Private Cloud Database Consolidation

Randy Hietter, Product Management - Oracle

Gartner on Database Consolidation

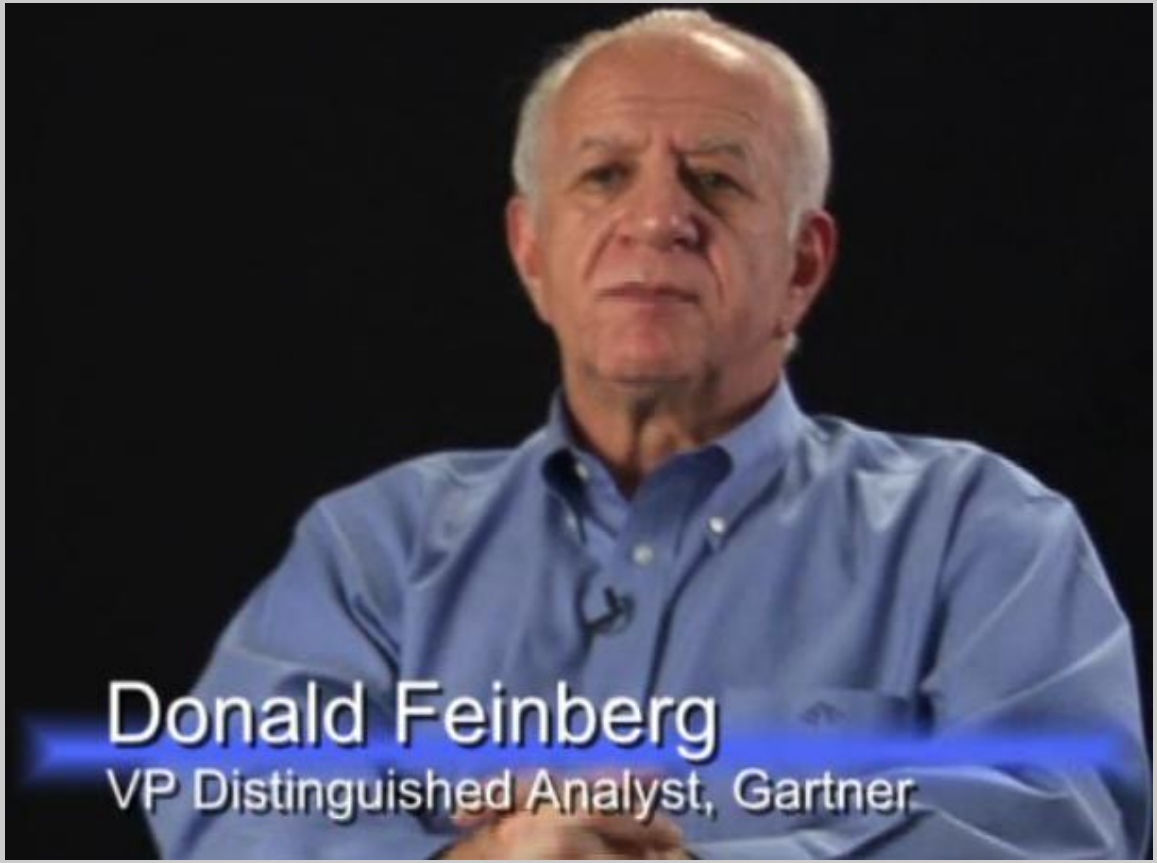

Topics

- What is Cloud computing?
- Business drivers for cloud computing
- Database deployment models
- Cloud deployment models
- Oracle enabling technologies
- Customer Examples
- Considerations for your Project
- Summary

NIST Definition of Cloud Computing

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

This cloud model promotes availability and is composed of:

5 Essential Characteristics
- On-demand self-service
- Resource pooling
- Rapid elasticity
- Measured service
- Broad network access

3 Service Models
- SaaS
- PaaS
- IaaS

4 Deployment Models
- Public Cloud
- Private Cloud
- Community Cloud
- Hybrid Cloud

Source: NIST Definition of Cloud Computing v15

Database Cloud definition varies

- **All customers embracing**
  - **Resource pooling**
  - **Rapid elasticity**
- **Other characteristics adopted selectively**
  - **On-demand self-service**
  - **Measured service**
  - **Broad network access**
- **Standardization key for most customers**
  - **Cloud services**
  - **Cloud building blocks**

Database as a Service

Two alternatives to deploy DBaaS

Software-as-a-Service

ORACLE ON DEMAND

salesforce

Google Apps

Platform-as-a-Service

Database Cloud

Infrastructure-as-a-Service

Database in a Cloud

Database-as-a-Service

amazon web services™

SimpleDB, Oracle DB in OVM

Database as a Service

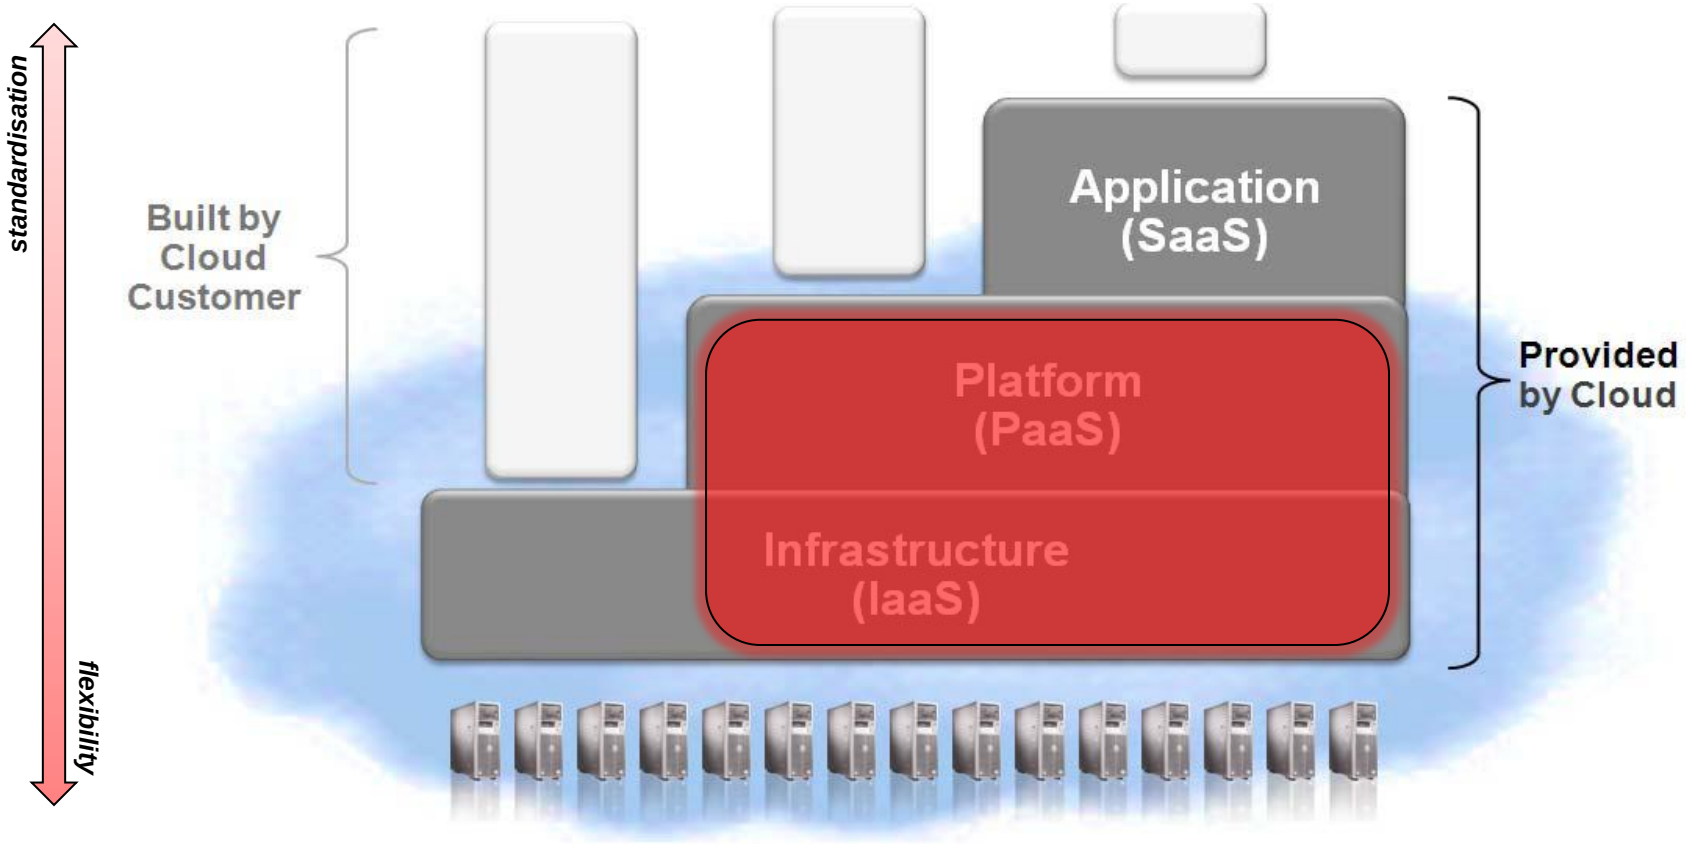

- Infrastructure Platforms are pre-integrated software assemblies
- A set of standard, pre-built containers into which we build and run applications and services, delivered as-a-Service.
- Centralized, Unified management software which provides a single point of control over all our Infrastructure Platforms

Database Cloud Business Drivers

Lower:
- CapEx
  - Servers
  - Storage
  - S/W licenses
- OpEx
  - Maintenance
  - Management

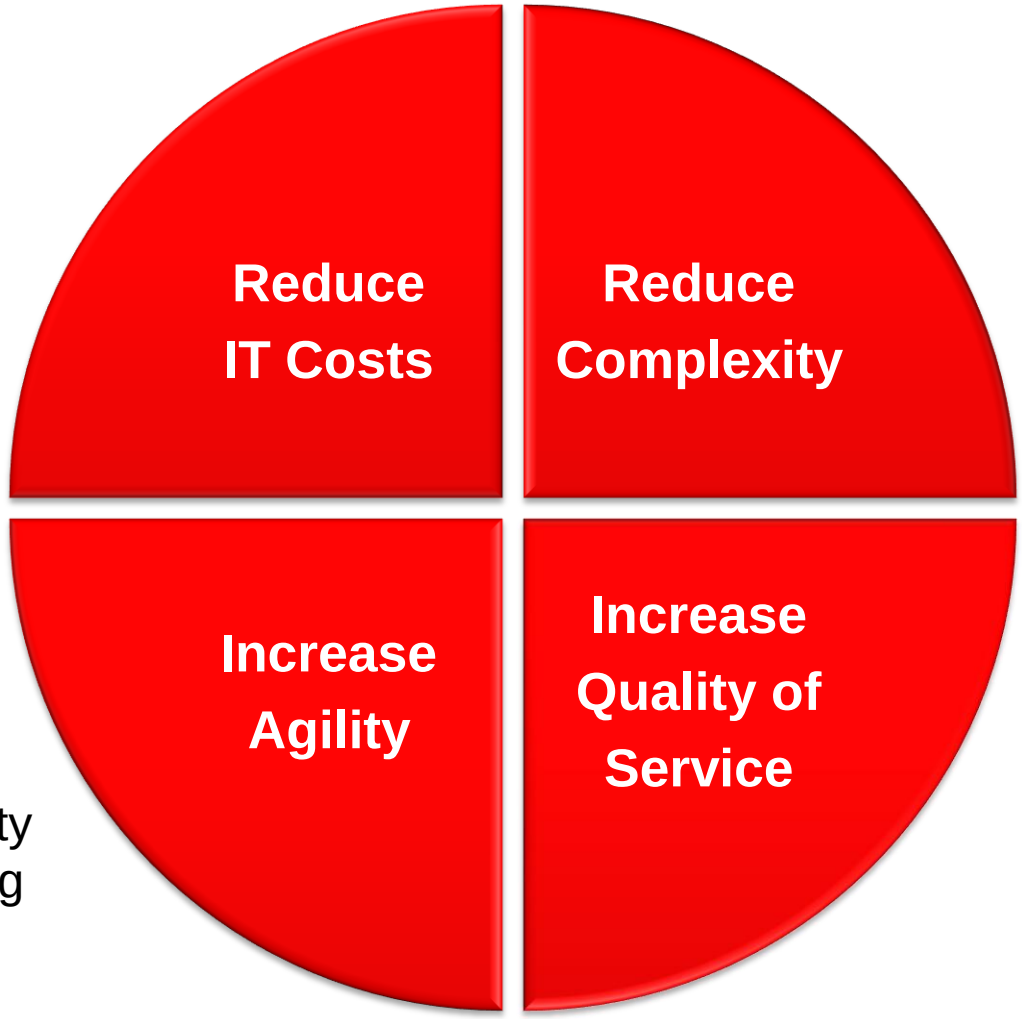

Reduce:
- Configurations
- Services

Standardize:
- OS
- DB Versions

Enable:
- Resource Elasticity
- Rapid Provisioning
- Fast Deployment

Enhance:
- IT service time
- Availability
- Security

Oracle Customer Survey

- Survey of 400+ customers
- About 50% response
- Reducing costs top busi driver
- Most customers interested in hybrid deployment model

