

ORACLE®

Oracle Database 11g Release 2

Major New Features

Costantinos Bourboulas

Master Principal Sales Consultant

Oracle Database 11g Release 2

Specific Areas


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

Grid Computing



Real Application Clusters


Virtualizes server resources



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

Server Pools


Dynamic Cluster Partitioning



- Server Pools
 - Dynamically assigns the server resources required to run specific workloads
- Both Application and Database Pools
- Policy Managed
 - Min and Max Servers
 - Relative Importance
- Unassigned Servers go to Free Pool

Server Pools


Dynamic Cluster Reconfiguration for High Availability



- If a server pool falls below its minimum the cluster reconfigures

Server Pools

Dynamic Cluster Reconfiguration for High Availability



- If a server pool falls below its minimum the cluster reconfigures
- Moves a server from
 - A server pool that is less important
 - A server from pool with the same importance with more servers than its min

Server Pools


Simplified Provisioning



- New intelligent installer
 - 40% fewer steps to install RAC
- Nodes can be easily repurposed
 - Network and storage information read from profile and configured automatically
 - No need to manually prepare a node
 - Nodes can be dynamically added or removed from the cluster
- SCAN - Single cluster-wide alias for database connections
 - Clients need not be aware of cluster configuration changes

Server Pools

Better Virtualization for Databases



- RAC One extends benefits of server virtualization to single-instance databases on physical hardware
 - Consolidation
 - Live Migration
 - Rolling Patches
 - Server failover
 - Standardized DB environment
- Online upgradeable to full RAC

Server Pools

Resources protected with Oracle Clusterware



SOA
IMDB
OC4J
APACHE

- New agents for IMDB Cache, Hyperion, Siebel
- Resource Dependencies can be specified
 - Hard / Weak / Attraction / PullUp / Dispersion

Start Dependencies

This section specifies a set of relationships that govern the start of the resource.

Hard	ApacheVIP Dependent resources must be ONLINE before the resource can be started
Weak	n/a An attempt will be made to start dependent resources if they are not ONLINE. Failure to start dependent resources will have no impact on starting this resource
Attraction	n/a Resource will be started preferably on the servers where dependent resources are ONLINE
Pullup	n/a All dependent resources will be auto started before starting the resource
Dispersion	active:n/a Resource will try not to be co-located with the dependent resources


Stop Dependencies

This section specifies a set of relationships that govern the stop of the resource.

Hard	ApacheVIP Resource will be stopped if one or more dependent resources go into INTERMEDIATE, OFFLINE or UNKNOWN state
------	--

Automatic Storage Management


Lowers the cost of storage management



- Virtualize and share storage resources
- Advanced data striping for maximum I/O performance
- Online addition and migration of storage

ASM Cluster File System


ASM Supports All File Types



- ASM now supports all types of files
 - Database files
 - Shared Clusterware files (OCR/Vote Disk)
 - File system files
 - ASM Cluster File System (ACFS)
 - 3rd-party file systems

ASM Cluster File System

ASM Supports All File Types




- General purpose clustered or local file system
- Data stored in ASM
- Inherits all ASM manageability benefits
 - Optimized disk layout
 - Online disk add/drop/rebalance
 - Integrated mirroring
- Read-Only Snapshots
 - up to 64 point-in-time space efficient copies of file system

Automatic Storage Management

Other ASM Enhancements


- Improved Management
 - ASM Install & Configuration Assistant (ASMCA)
 - Full Featured ASMCMD
 - ASM File Access Control
- Tunable Performance
 - Intelligent Data Placement



High Availability



Oracle Maximum Availability Architecture



Oracle Data Guard

Enhancements

- Active Data Guard - Quality of Service
 - Automatic block repair
 - Standby Query SLA configurable from zero to “n” seconds
- Data Guard - network utilization
 - Transport compression
 - Not just to resolve gaps after propagation interruptions

Eliminating Downtime

Online Application Upgrade

- Large, mission critical applications are often unavailable for tens of hours while a patch or an upgrade is installed
- Oracle Database 11g Release 2 introduces revolutionary new capabilities that allow online application upgrade with uninterrupted availability of the application
- The pre-upgrade application and the post-upgrade application can be used concurrently

Eliminating Downtime

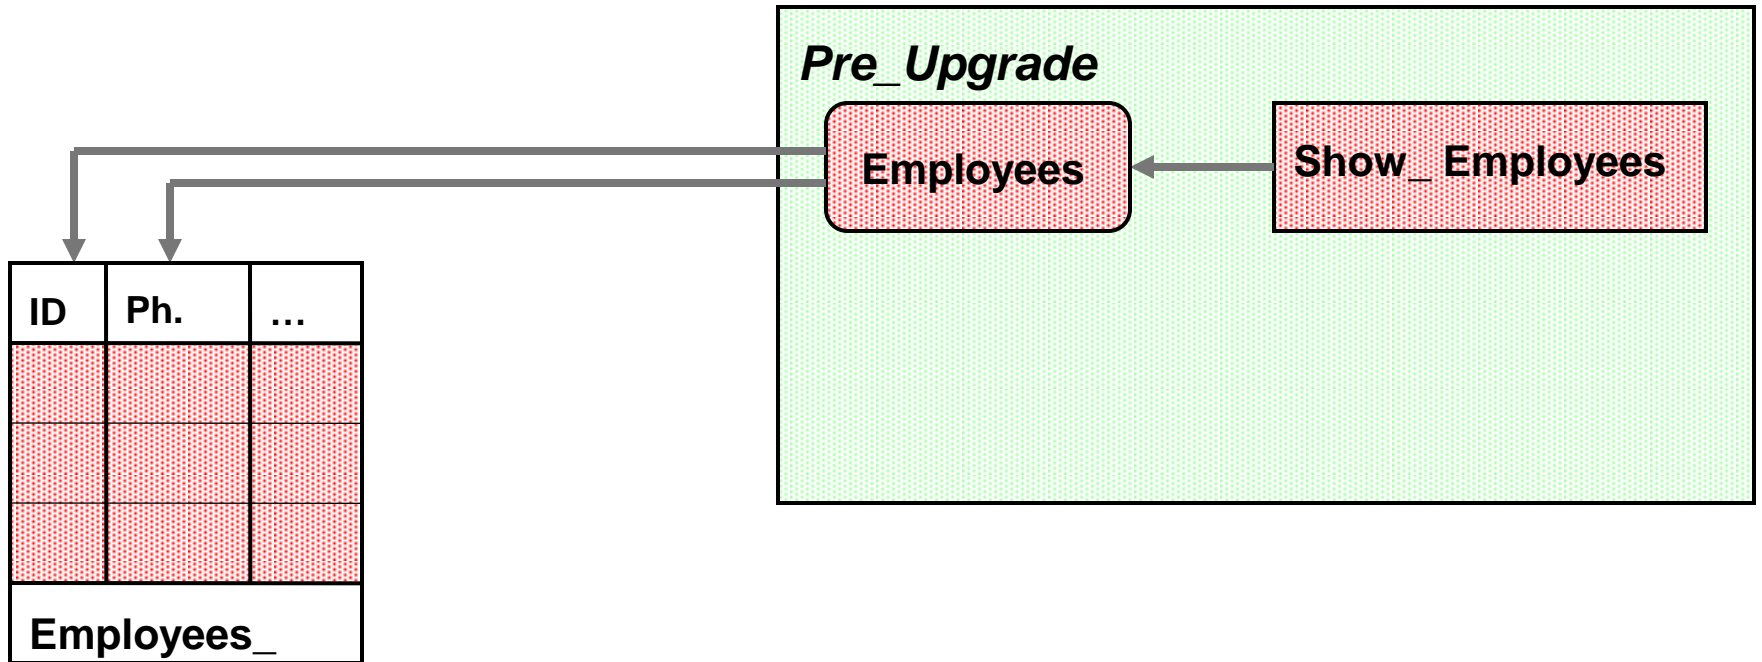
Edition-based redefinition

- Revolutionary new features:
 - Editions
 - Editioning View
 - Crossedition Trigger
- Code changes are installed in the privacy of a new edition
- Data changes are made safely by writing only to new columns or new tables not seen by the old edition
 - An **Editioning View** exposes a different projection of a table into each edition to allow each to see just its own columns
 - A **Crossedition Trigger** propagates data changes made by the old edition into the new edition's columns, or (in hot-rollover) vice-versa

An example

- The HR sample schema, as shipped by Oracle Corp, has a single column for phone number
 - Diana Lorentz 590.423.5567
 - John Russell 011.44.1344.429268
- It must be split into two new columns
 - Country code
 - Number within country
 - Diana Lorentz + 1-590-423-5567
 - John Russell + 44-134-442-9268

Starting point.
Pre-upgrade app in normal use.



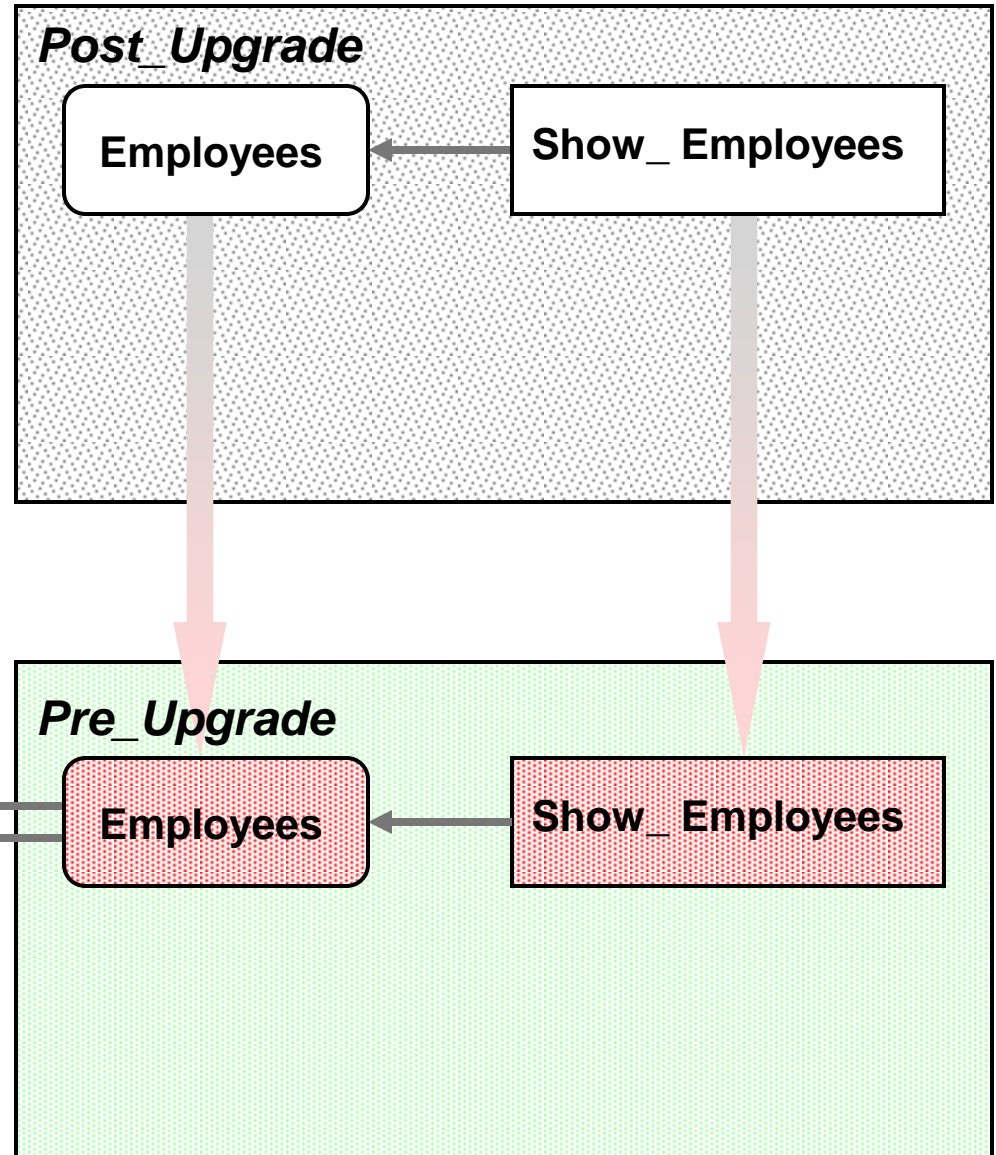
Start the edition-based redefinition exercise.

