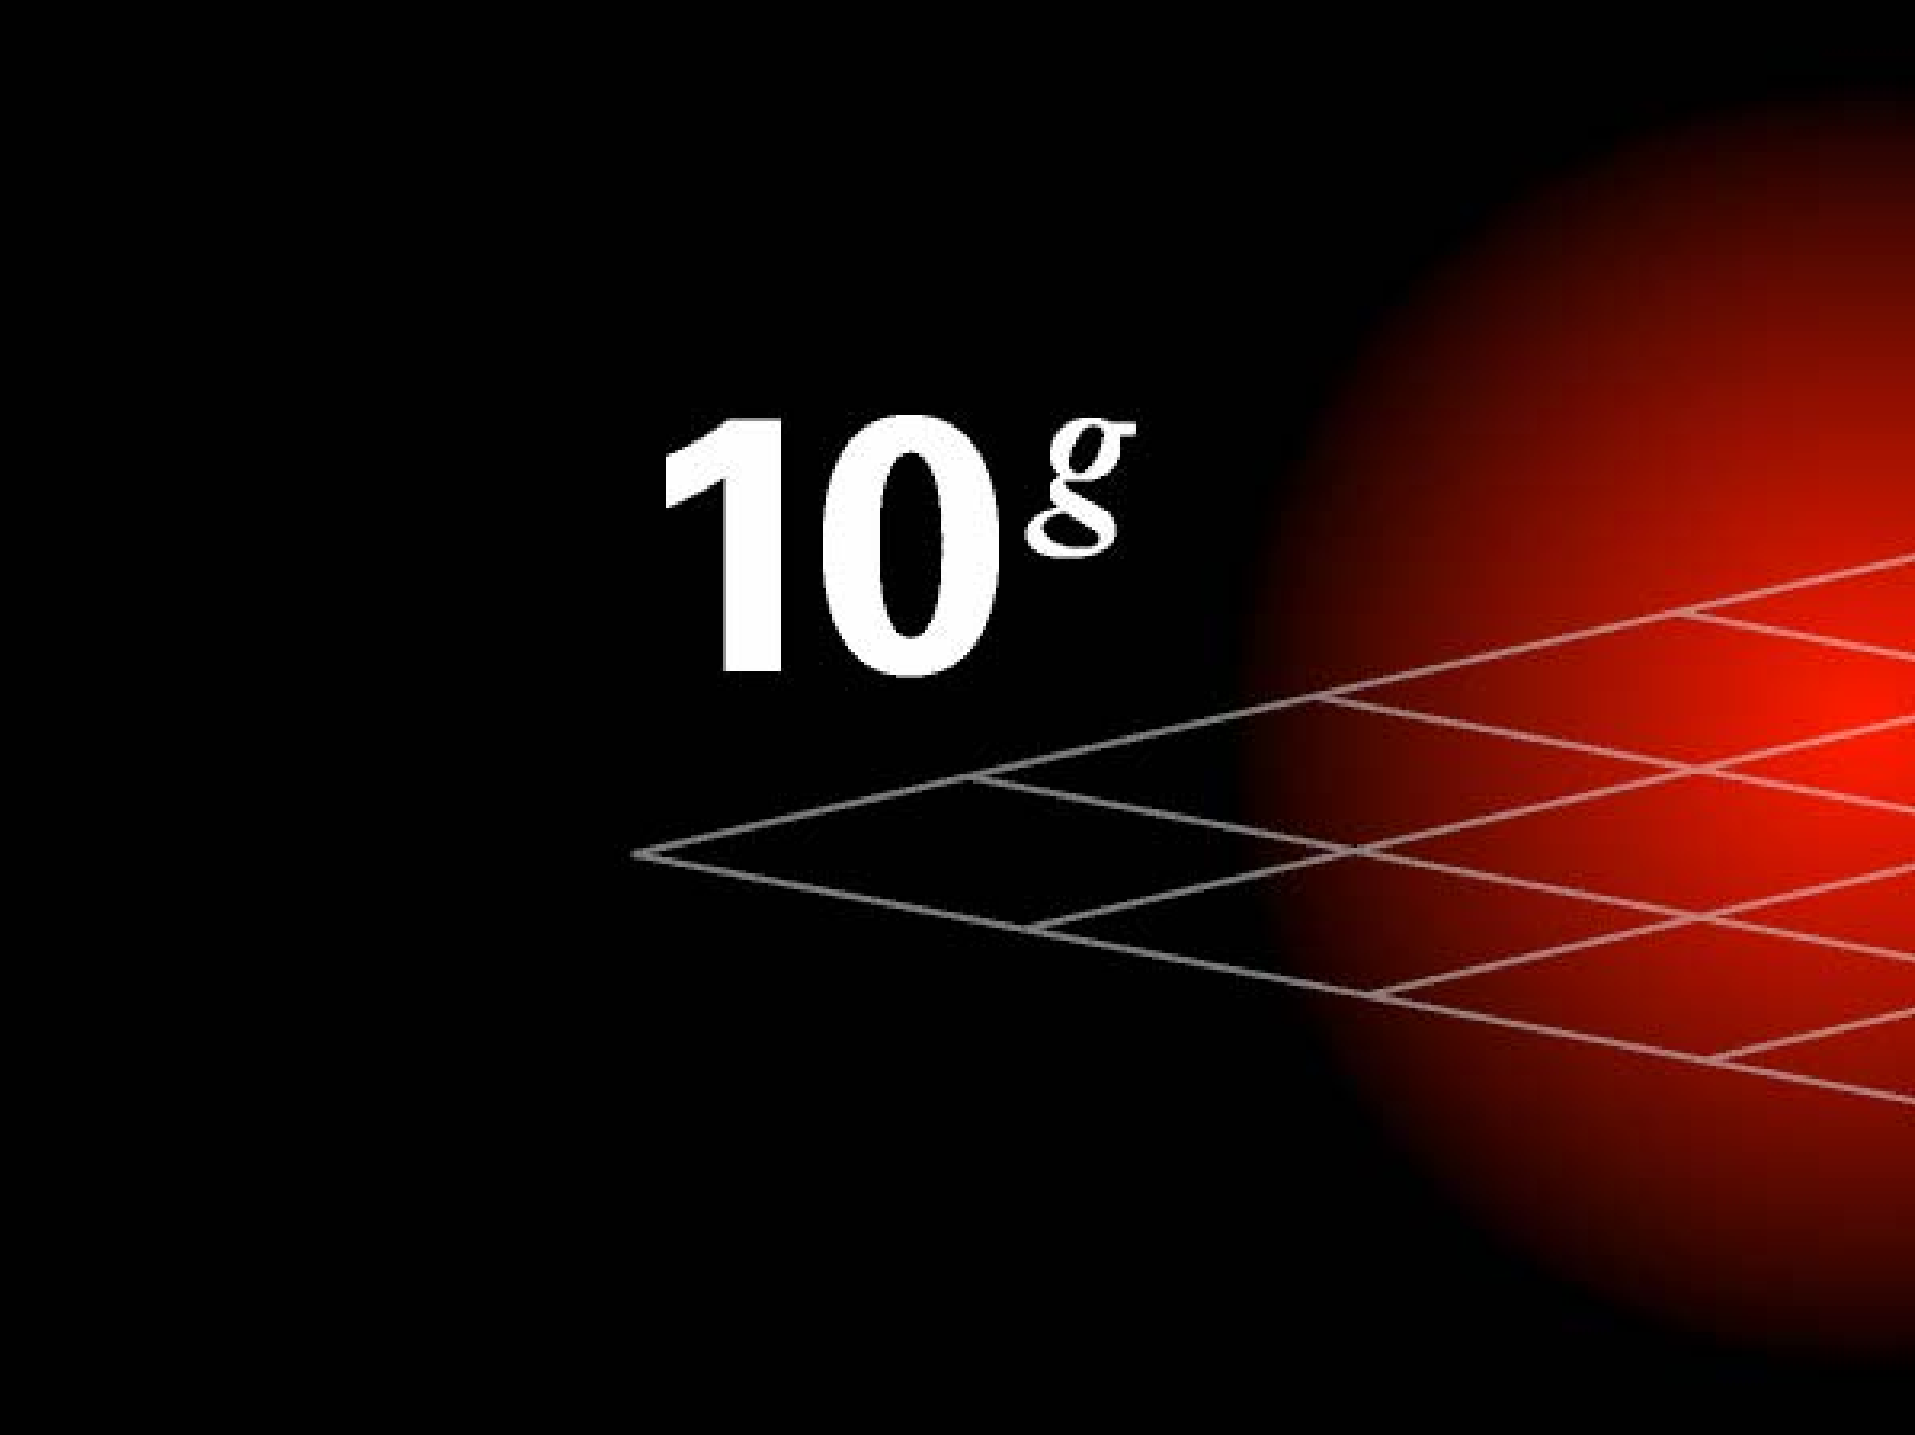
10g


QUESTIONS
ANSWERS