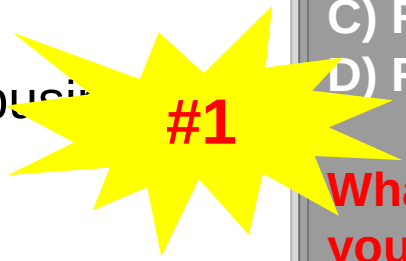

What's driving your interest in Private Clouds?

A) Higher quality of service 20.2% (44)
B) Improve agility 22.9% (50)
C) Reduce complexity 16.5% (36)
D) Reduce IT Costs 40.4% (88)

What type of Consolidation are you interested in?

A) Server 10.1% (22)
B) Operating System 11.5% (25)
C) Database 26.3% (57)
D) Mixture of Consolidation types 52.1% (113)

Private Clouds Are Gaining Momentum

2010 Survey* of Independent Oracle User Group Members

Private vs Public Cloud Adoption

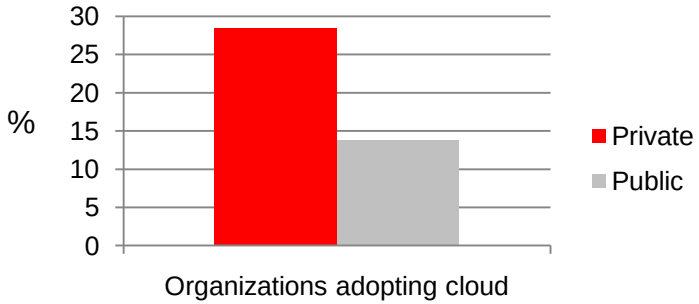

Benefits of Private Cloud Computing

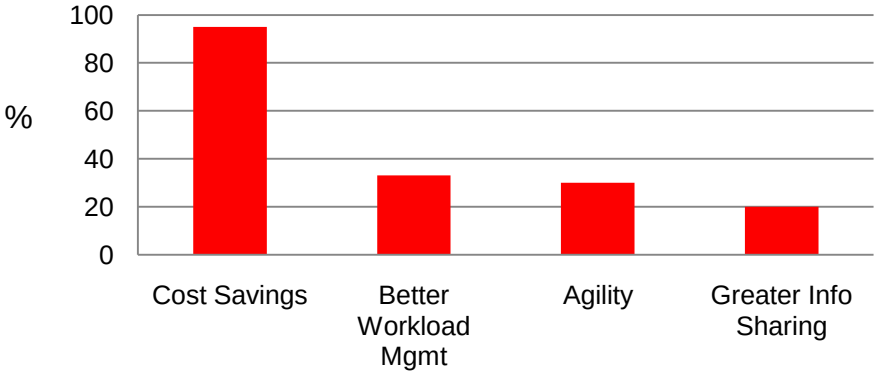

Concerns Over Public Cloud

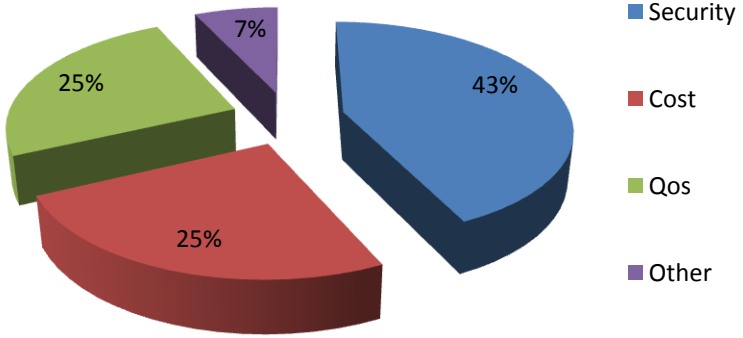

* Findings from IOUG ResearchWire member study on Cloud Computing, conducted in August-September 2010.

ORACLE

Private Database Cloud

Greatest consolidation, maximum ROI

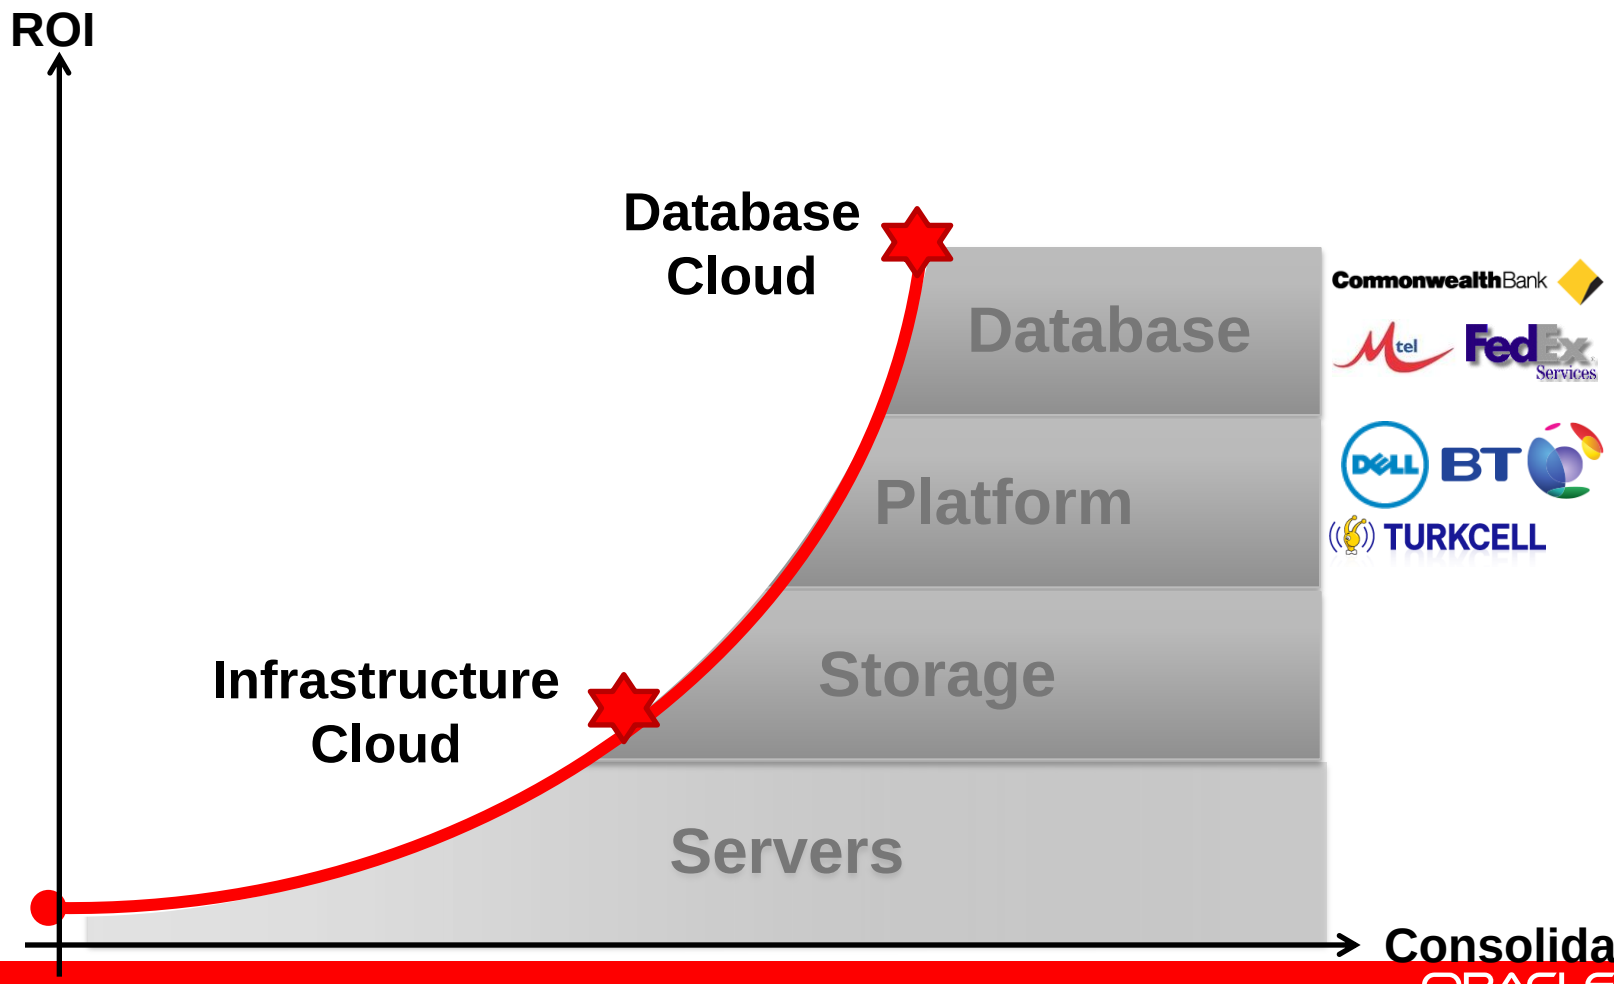

Consolidation Plan

Hypothetical Example of a Large Financial Institution

Consolidation → Rationalization

Environment	Start	Server	Storage	OS	DB
# of Databases	350	350	350	350	3
Homes	350	350	350	350	6
Versions	15	15	15	15	3
Configurations	20	20	20	20	3
# of OSes	2	2	2	1	1
Versions	10	10	10	1	1
Configurations	3	3	3	1	1
# of Servers	350	15	15	15	15
# of Storage Pools	350	350	15	15	15