Create the new edition as the child of the existing one.

This is fast because initially all the editioned objects are just inherited.

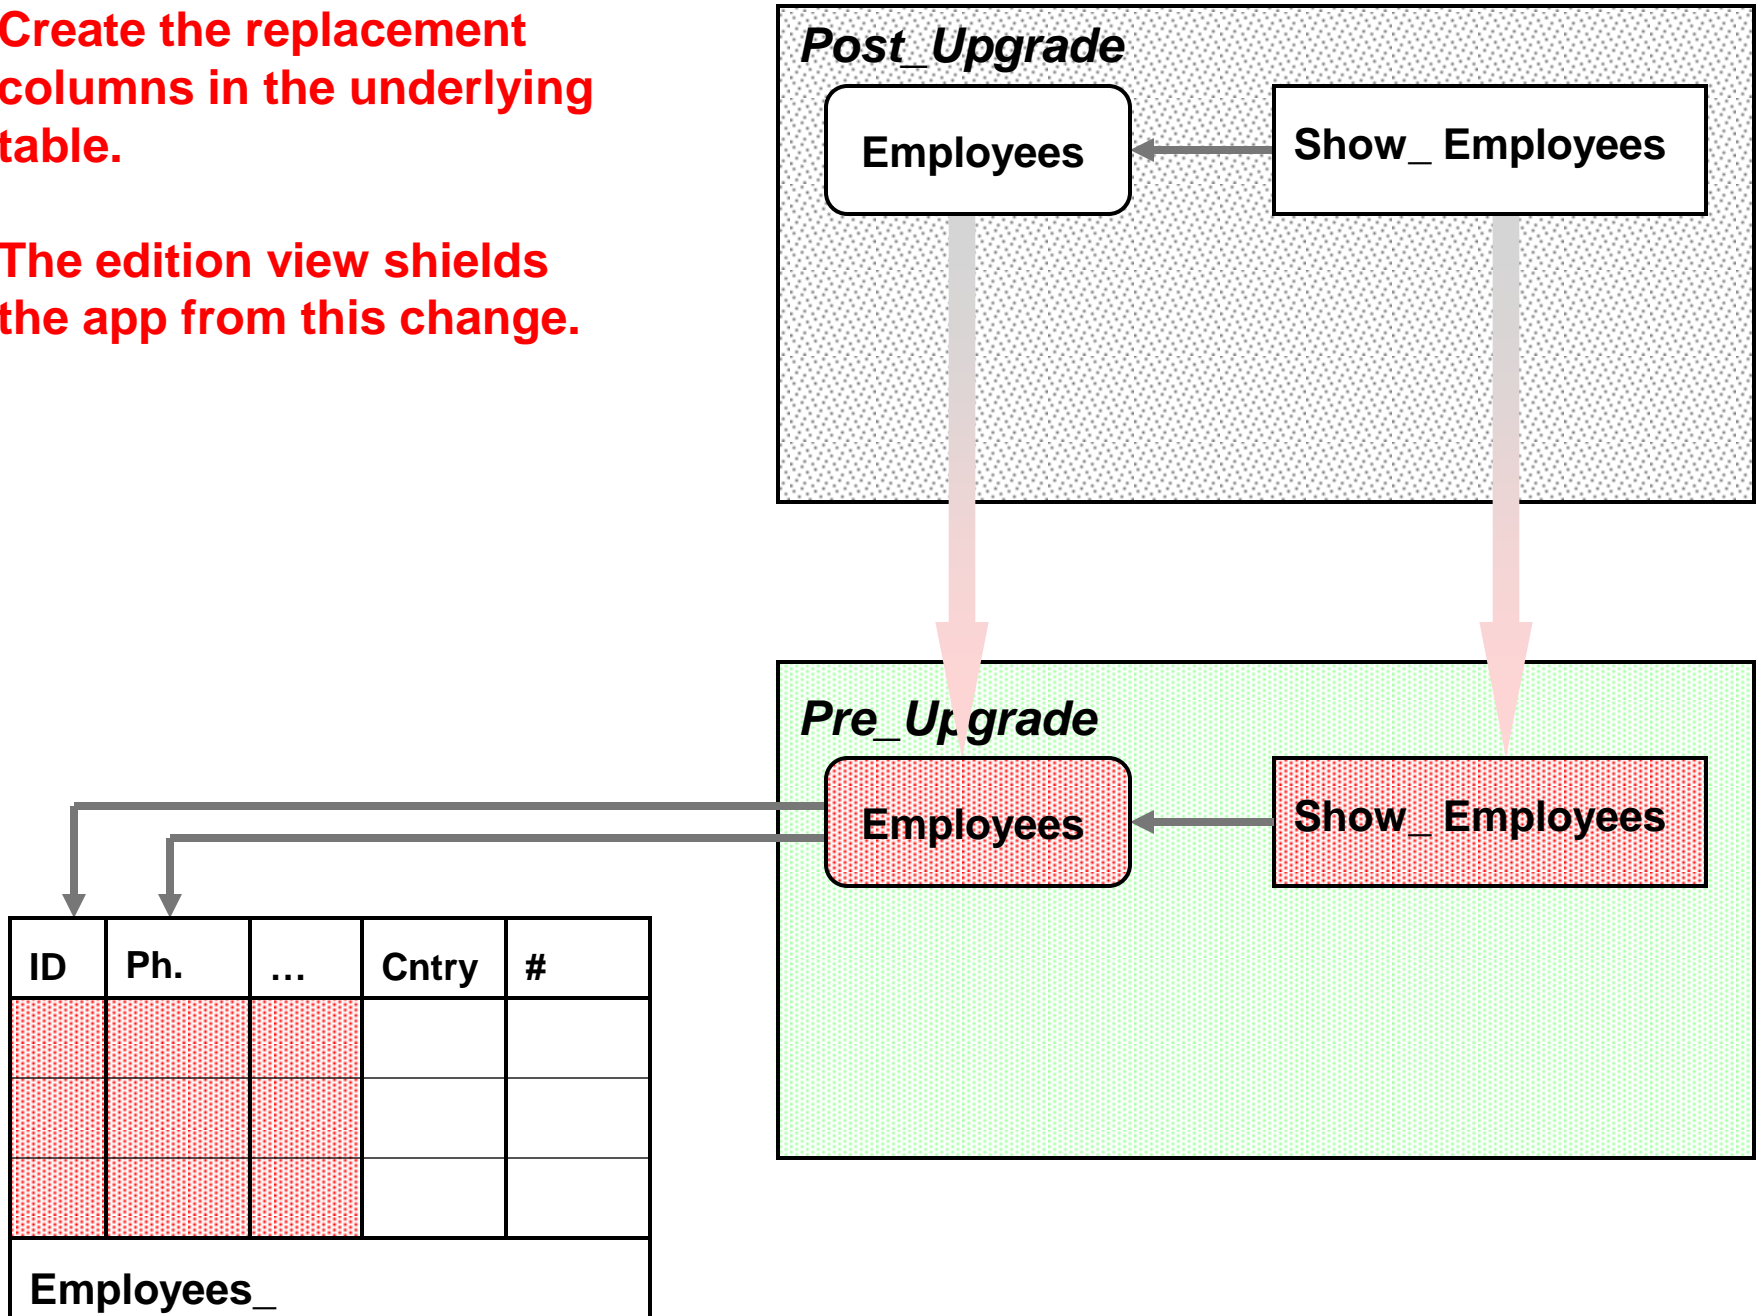
ID	Ph.	...

Employees_



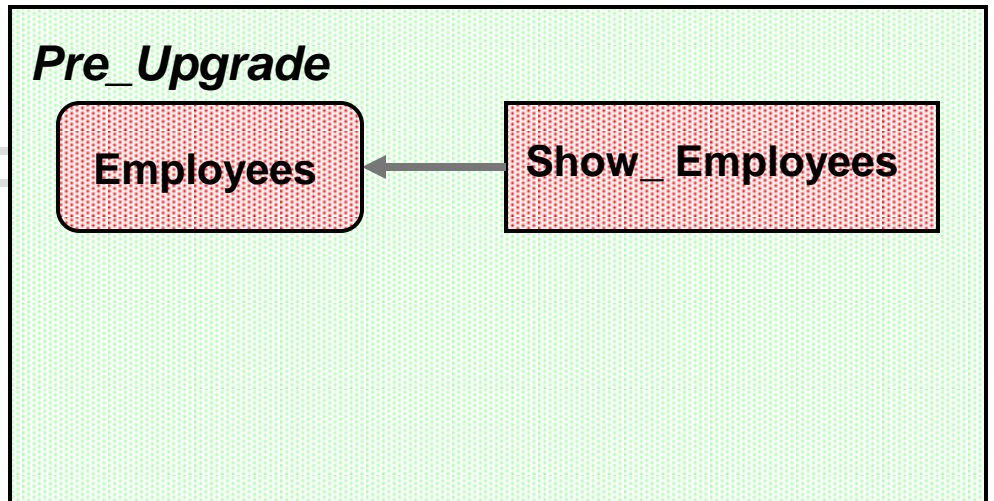
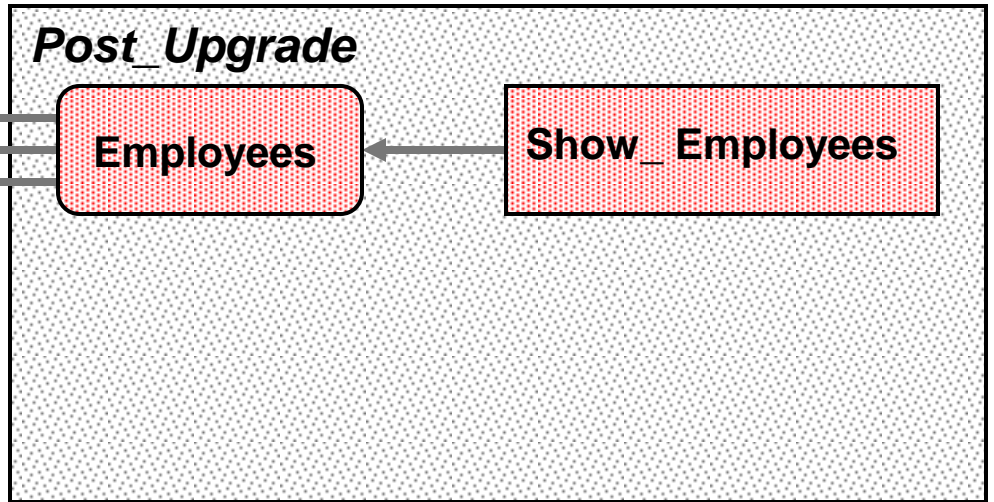
Create the replacement columns in the underlying table.

The edition view shields the app from this change.



Change *Employees* to select the new columns.

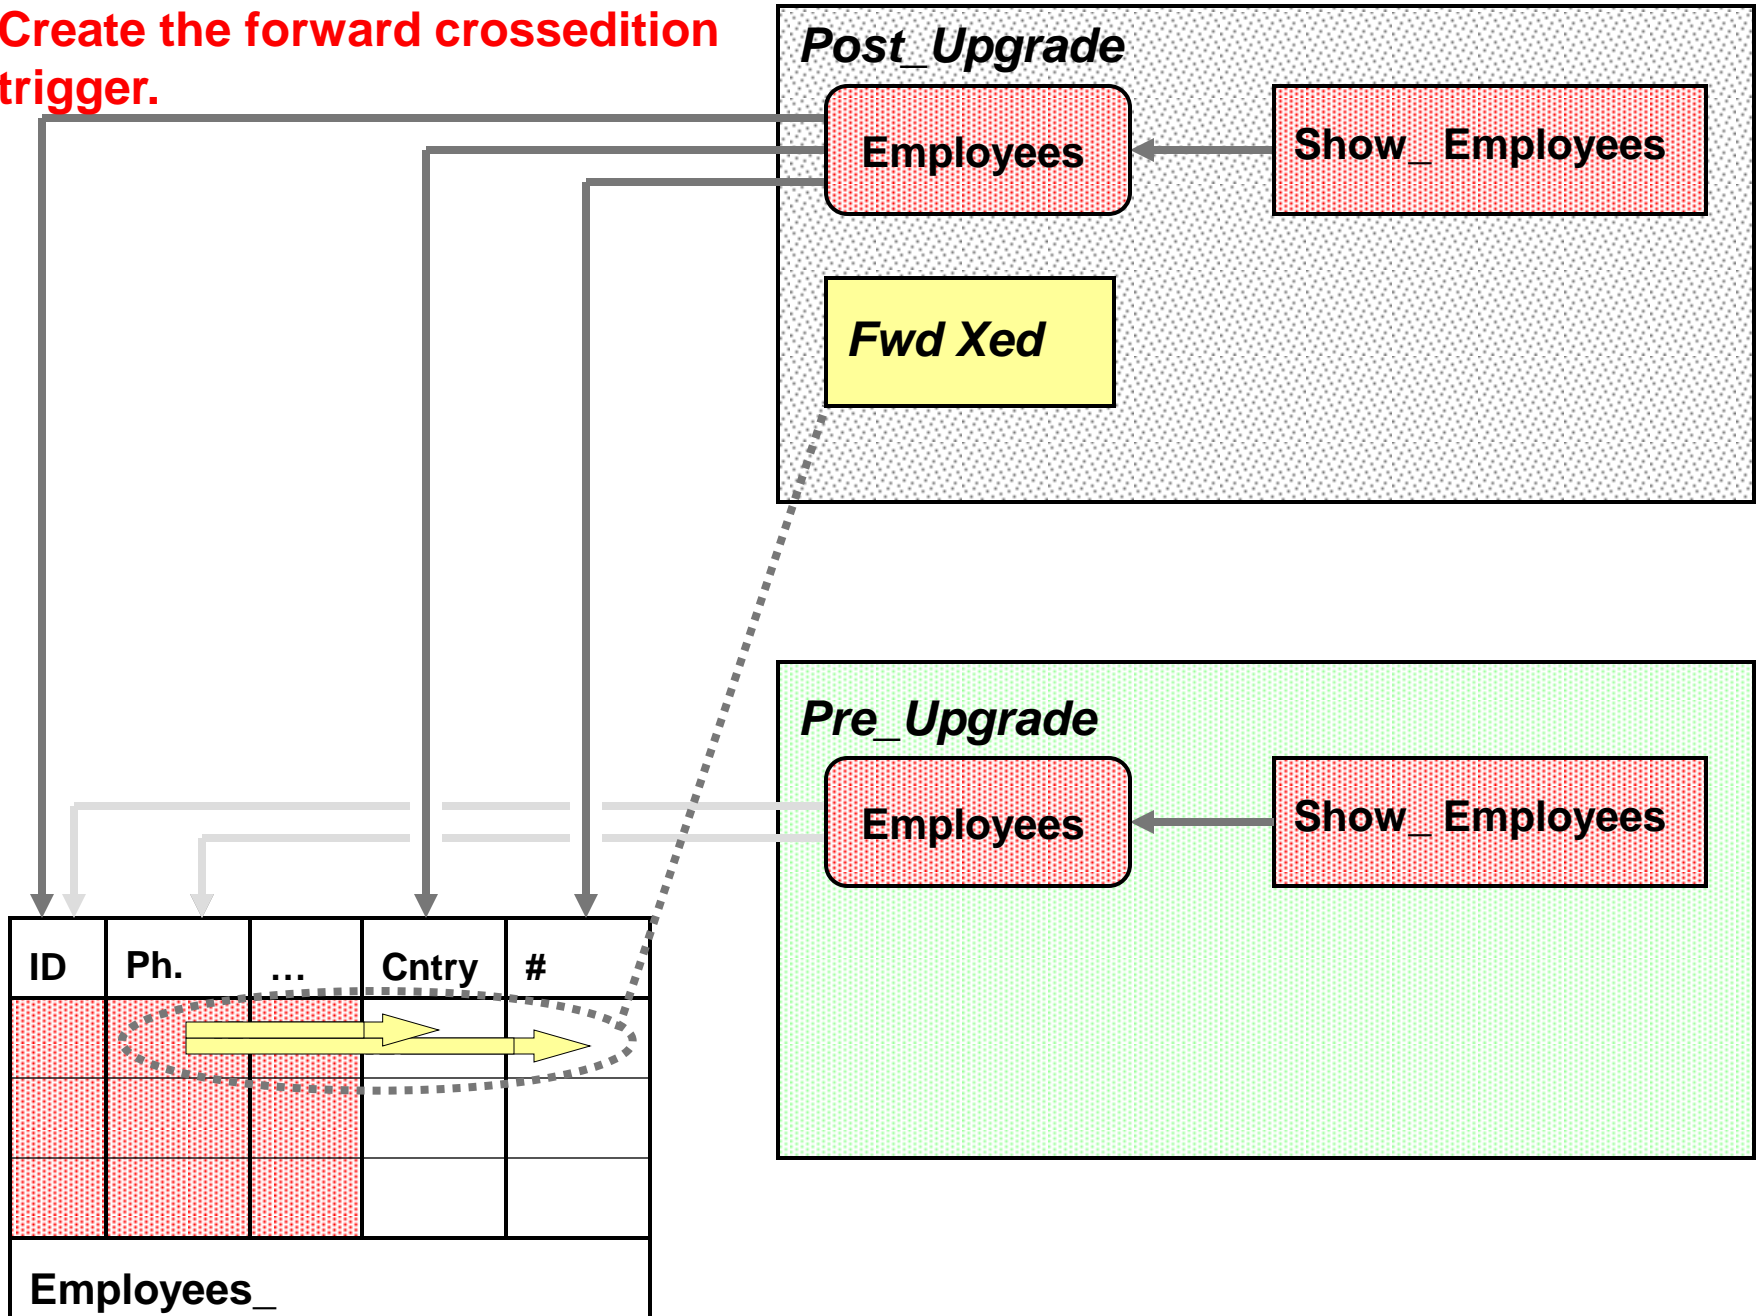
Change *Show_Employees* to implement the new behavior.