Database Cloud Architectures

Common building blocks are shared server and storage pools

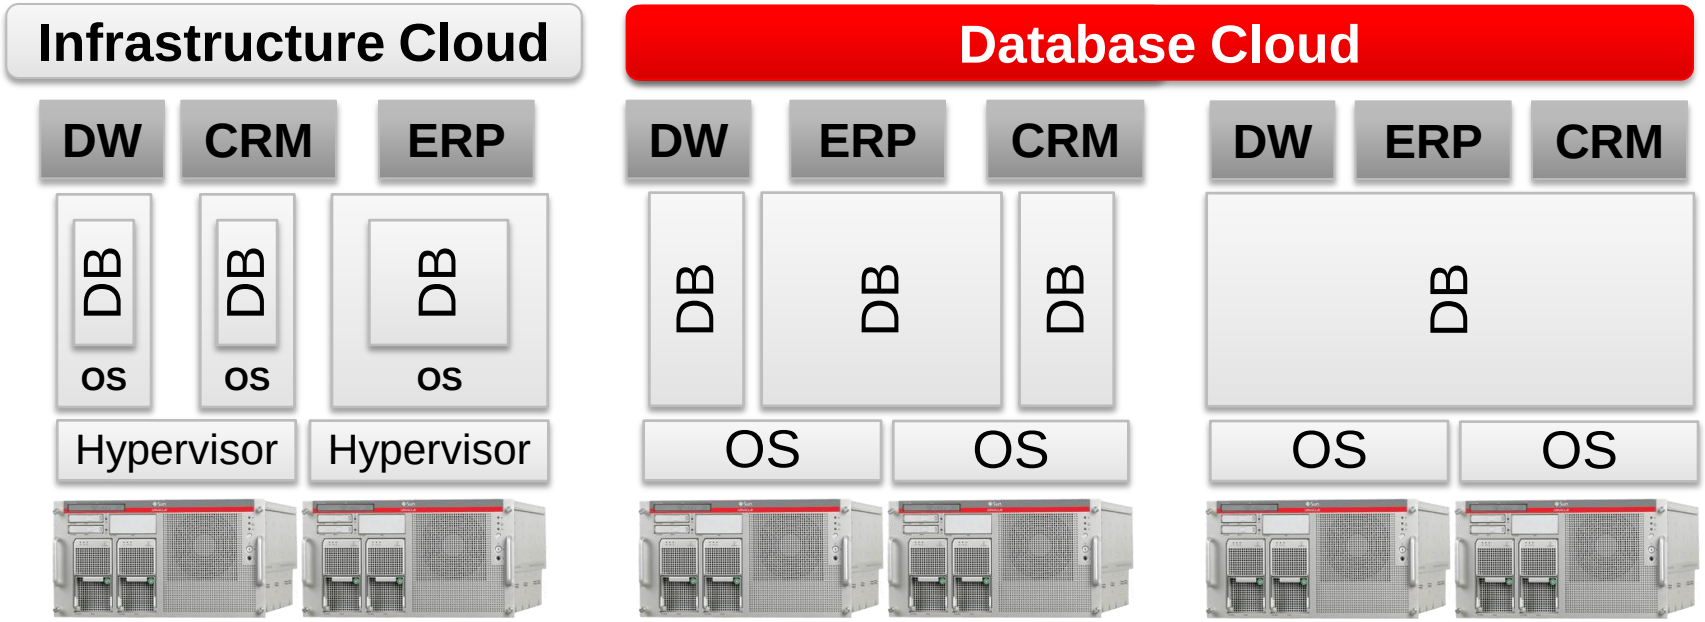

Server

Deploy in dedicated VMs
Server virtualization

Platform

Share server pool
Real Application Clusters

Database

Share database instances
Real Application Clusters

Infrastructure Cloud

Server - Provision a Database in a VM

- **Reasons for adoption**
  - Simple to implement
  - Excellent isolation
  - Mixed workloads
  - As-is consolidation
  - Legacy support
- **Customer concerns**
  - Lower consolidation density
  - Lower ROI
  - Performance (latency)
  - Managing sprawl
  - Not suitable for all deployments

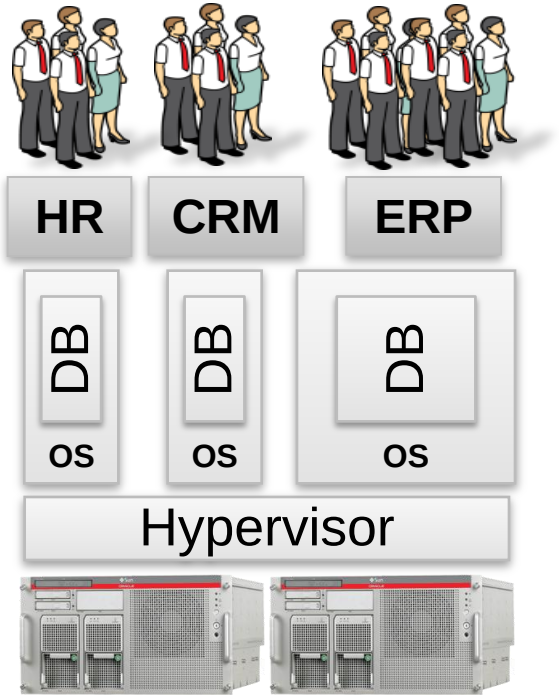

Private Database Cloud

Platform – Provision Database

- Reasons for adoption
  - Consolidation density
  - Good ROI
  - Performance
  - Supports any app
- Customer concerns
  - Requires OS standardization
  - Database only

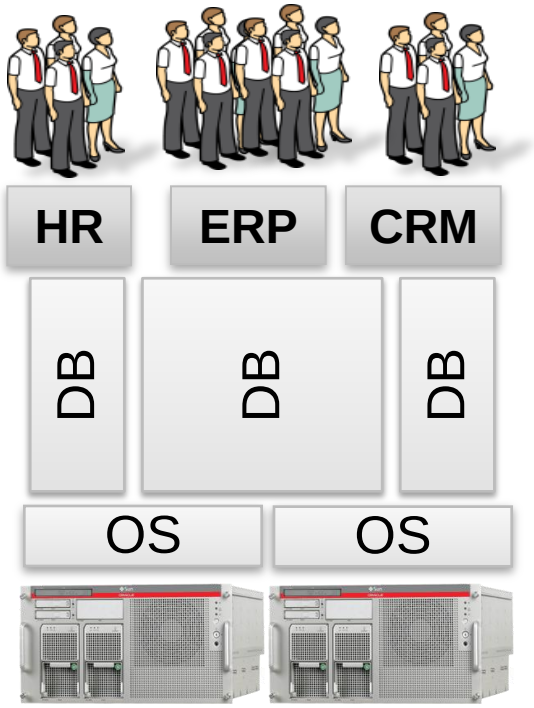

Private Database Cloud

Database – Provision Schema

- **Reasons for adoption**
  - Most efficient
  - Extremely fast provisioning
  - Best ROI
  - Performance
  - Efficient memory use
- **Customer concerns**
  - App qualification required
  - Requires OS and DB standardization
  - Isolation

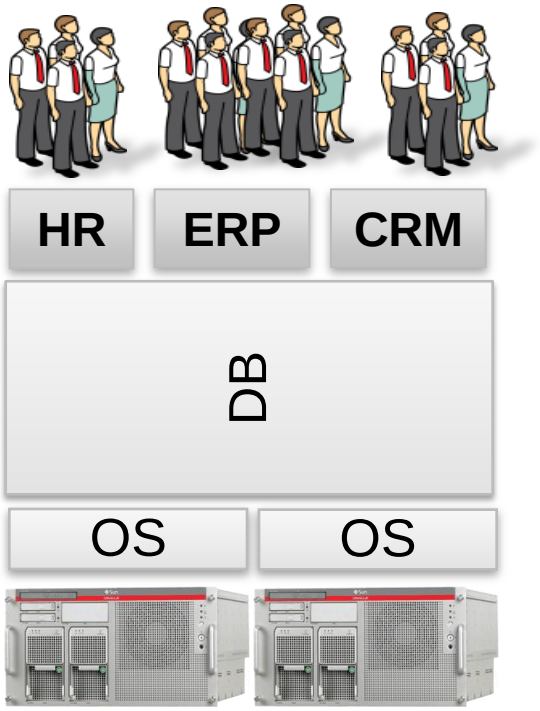

A few new terms

- Database Cloud: The preferred way to implement database consolidation and deploy DBaaS. Leverages database capabilities and does not require an infrastructure cloud.

- Dynamic Database Services: A Database Service which supports online dynamic changes (grow, shrink, move) and transparent failover.

Terminology

- **Server** – An 11gR2 CRS modeling construct that represents a cluster member host.
- **Server Pool** – An 11gR2 CRS modeling construct that represents a set of servers as a placement entity.
- **Cloud Pool (Flex Pool)** – In Exadata environment, it refers to a pool of servers and storage cells carved out into a separate cluster. Generically, it refers to an Oracle11gR2/+ cluster deployment.

Terminology

- **Policy Managed Configuration** – A RAC deployment model where servers are dynamically assigned to the pool(s) hosting the RAC database/services.
- **Administrator Managed Configuration** – A pre-11gR2-like RAC deployment model where servers are manually assigned (aka preferred/available) by the administrator to RAC database/services.
- **Dynamic Service** – A term that refers to RAC service related technologies, such as FaN/FCF/UCP, etc.