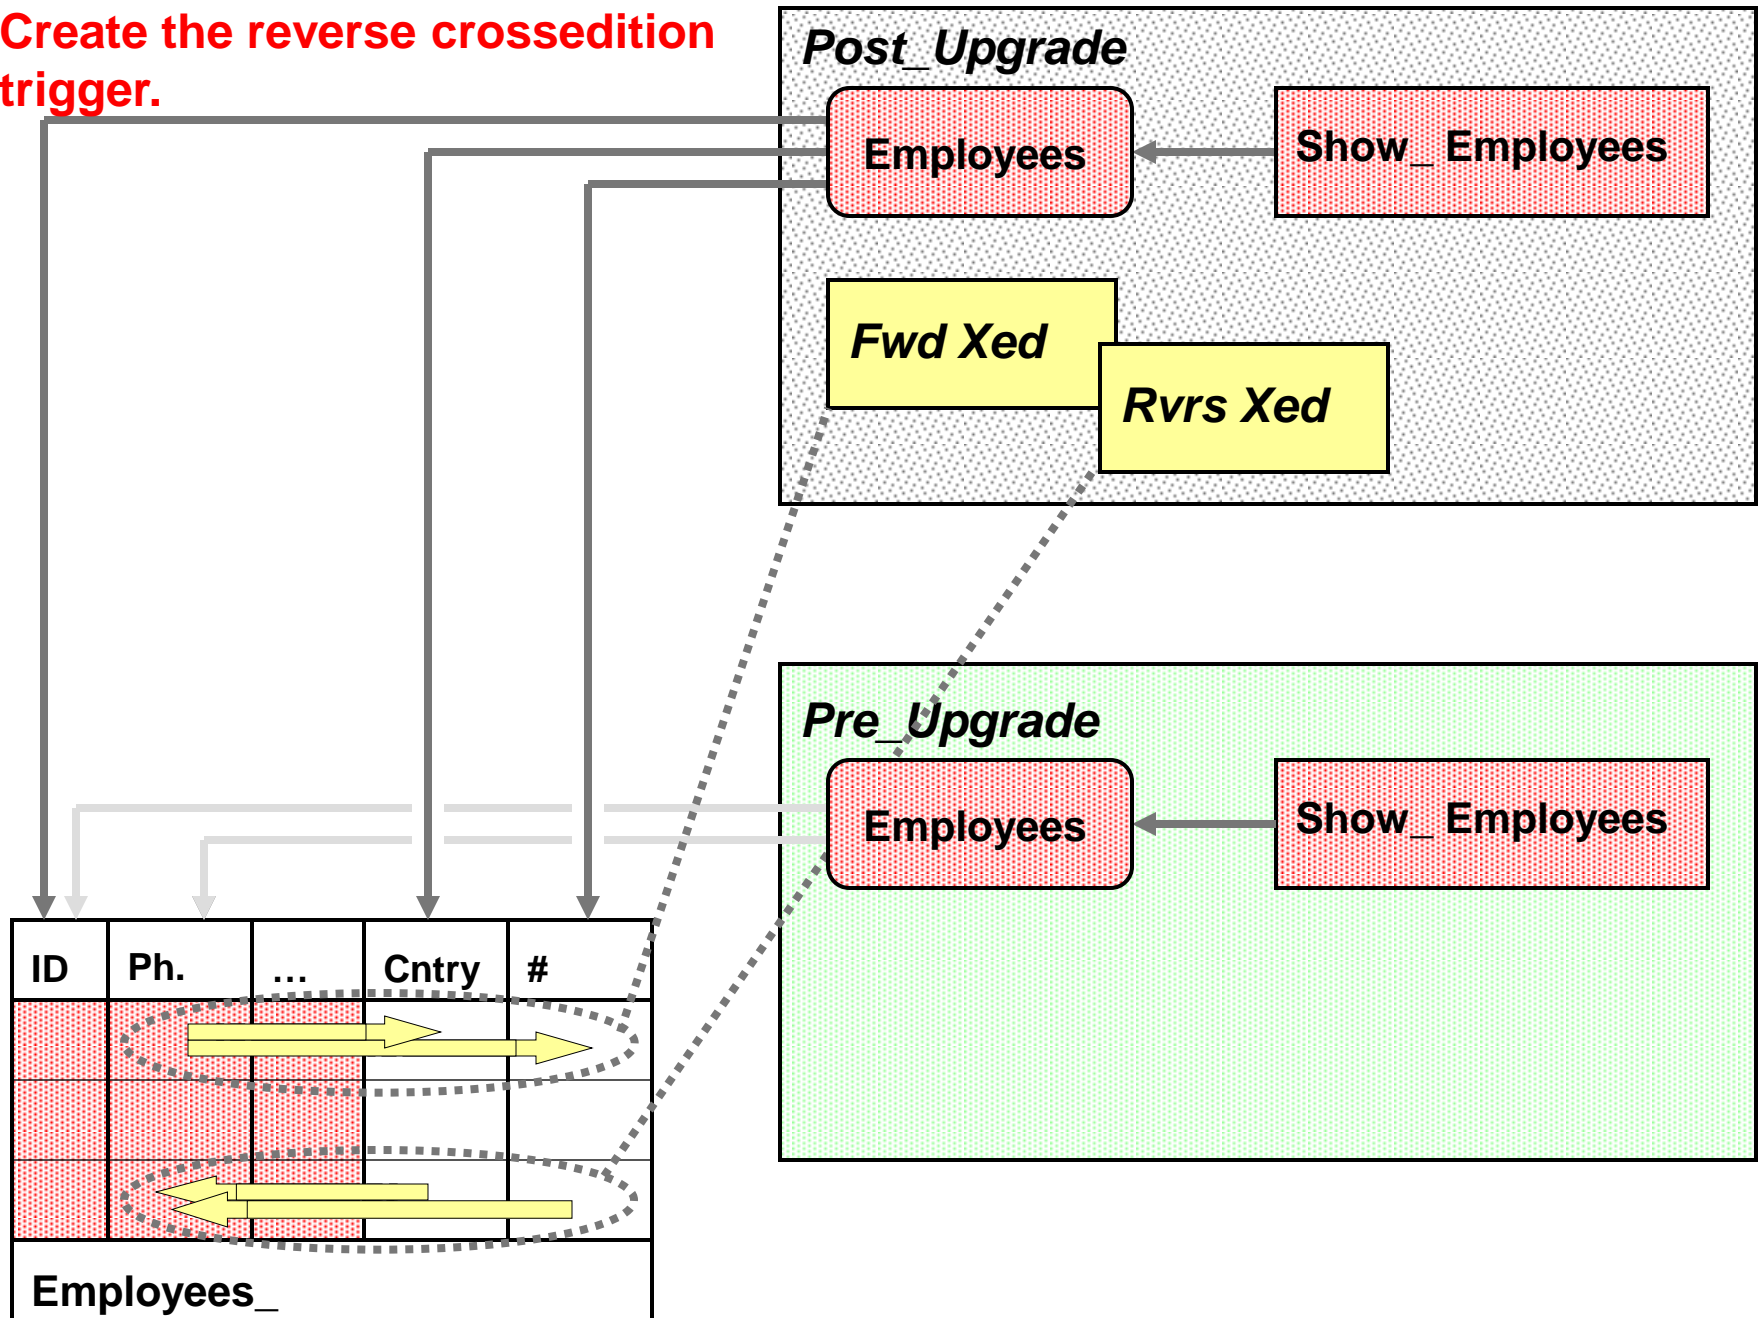
ID	Ph.	...	Cntry	#

Employees_

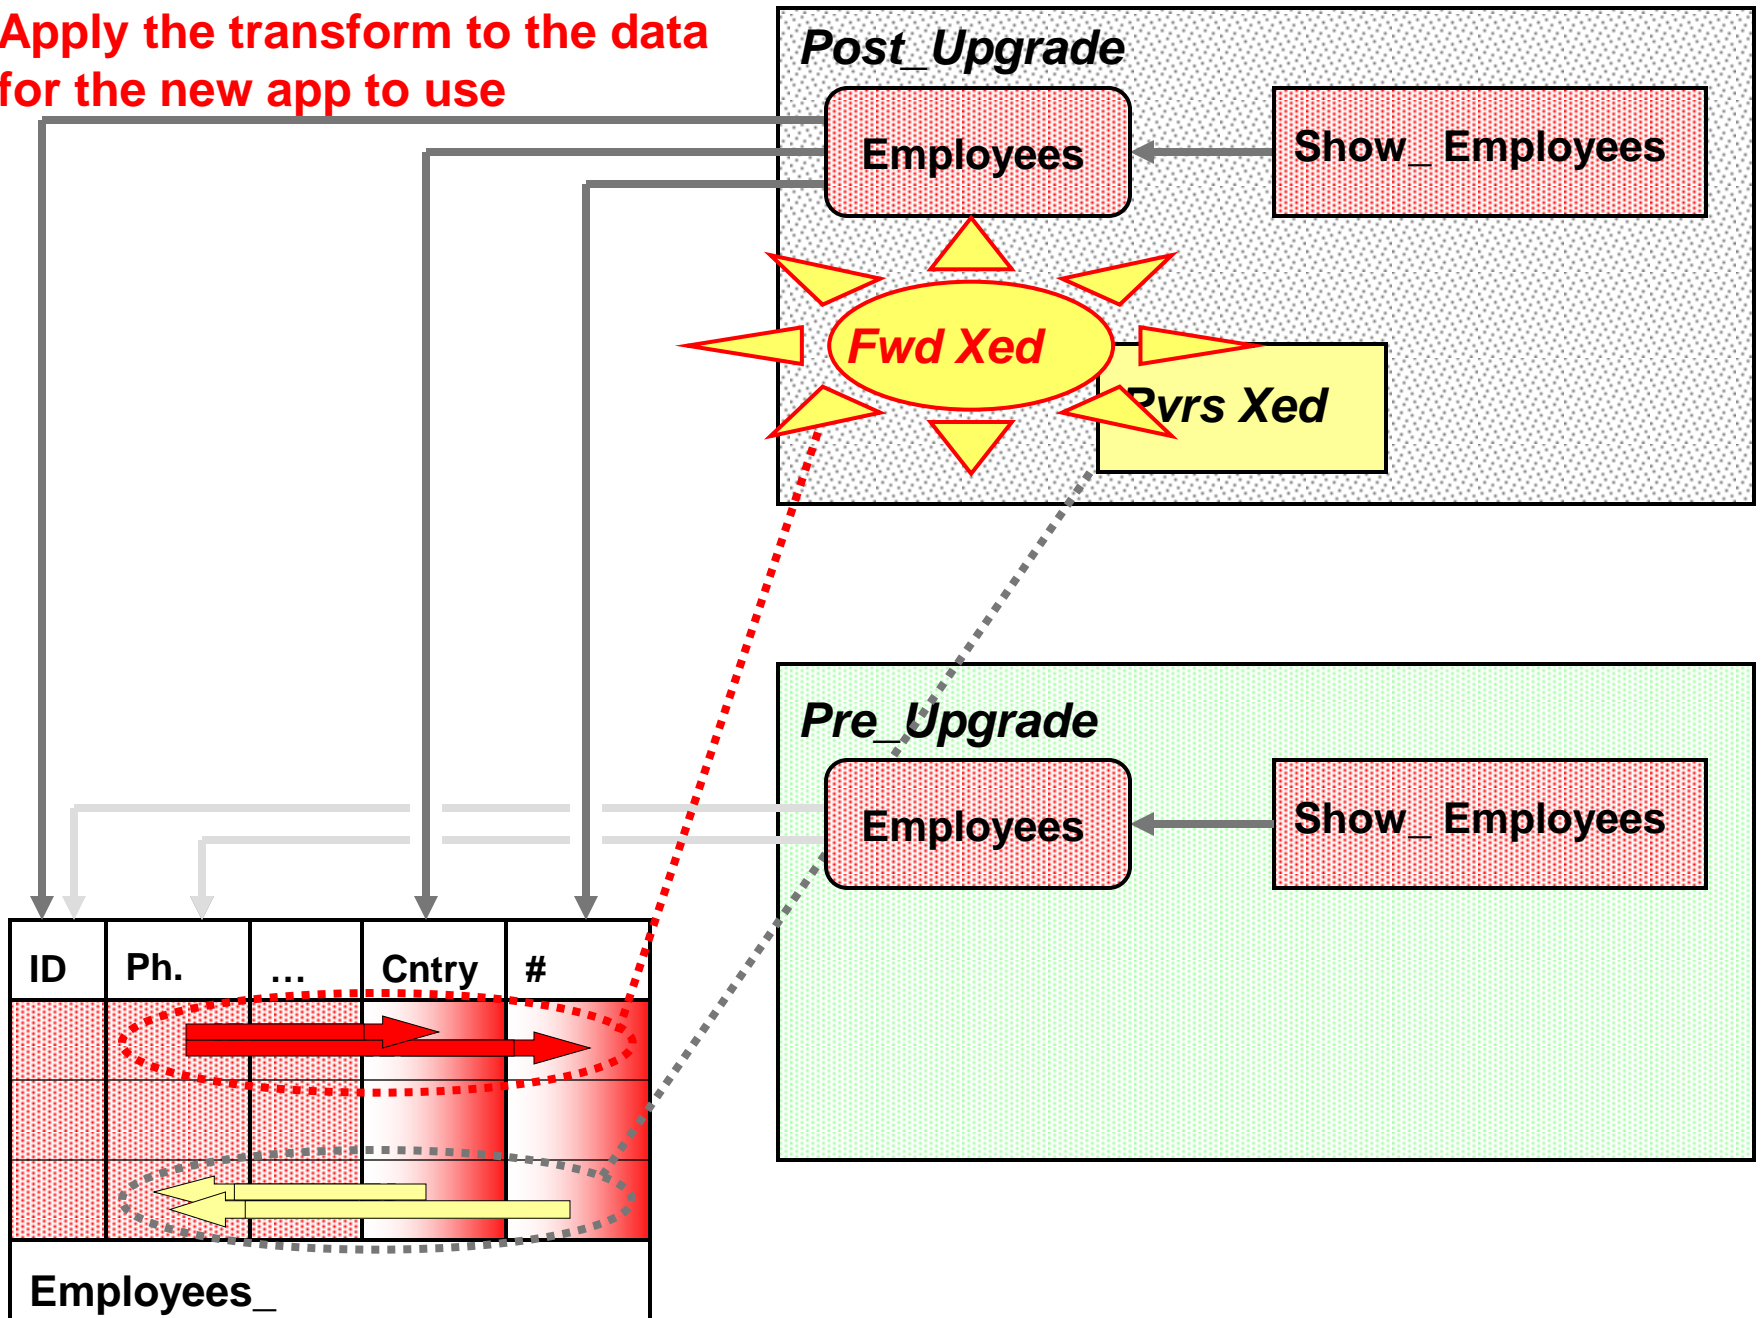
Create the forward crossedition trigger.



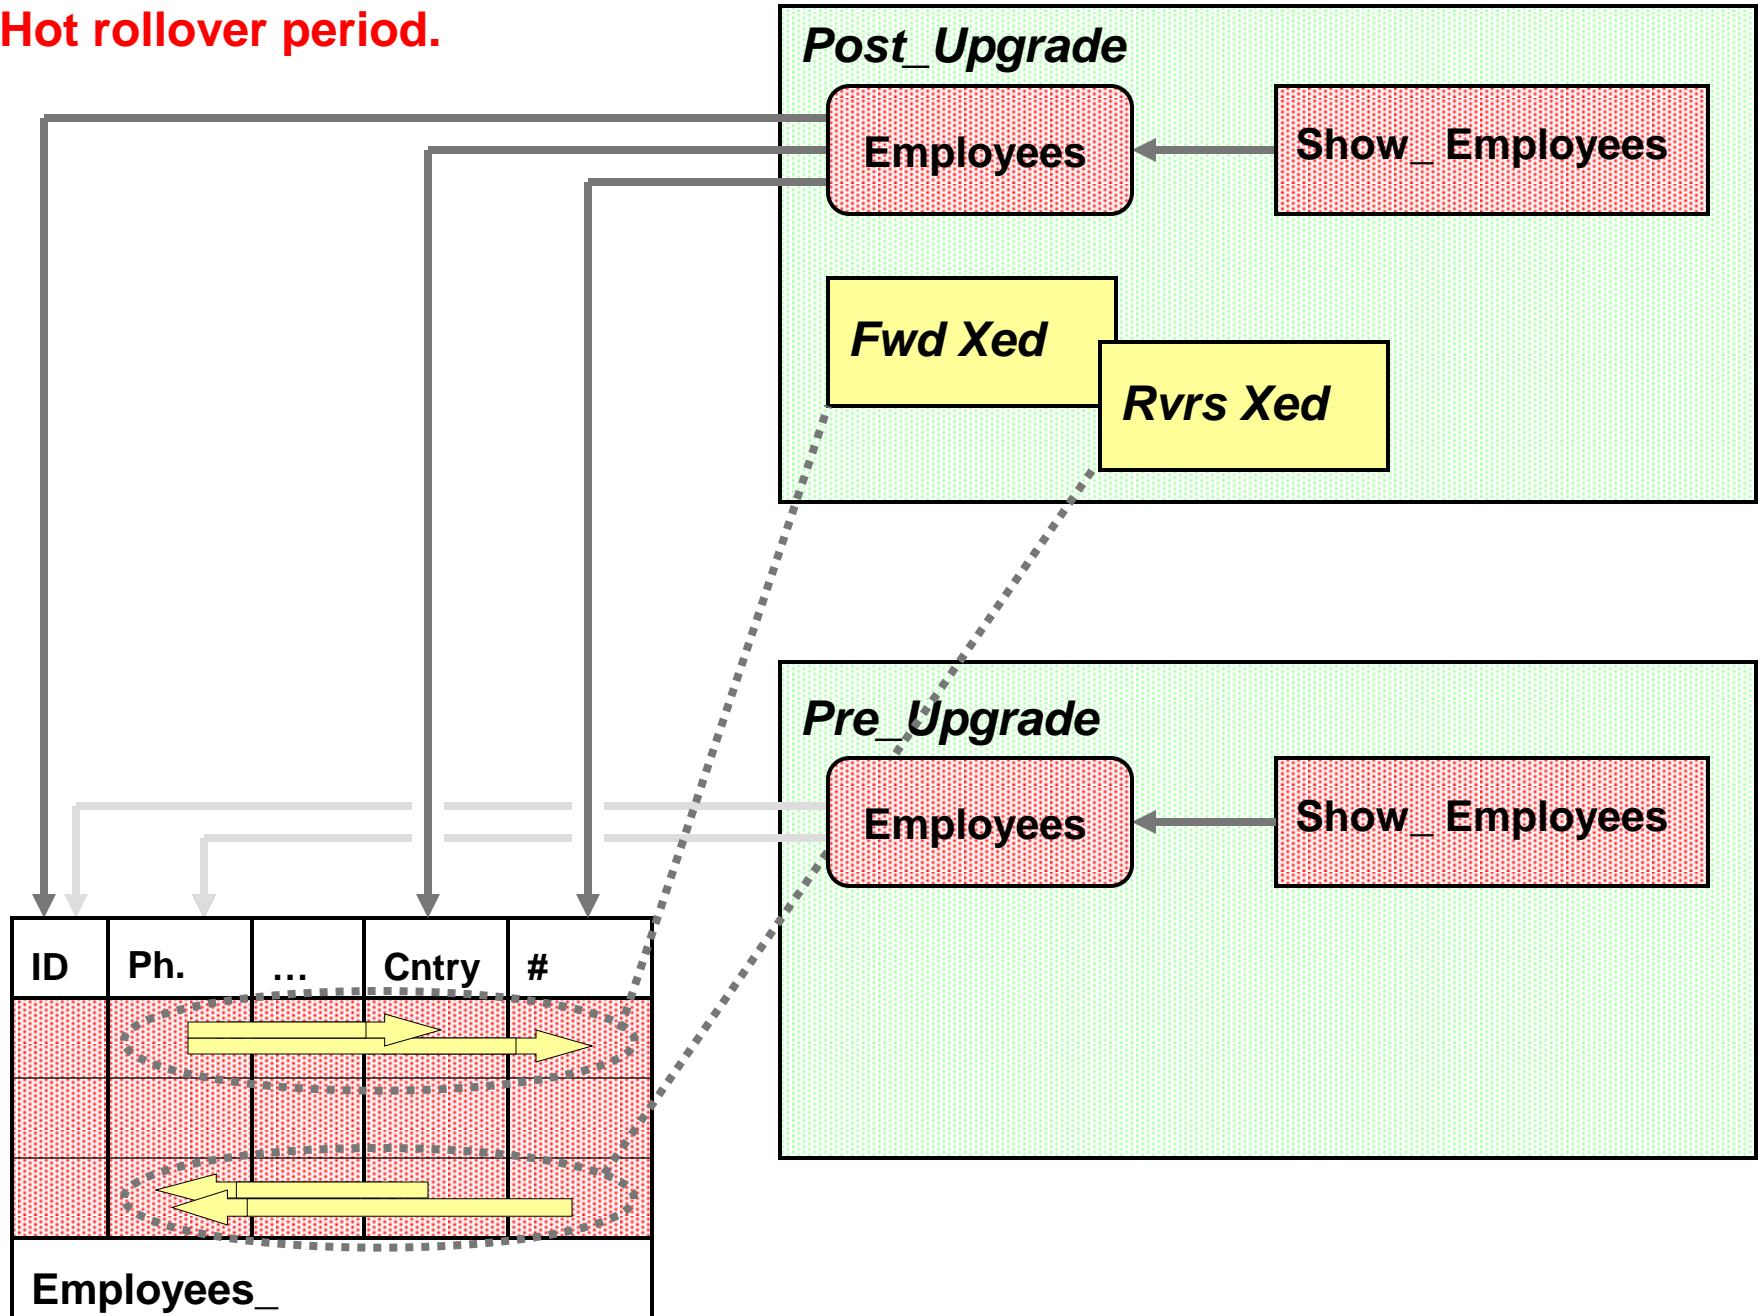
Create the reverse crossedition trigger.