Cloud Deployment Models

Typical Infrastructure Cloud

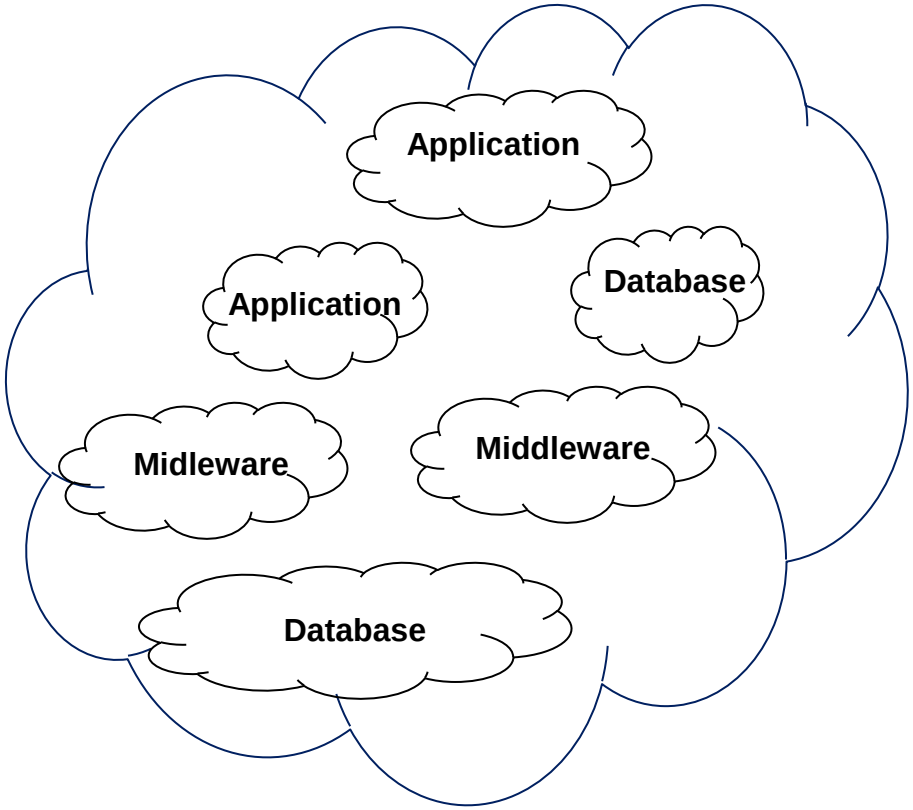

Oracle Cloud Model

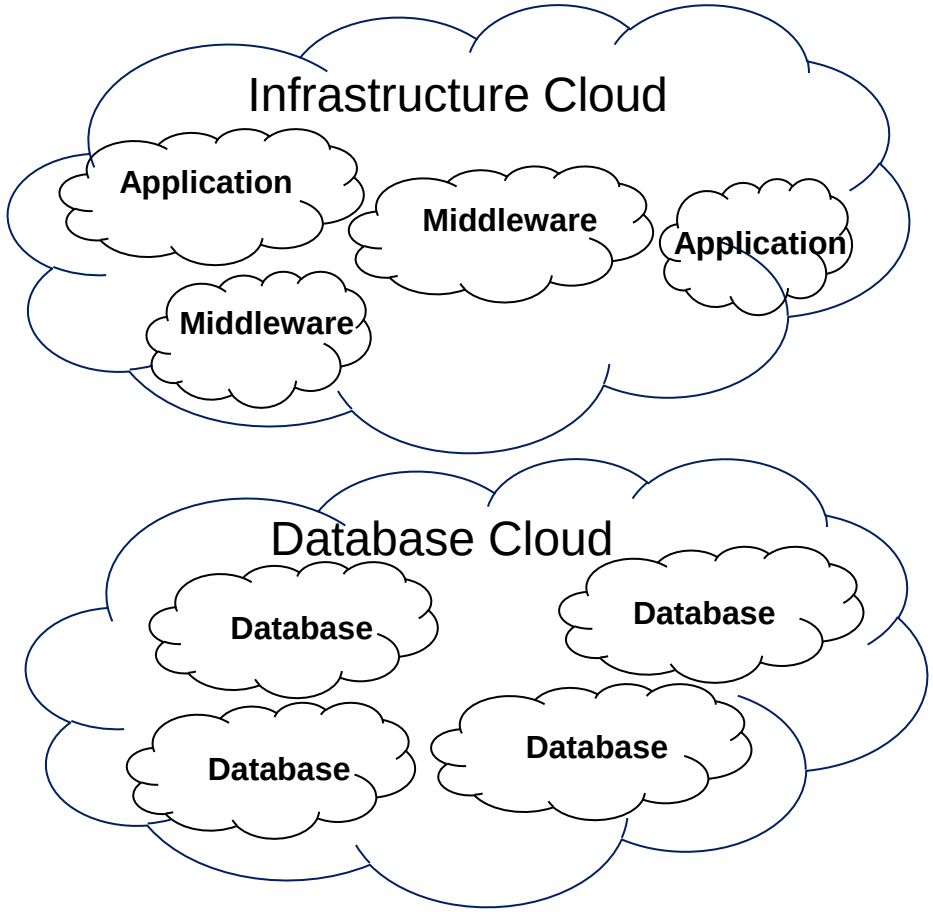

Management Scope

Cloud Zones

- Defines an administrative domain
- Can span multiple cloud pools
- Likely aligned with business unit

Cloud Farm

- Not an EM concept
- Physically connected resources
- Can be configured into pools
- Resources can be easily migrated between (Flex) pools
- Exadata provides this

Cloud Pools

- Equivalent to a cluster
- Access to dedicated or shared storage

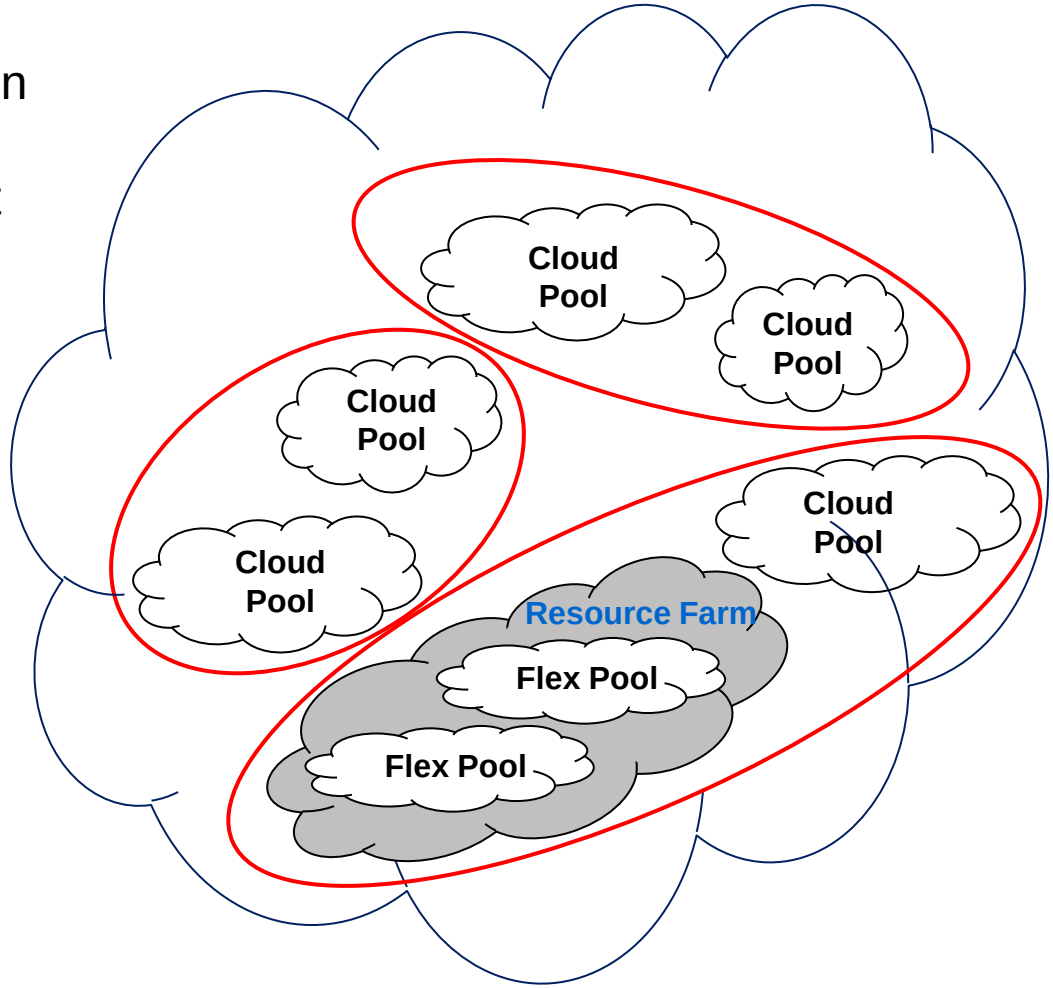

Oracle Enabling Technologies

Years of continuous Oracle innovation

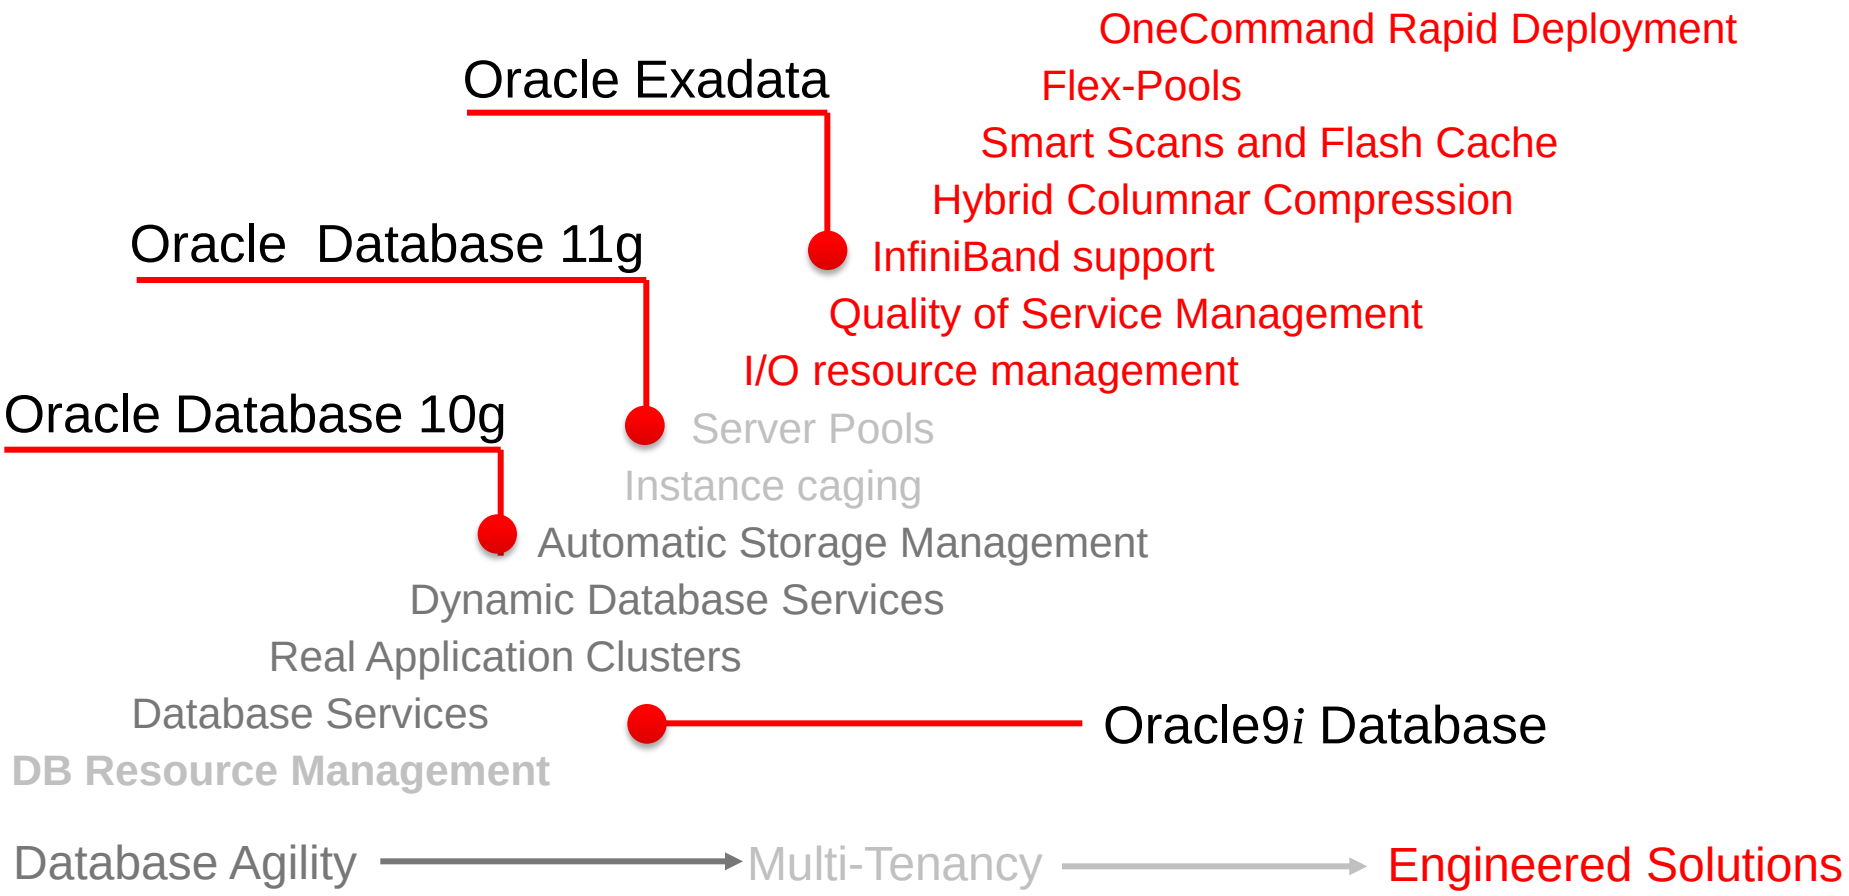

Dell IT

Consolidated 450 Databases onto 3 Private Clouds

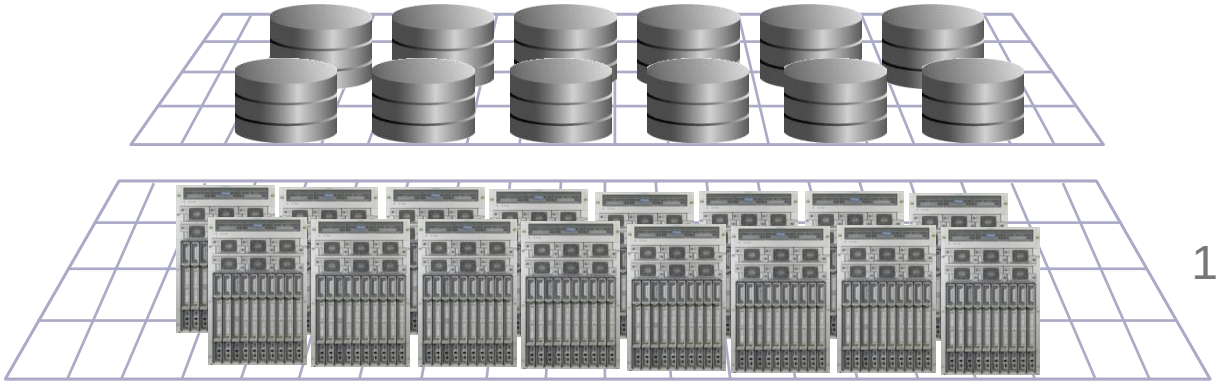

Up to 150 databases per server pool

- P & L break even in 19 months
- Saved $3.3M from labor productivity and cost avoidance
- 50% faster delivery of BI to decision makers
- Time-to-wire reduced from 8 weeks to 2 days