Apply the transform to the data for the new app to use

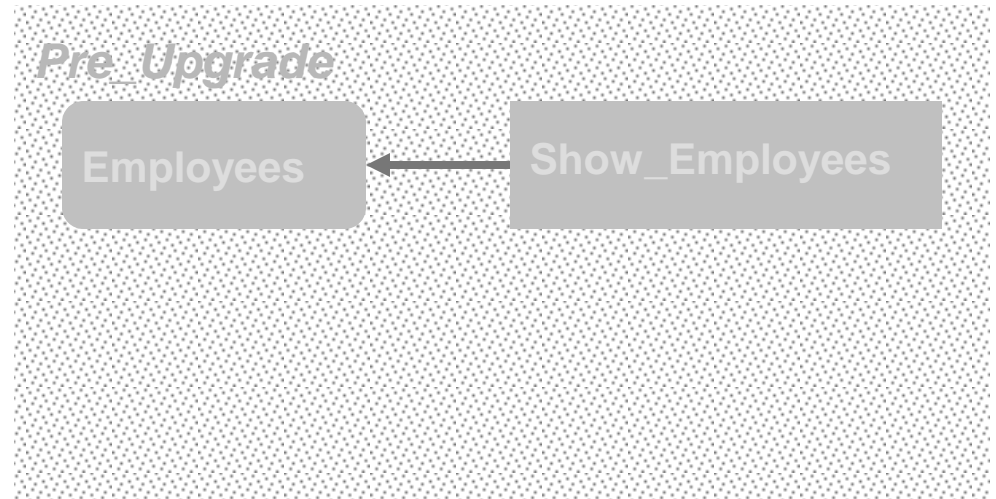
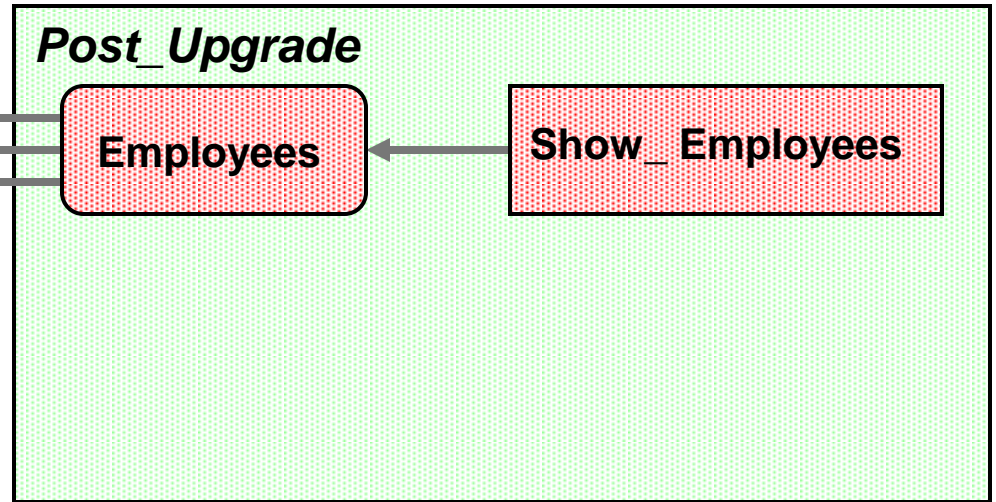


Hot rollover period.



The *Pre_Upgrade* edition is retired.

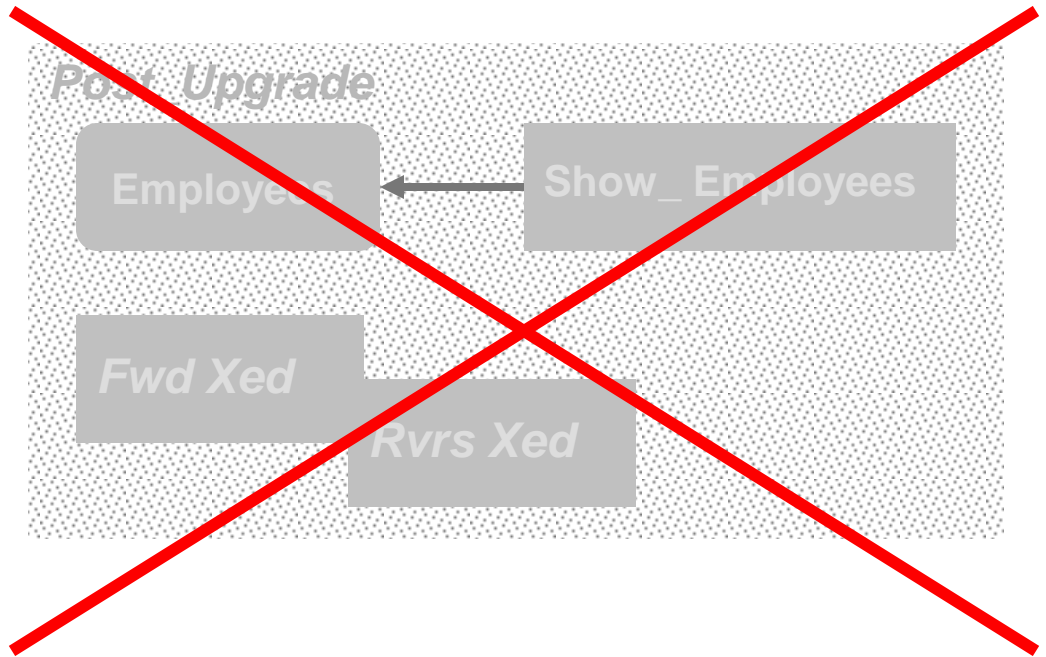
The edition-based redefinition exercise is complete.



ID	Ph.	...	Cntry	#

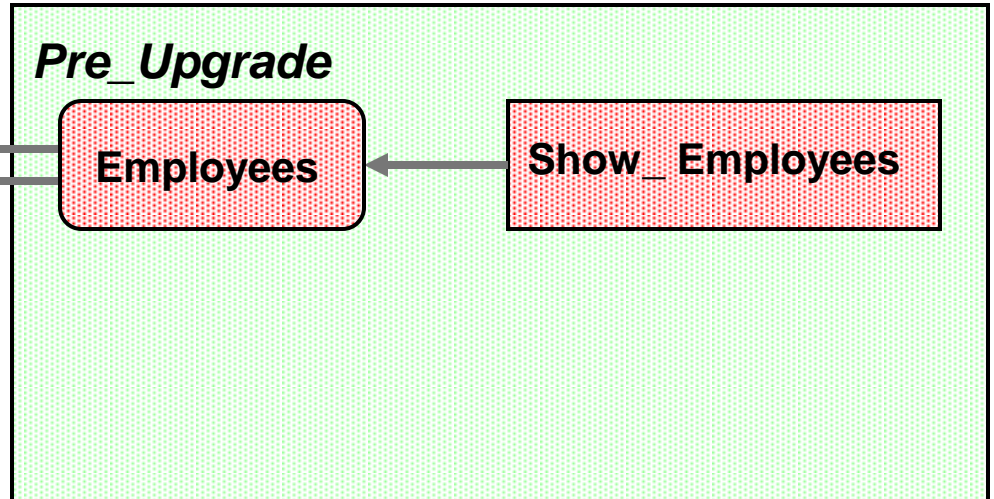
Employees_

The pristine Pre_Upgrade is intact !