FedEx Services

Re-hosted 400 Databases onto Private Cloud

4-node Cloud pool for smaller databases

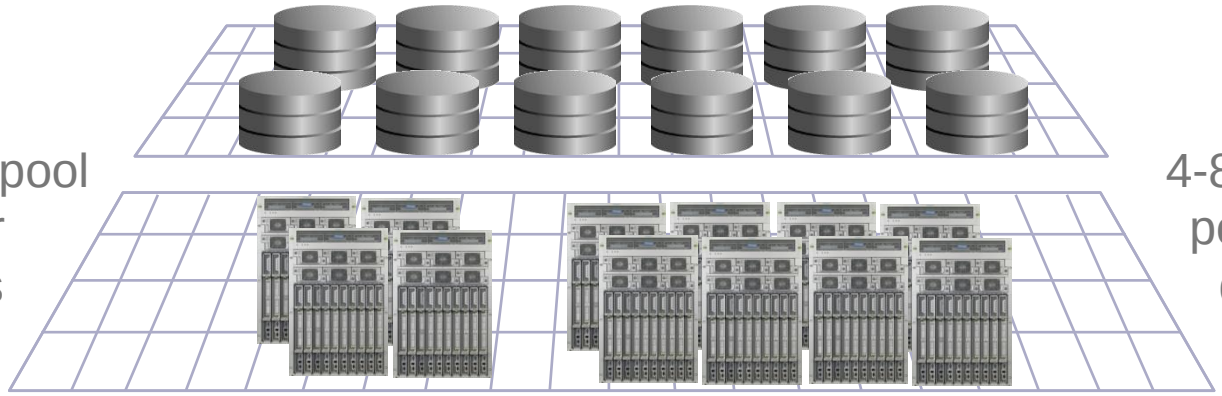

4-8 node Cloud pool for large databases

- Economies of scale consolidating small database schemas (10-12)
- Standardization provides better service at lower cost
- 90% of OLTP workload leverages dynamic infrastructure
- Apps requiring 50% of resources get dedicated pool sized to need

Considerations for Your Project

Planning

- It is critical to carefully plan for the deployment.
- Many of the Cloud benefits are delivered from rationalizing the existing environment and standardizing the deployment model.

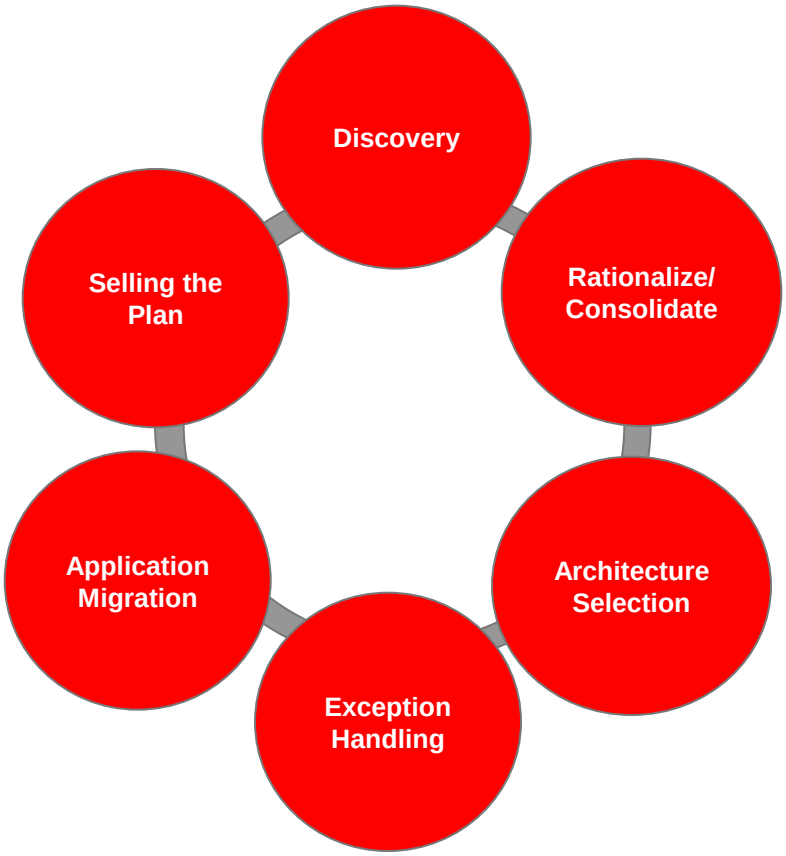

Cloud Planning Process

Discovery

IT Environment

- What is the hardware platform(s)?
- What software stack is being used?
  - What software stack could potentially be used?
- What are the existing databases?
  - DB Versions, OS, OS Versions, Applications?
  - What is the resource consumption of these databases?
- What are the applications' characteristics?
  - Do they support multi-node DB deployment, or singleton only?
  - Internal or Third party? Can the applications be modified?
  - Do they support dynamic DB services? QoS?

Discovery (cont'd)

Business Requirements

- What are the service level requirements?
  - Performance and availability
- What level of Isolation is Required?
  - This may vary by business unit
  - **Isolation drives the solution in most customers.**
- Do tenants need privileged DB access?

Corporate Standards

- What are the corporate standards for IT?
- How will these affect database deployments?

Experience

- Prior virtualization, RAC experience?

Rationalization

Why Rationalize?

- IT rationalization determines the best use of IT services to **reduce non-productive redundancy** in enterprise IT solutions

- By standardizing on a set of building blocks, IT departments can easily deploy pre-defined configurations and scale-out using modular components.

- Standardization results in a more homogeneous environment that is easier to manage, lower cost, less complex, and more agile.

Application Migration

Readiness

- Applications may be in various states of **technical readiness** for a cloud environment, based on application characteristics and requirements
  - Name-space conflicts (relevant for schema consolidation)
  - Character set requirements
  - I/O rate requirements

- Applications may have **business constraints** on their ability to be migrated to the cloud
  - Service level agreements
  - Security
  - Compliant restrictions – PCI-DSS, HIPPA, etc. datasets cannot be co-mingled