ID	Ph.	...	Cntry	#

Employees_



Performance



In-Memory Parallel Execution

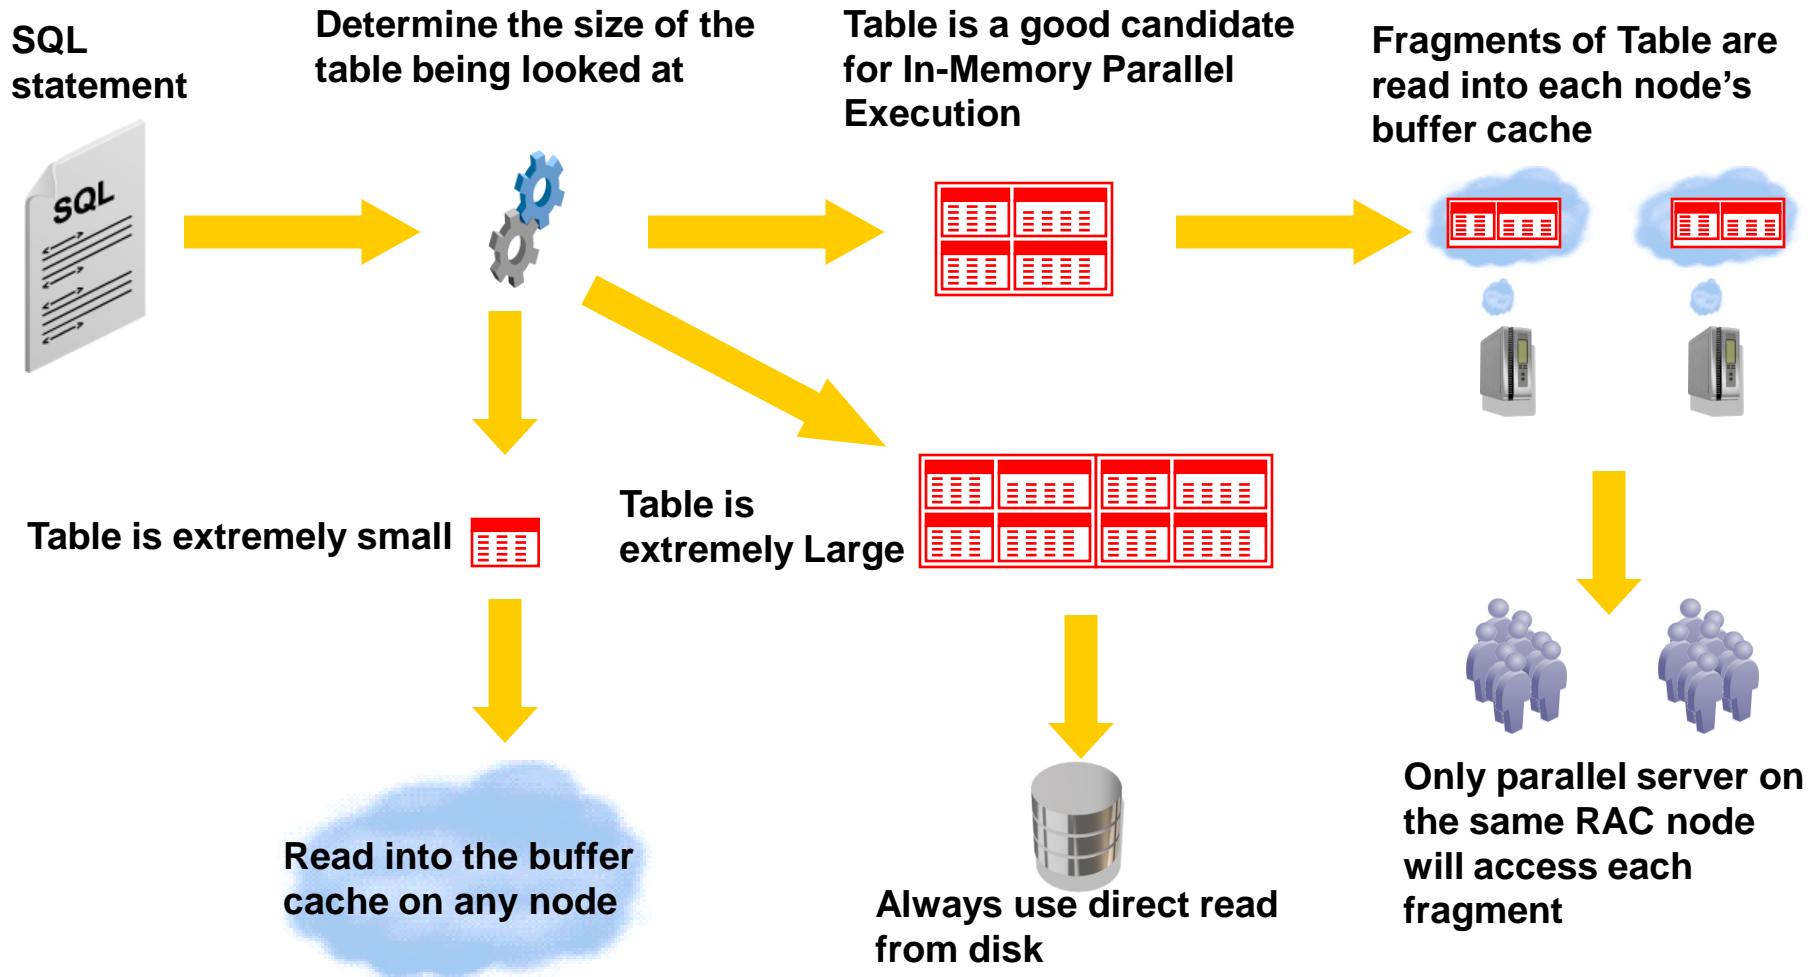
Parallel Execution

- Apply multiple CPU and IO resources to the execution of a single database operation
- Traditionally, Parallel Execution takes advantage of the IO capacity of a system
- Modern grids feature massive amount of low cost memory

Solution

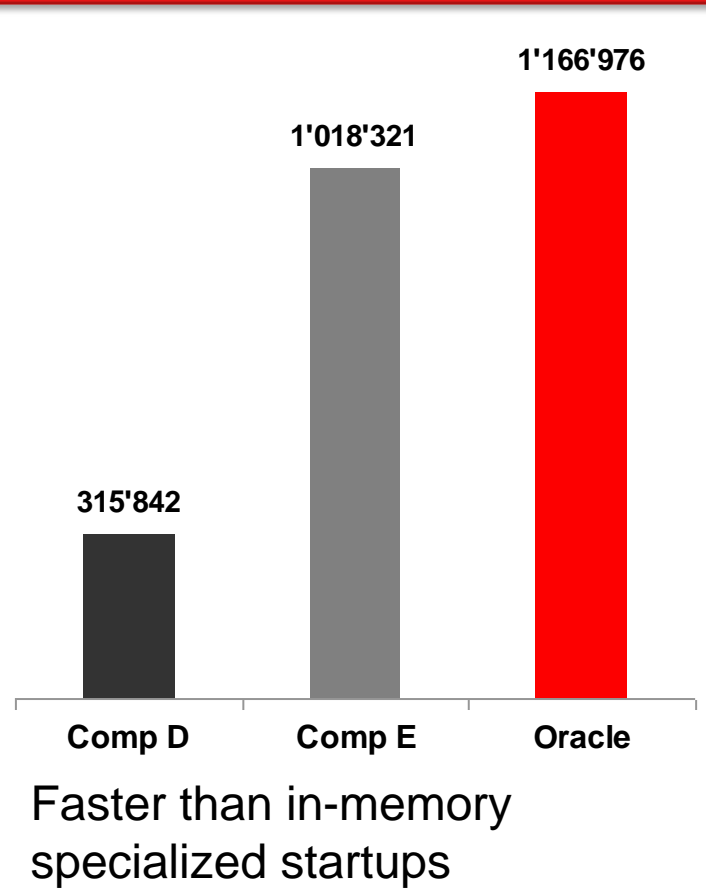
- In-Memory Parallel Execution harness the memory capacity of the entire system
- An algorithm places fragments of a object in memory on different RAC nodes

In-Memory Parallel Execution



In-Memory Parallel Execution

QphH: 1 TB TPC-H



- Parallel query processing on memory cached data
 - Harnesses memory capacity of entire database cluster for queries
 - Foundation for world record 1TB TPC-H

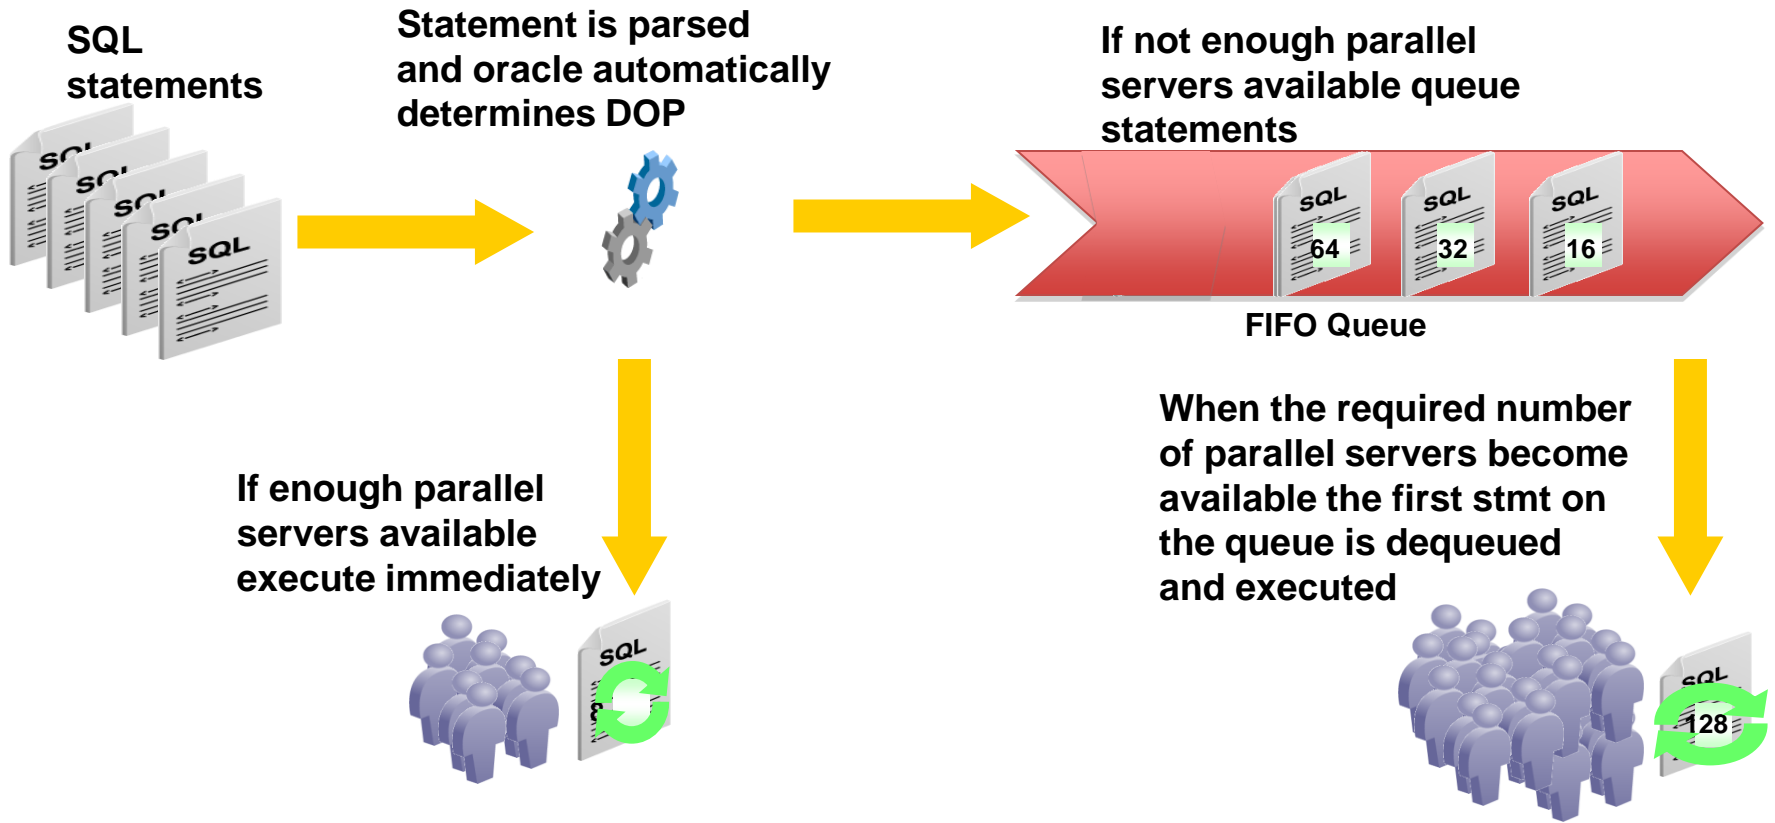
Memory has 100x more bandwidth than Disk

Automatic Degree of Parallelism

Auto DOP

- Optimizer derives the DOP from the statement based on resource requirements for all scans operations
- Applies to all types of statements Query, DML, or DDL
- Explain plan has been enhanced to show DOP selected
- SQL Tune now uses Auto DOP to recommend parallelism

Parallel Statement Queuing



Exadata Database Machine



- **Version 1**

- World's Fastest Machine for Data Warehousing
- Extreme Performance for Sequential I/O
- 10x Faster than other Oracle D/W Systems

- **Version 2**

- 5x Version 1 Data Warehousing Performance
- World's Fastest Machine for OLTP
- Extreme Performance for Random I/O
- Exciting new Exadata Software Capabilities

Sun Oracle Database Machine

Hardware Improvements

- Same architecture as previous Database Machine
 - Same number and type of Servers, CPUs, Disks



Latest Technologies

Faster

80% Faster CPUs

Xeon 5500 Nehalem

100% Faster Networking

40 Gb InfiniBand

50% Faster Disk Throughput

6 Gb SAS Links

200% Faster Memory

DDR3 DRAM

Better

33% More SAS Disk Capacity

600 GB SAS Disks

100% More SATA Disk Capacity

2 TB SATA Disks

125% More Memory

72 GB per DB Node

100% More Ethernet Connectivity

4 Ethernet links per DB Node

Plus Flash Storage!