Application Selection

Commonwealth Bank of Australia - Customer Example

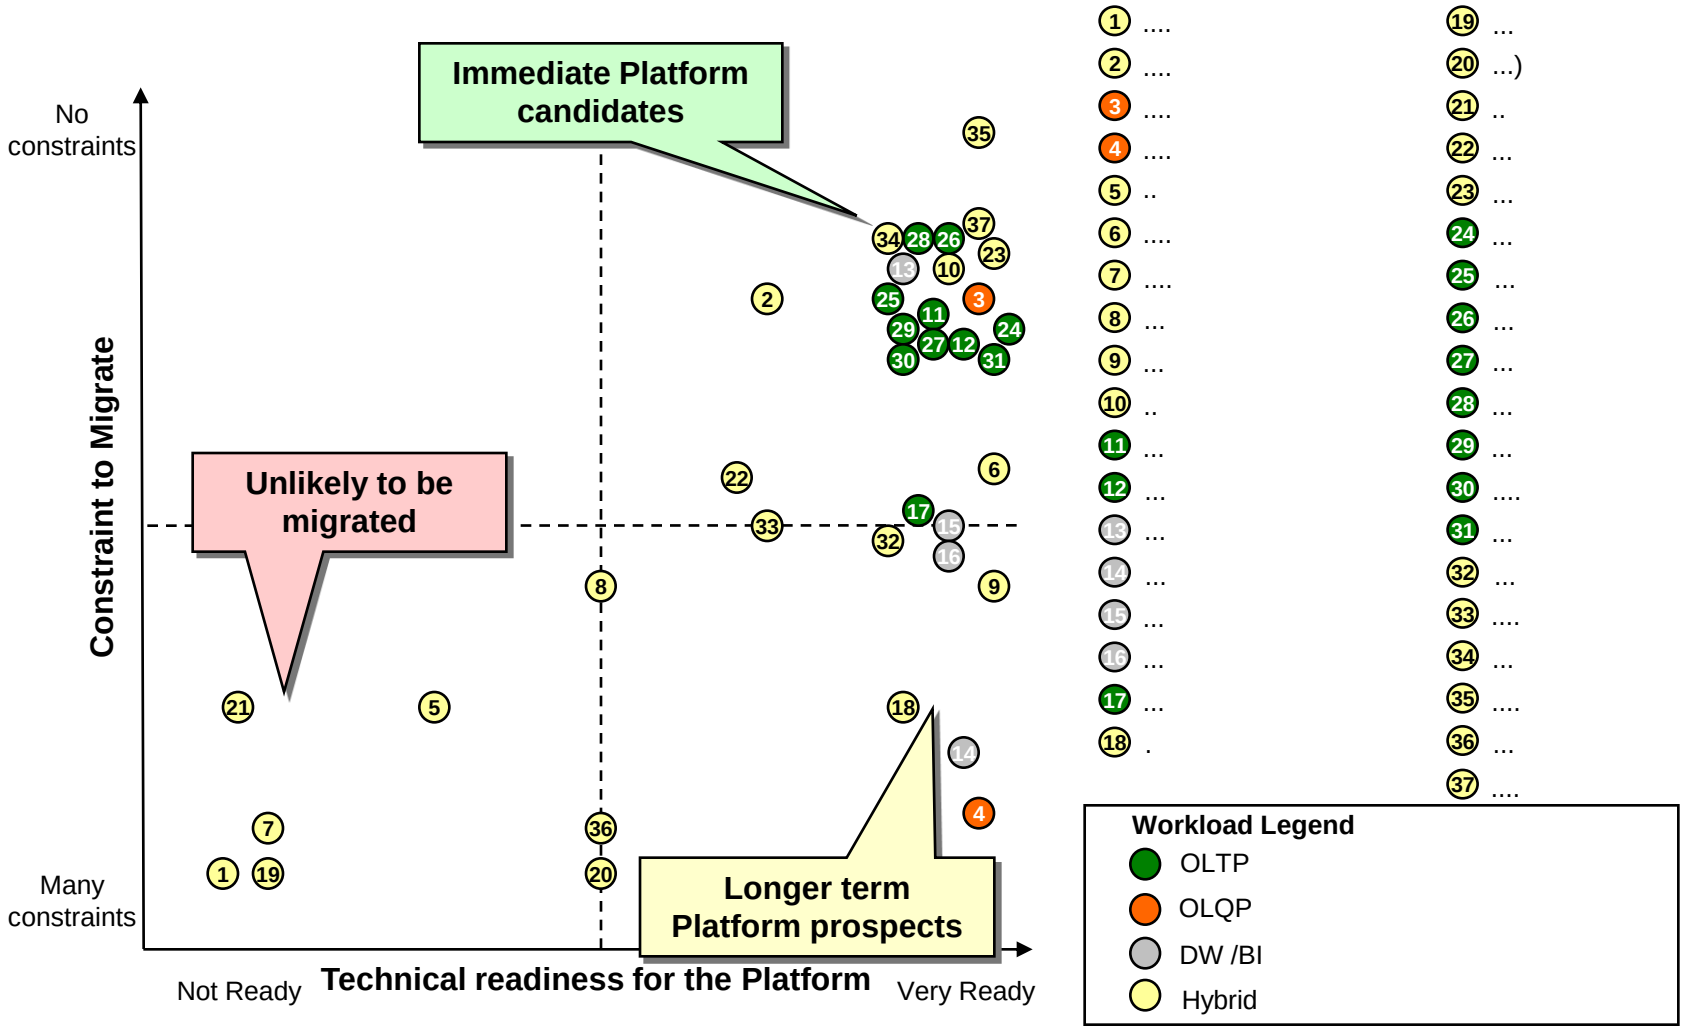

Application Migration Planning

- Special requirements for a given application may make it difficult to migrate
- Decide whether the cloud will be used for new deployments only, or existing deployments as well
- Determine when applications will be migrated
  - During an existing maintenance window?
  - As part of a stand alone event for the application to be migrated?
- Plan a migration process that will cause minimal disruption
- Choose applications that will provide the biggest benefit
- Choose “low hanging fruit” first; Go for quick wins
- Ensure that early migrations are successful!

Exception Handling

Plan for exceptions that don't fit into the cloud

- No matter what architecture or product is chosen to implement the cloud solution, there will be customer workloads and applications that are not well suited for it.
- Develop a plan to handle exceptions rather than to force fitting them into an environment.
- All of the customers that have currently deployed a database cloud solution have a plan in place to handle these exceptions.

Cloud Layout

How will you group applications together?

- **Business**
  - Separate Cloud Pools for different lines of business (LOB) or departments
  - Separate Cloud Pools for different application service levels or governance compliance
- **Functional**
  - A pool for similar functions; e.g., Database pool, Middleware, Applications, Internal/External
- **Technical**
  - Separate pool based on OS type or database version, or isolation requirements
  - Group applications with complementary workloads
  - Pool built around very specific High Availability goals
- **Most customers have chosen both Functional and Technical**
  - Business factors that are considered are Audit/Governance

Consolidating Workloads

New Workload

A

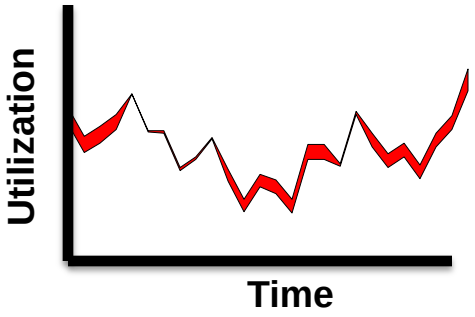

- or -

B

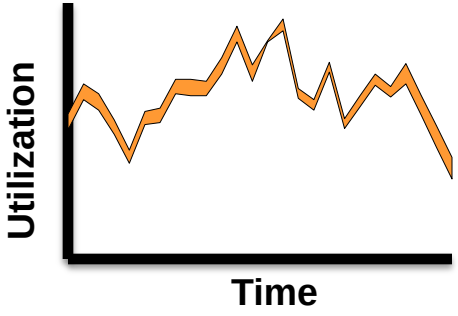

Consolidating Workloads

New Workload

A

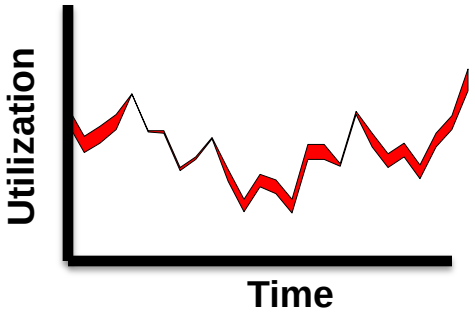

- or -

B

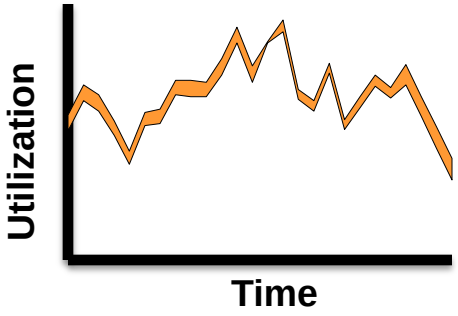

Existing Workload

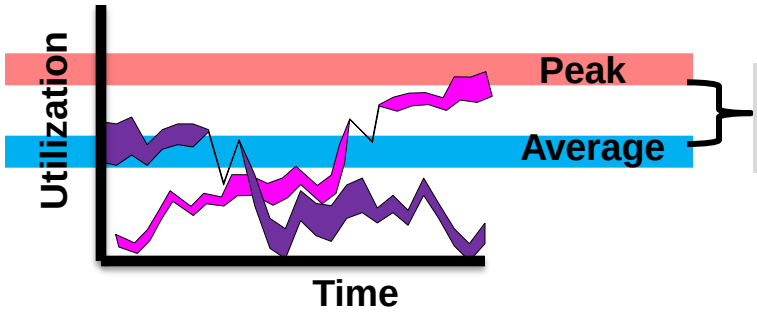

The smaller this gap the better

Consolidating Workloads – Poor match

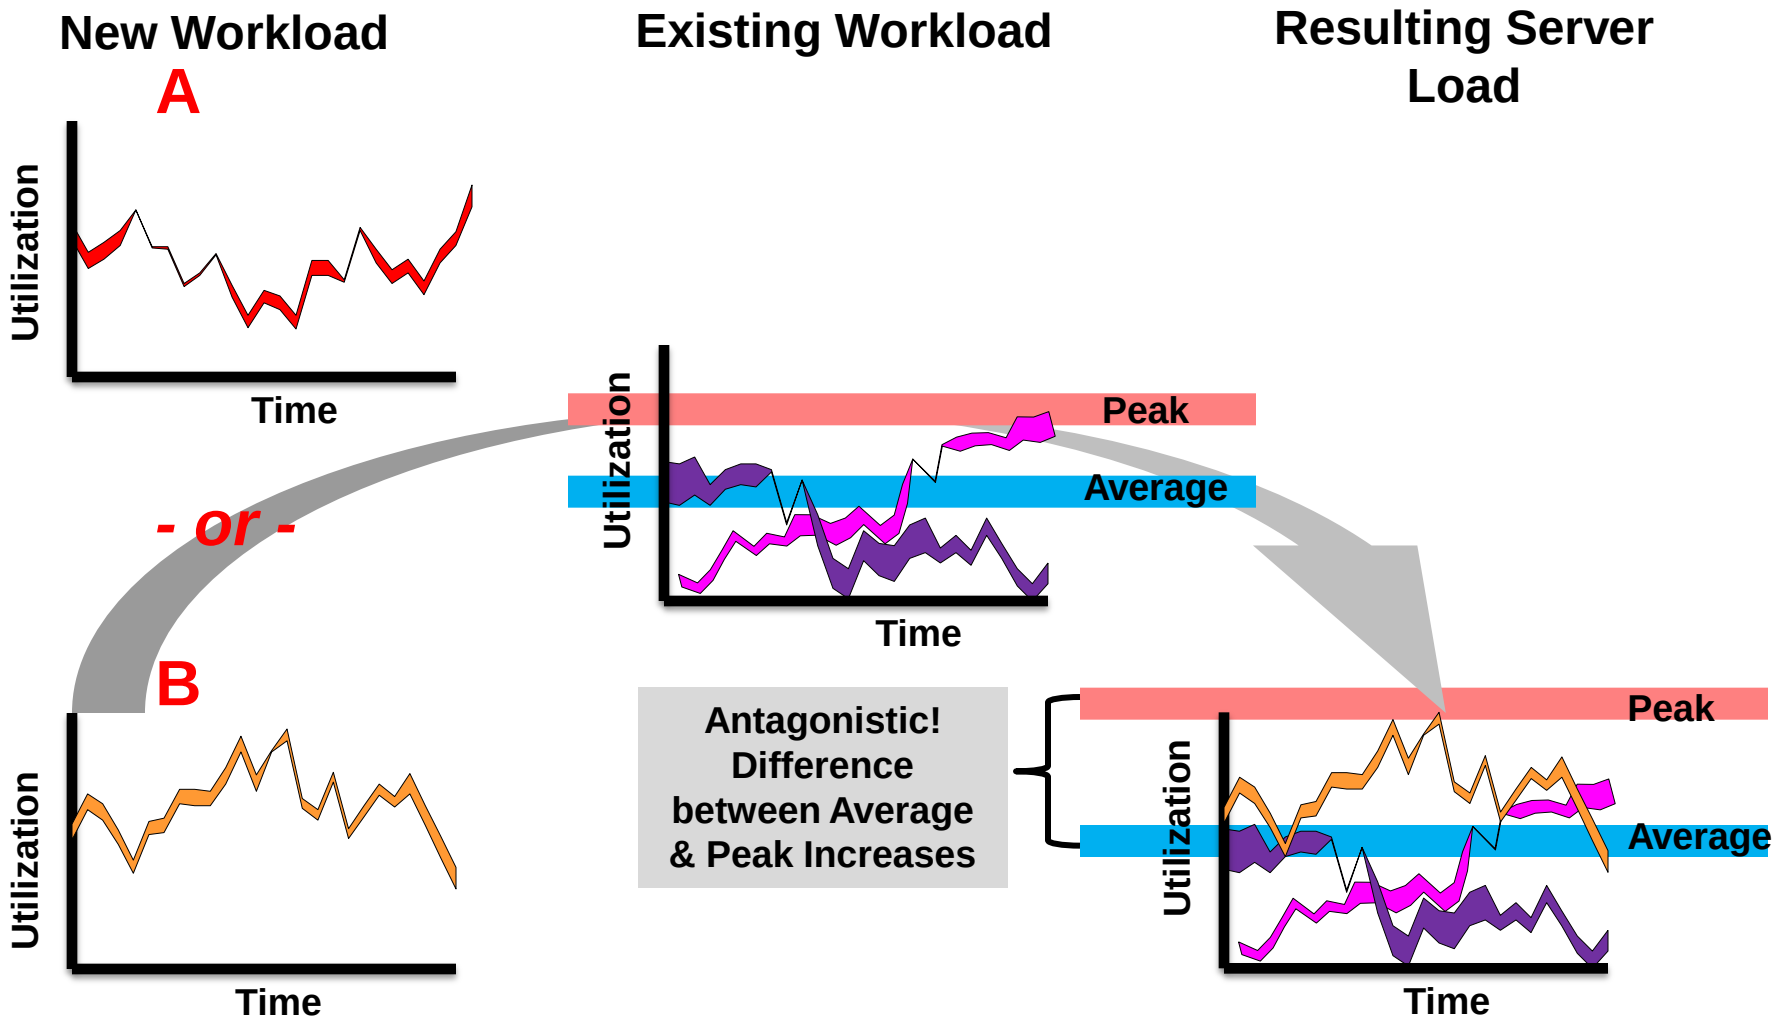

Cloud Pool Layout

What size is the pool?

- Typical customer deployments between 2-8 nodes
  - Some in the high 20s
  - Recommend a minimum size of 3 nodes in Cloud Pool
    - Protects from unplanned outage during planned events.
    - 2 nodes are manageable if a DR cluster can be used for planned outages
  - Can be larger for RAC One deployments in a Cloud Pool
    - Because of one active/one stand-by node requirement

Cloud Pool Layout

Considerations

- Consideration for application co-existence
  - Applications with similar SLA requirements are best suited to co-exist in a consolidated environment.
  - Do not mix mission critical applications with non-mission critical applications in the same consolidated environment.
  - Do not mix production and test/dev databases in the same environment.

- Over-subscription
  - It is possible to “over-subscribe” an application’s resource requirements in a consolidated environment.
  - Do not “over-subscribe” in a consolidated environment that contains mission critical applications.

Application Considerations
Custom Applications

- **Technical Questions**
  - Do you have unique characterset requirements?
  - Do you have same schema names for ease of administration?
  - Do you use public synonyms for convenience?
  - Are there any public DB links?
  - Do you have users with DBA privileges such as select_catalog_role, DBA etc?
  - Do you have same usernames, roles across your applications?
  - Do you need system privileges for any of your users?
  - Do you need special Configuration Parameters?
  - Do you need a specific Block Size
  - Do you need specific optimizer settings?

- **Business Questions**
  - Are there any data security requirements
  - Are there any regulatory compliance requirements that need to be met?

Summary

- Database Consolidation on Private Clouds offers enormous value to an IT organization
- Oracle technologies uniquely enable DB Consolidation
- Lots of Customers are already doing this

Additional Slides on Resource Management

Resource Management Considerations

- **Customer use cases**
  - Load: Start/Stop, seasonal, growing fast, fast bursting, slow peaks, batch, IT operations, …
  - SLA management: Response time
  - Problems: Runaway query, logon storm, Plan Flip
- **Resources**
  - Physical: CPU, memory, I/O, network, …
  - Logical: Concurrency, Deadlock
- **Management actions**
  - Planning: Headroom, connections, configuration, …
  - Run-time: Resource allocation, protection, planned, reactive
  - Type: Automatic, Manual, Disruptive
- **Workload type**
  - OLTP, OLAP, …

Customer Use Cases

- Understand important customer use cases
  - Known workload profiles
    - Start/stop
    - Predictable/Unpredictable
    - Batch
    - Operations
    - Slow/Fast peaks
  - Unknown events
    - Runaway query
    - Rapidly growing
    - Service level attacks
    - …
  - SLA management required

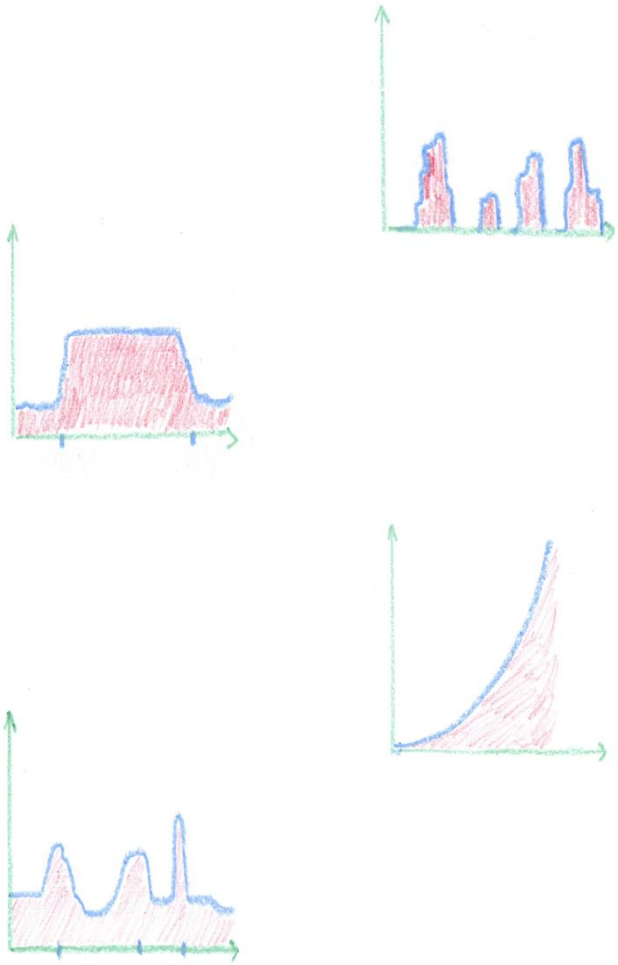

Resource Allocation – The Big Picture

Available intervention points

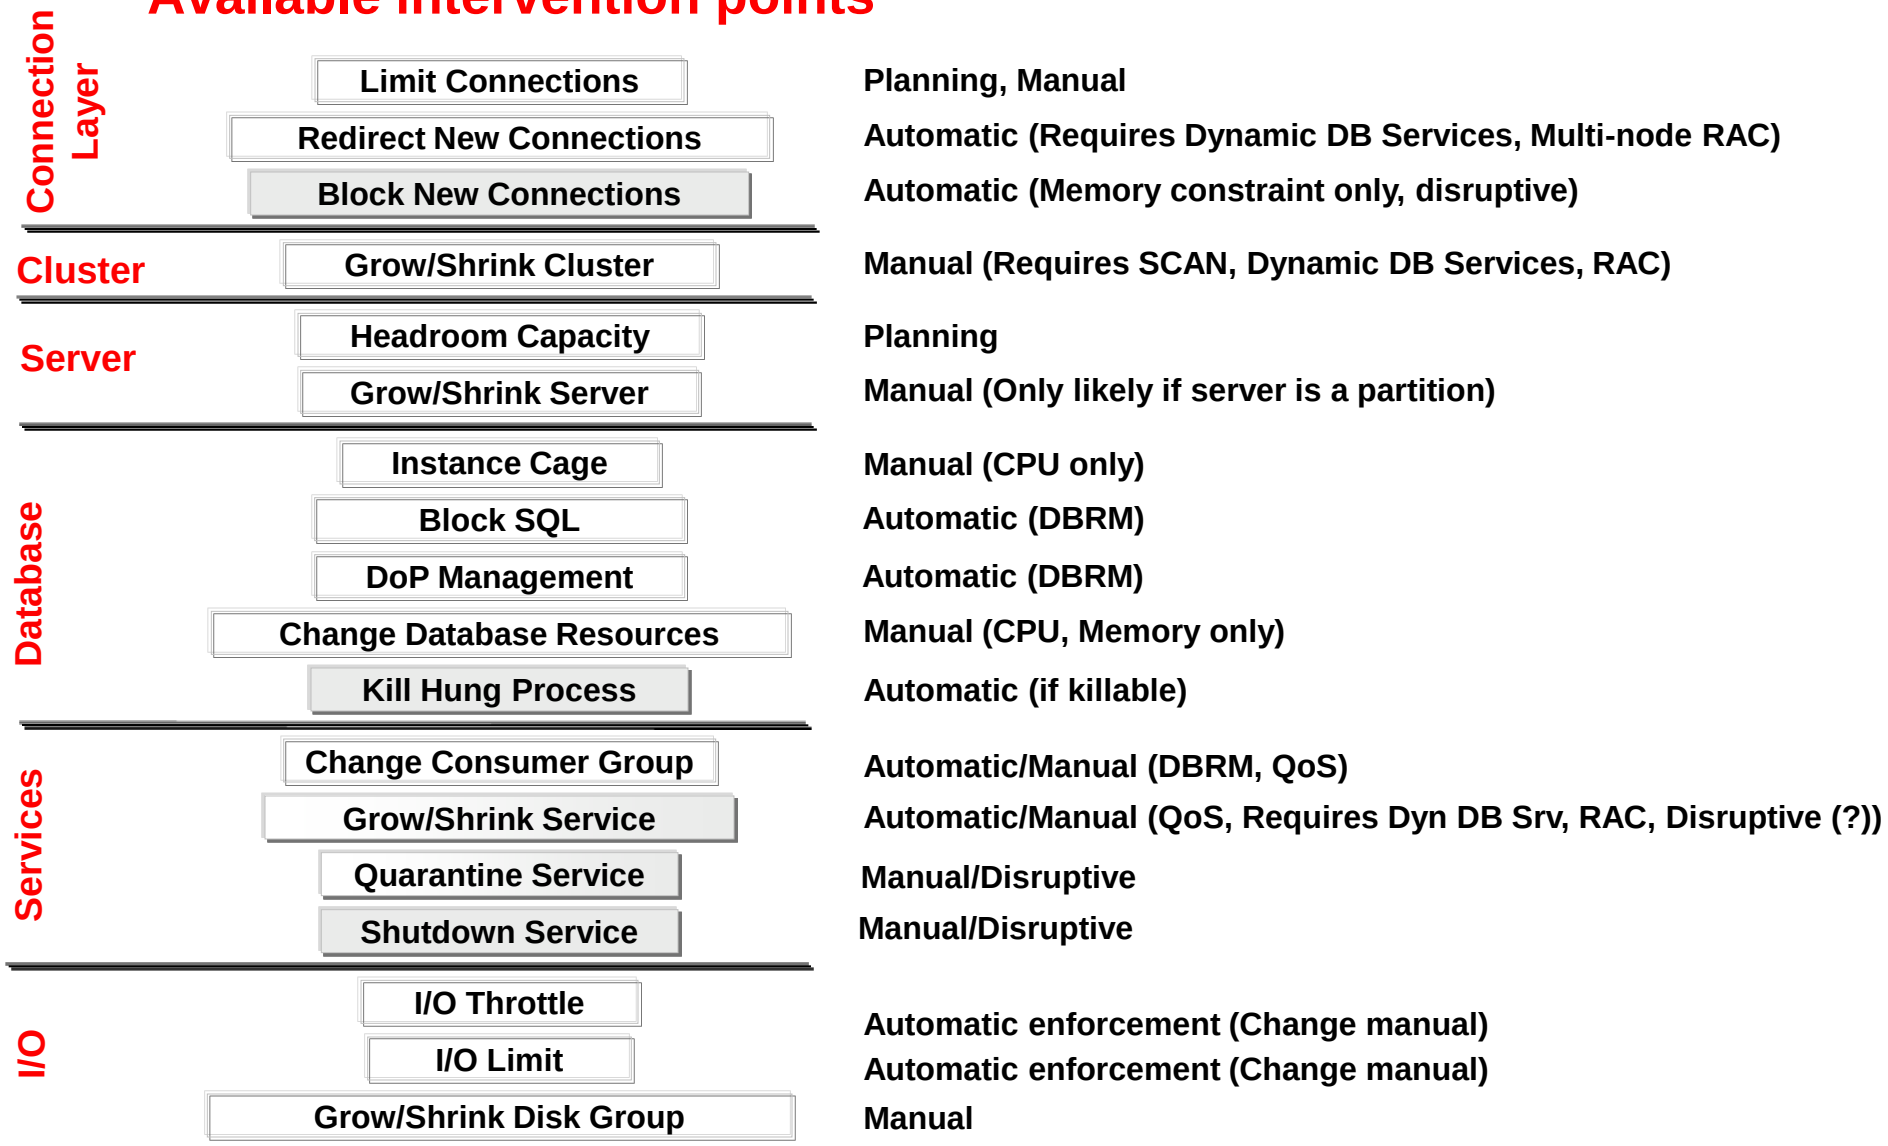

ORACLE

Disruptive operations | online operations

Keep it simple

- Lots of customers managing resources with headroom alone
  - Dell
- Or headroom + query management
  - CBA
- The tools are there for more active management
- Active management increases efficiency but at a management cost