New

Semiconductor Cache Hierarchy

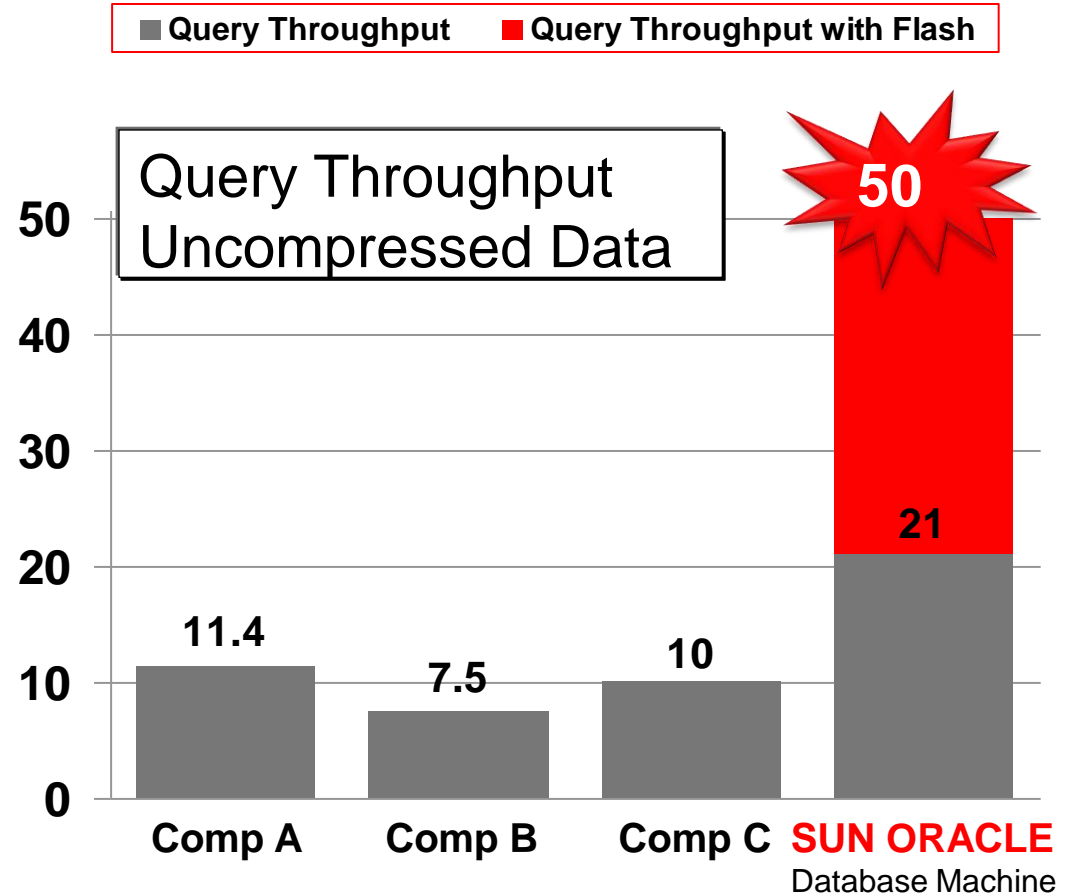
Massive throughput and IOs through innovative Cache Hierarchy



- Database DRAM Cache
 - 400GB raw capacity
 - Up to 4TB compressed user data
 - 100 GB/sec
- Exadata Smart Flash Cache
 - 5TB raw capacity
 - Up to 50TB compressed user data
 - 50 GB/sec raw scan
 - 1 million IO/sec
- Exadata disks
 - 100TB or 336TB raw
 - Up to 500TB compressed user data
 - 21 GB/sec scan
 - 50K IO/sec


Flash

- Flash storage more than doubles scan throughput
 - **50 GB/sec**
- Combined with Columnar compression
 - Up to 50 TB of data fits in flash
 - Queries on compressed data run up to **500 GB/sec**



Exadata Database Processing in Storage

- Exadata storage servers implement data intensive processing in storage
 - Row filtering based on “where” predicate
 - Column filtering
 - Join filtering
 - Incremental backup filtering
 - **Storage Indexing**
 - **Scans on encrypted data**
 - **Data Mining model scoring**
- 10x reduction in data sent to DB servers is common
- No application changes needed
 - Processing is automatic and transparent
 - Even if cell or disk fails during a query



Exadata Storage Index

Transparent I/O Elimination with No Overhead

Table

Index

A	B	C	D
	1		
	3		
	5		
	5		
	8		
	3		

Min B = 1
Max B = 5


Min B = 3
Max B = 8

- Exadata Storage Indexes maintain summary information about table data in memory
 - Store MIN and MAX values of columns
 - Typically one index entry for every MB of disk
- Eliminates disk I/Os if MIN and MAX can never match “where” clause of a query
- Completely automatic and transparent

Select * from Table where B < 2 - Only first set of rows can match


Exadata Hybrid Columnar Compression

- Data is stored by column and then compressed
- **Query Mode** for data warehousing
 - Optimized for speed
 - 10X compression ratio is typical
 - Scans improve proportionally
- **Archival Mode** for infrequently accessed data
 - Optimized to reduce space
 - 15X compression is typical
 - Up to 50X for some data




Up To
50X


Benefits Multiply




10 TB of user data
Requires 10 TB of IO



1 TB
with compression



100 GB
with partition pruning



20 GB
with Storage Indexes



5 GB Smart Scan on
Memory or Flash



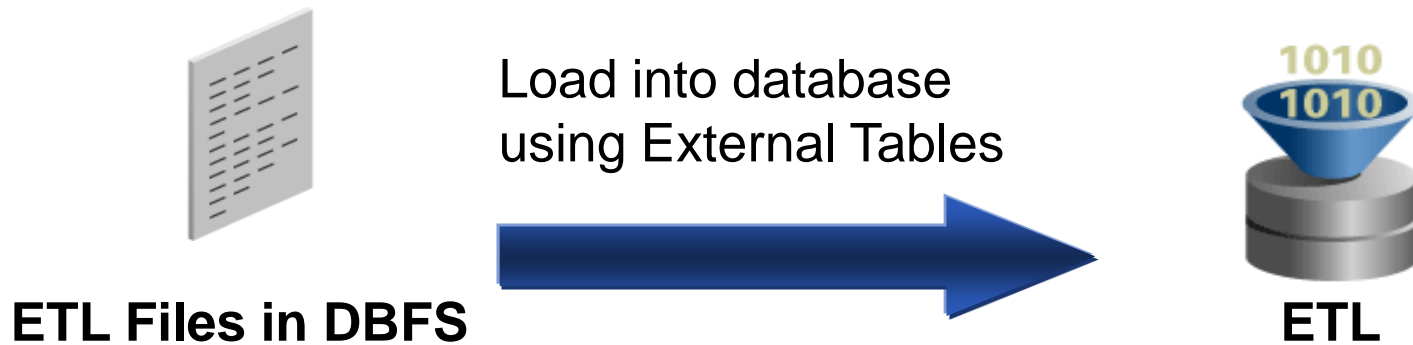
Subsecond
On Database
Machine



Data is 10x Smaller, Scans are 2000x faster

DBFS - Scalable Shared File System

- Database Machine comes with DBFS shared Linux file system
 - Shared storage for ETL staging, scripts, reports and other application files
- Files stored as SecureFile LOBs in database tables stored in Exadata
 - Protected like any DB data – mirroring, DataGuard, Flashback, etc.
- 5 to 7 GB/sec file system I/O throughput



More File Throughput than High-End NAS Filer

11g R2 major new features

Grid

- **Dynamic Cluster Partitioning via Server Pools**
- **Dynamic Cluster Reconfiguration for High Availability**
- **Simplified Provisioning**
- **RAC ONE – A grid of virtualized single instance databases**
- **ASM Supports All File Types**
- **Active Data Guard - Automatic Block Repair**

Manageability

- **Offline Performance Analysis with Active Reports**
- **Maximum Availability Architecture Advisor**
- **Compression Advisor**

Data Growth

- **Partitioning enhancements**
- **SecureFiles new compression: 'Low'**
- **Compress all data, 2x – 50x**
- **Hybrid Columnar Compression**
- **Automated Degree of Parallelism**
- **In-Memory Parallel Execution**
- **Database File System (DBFS)**
- **Streams Enhancements – Performance**
- **New RMAN backup compression capabilities**

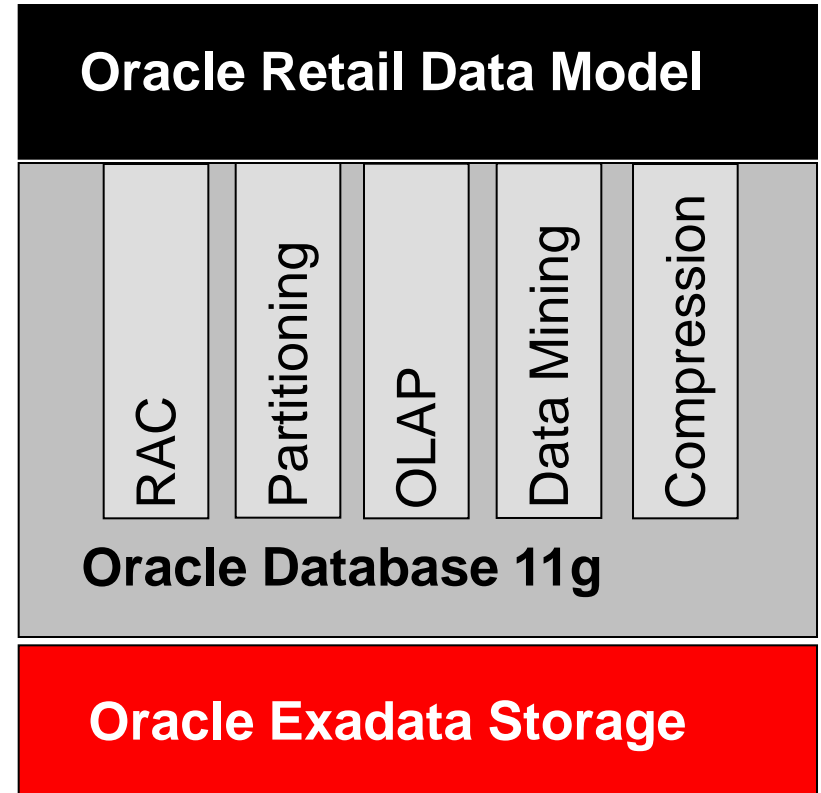
Misc.

- **Total Recall - Schema Evolution Support**
- **Segment Creation on Demand**
- **Resource Manager Instance Caging**
- **In-Memory Parallel Query Server Control**
- **Improved XML DB Performance and Scalability**
- **Database Replay Improvements**
- **SQL Performance Analyzer Improvements**




Oracle Retail Data Model

- Built using industry standards
 - 3NF logical data model
- Complete with pre-built
 - OLAP models
 - Data Mining models
 - Sample reports and dashboards
 - Retail-specific measures & KPIs
- Quick and easy to deploy
 - Speed time-to-value
- Optimized for Oracle
 - Best-in-class data warehousing technology
 - Including Exadata Storage



For More Information

<http://search.oracle.com>

or

www.oracle.com/database

**Oracle® Database New Features Guide
11g Release 2 (11.2) [Link](#)**