Resource Management- Four Ways to Manage

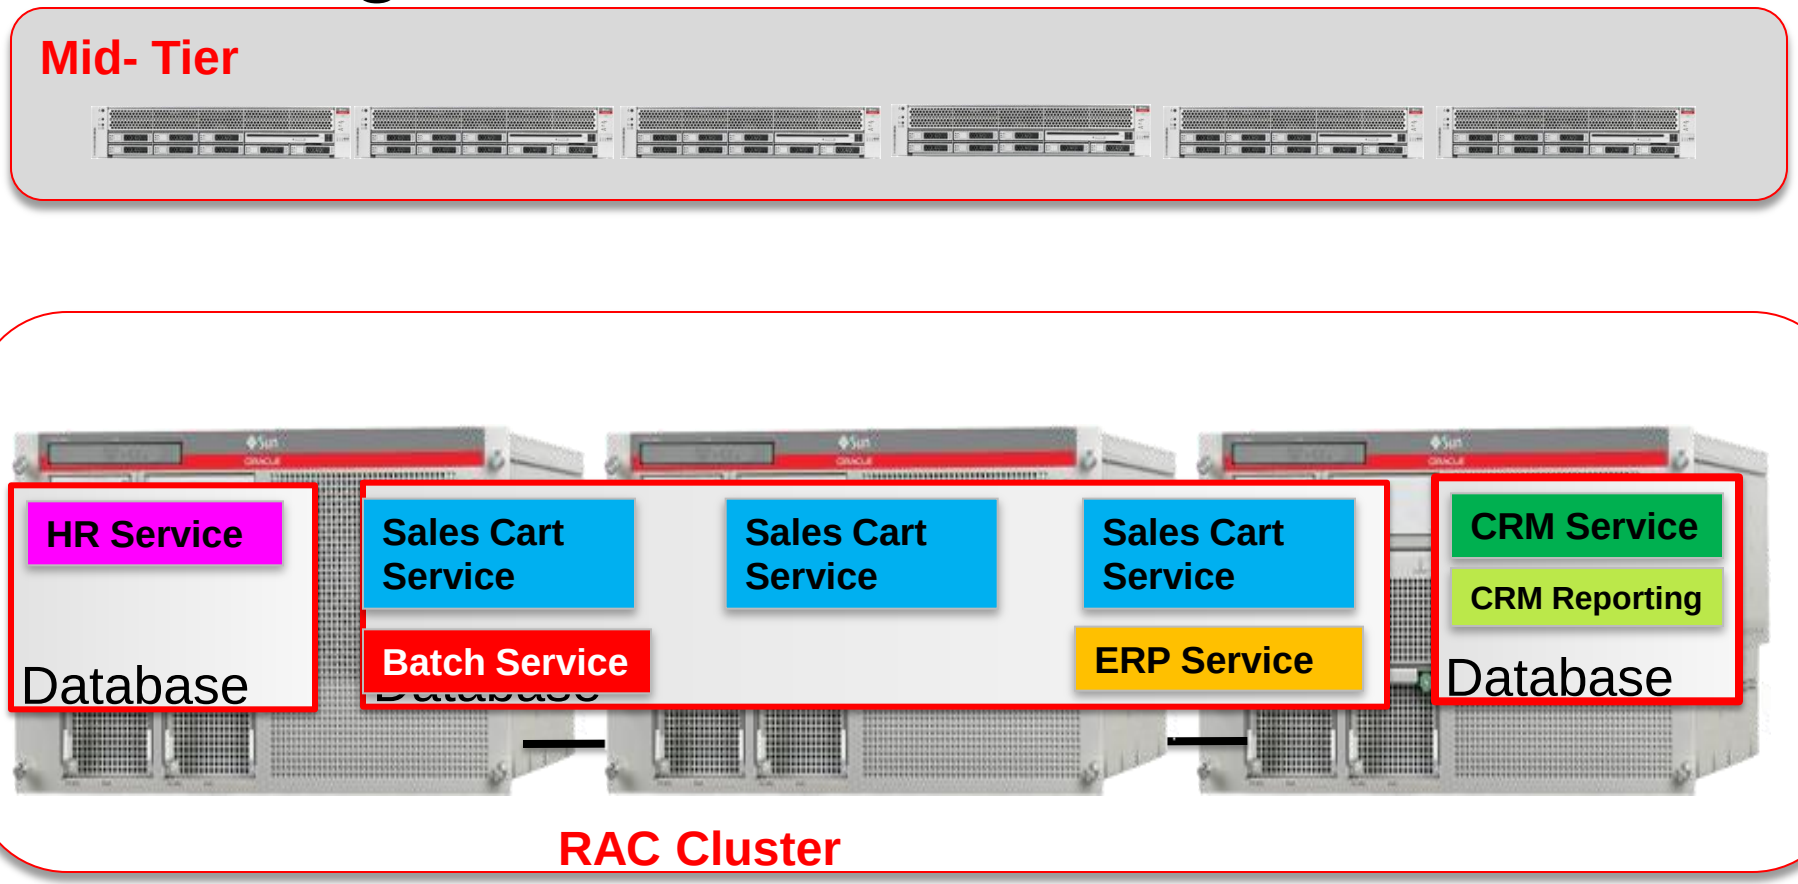

Resource Management- Four Ways to Manage

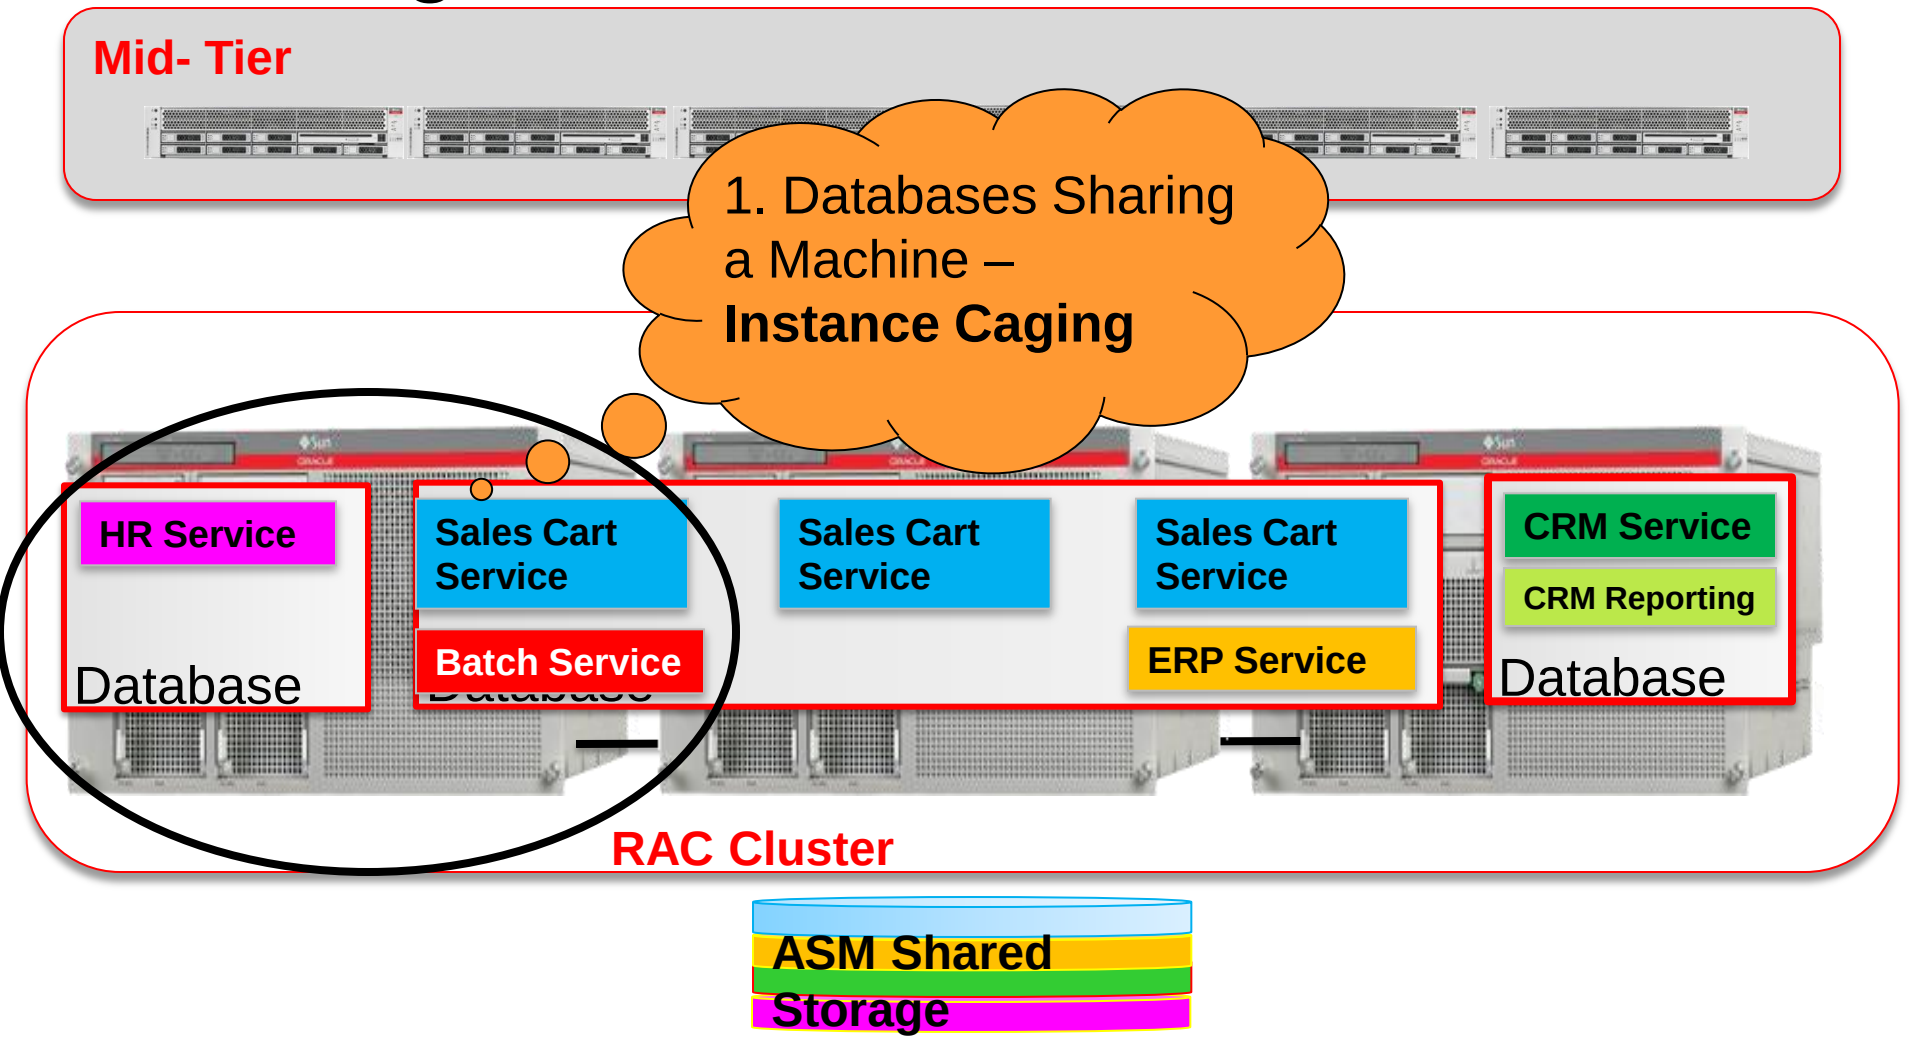

Resource Management- Four Ways to Manage

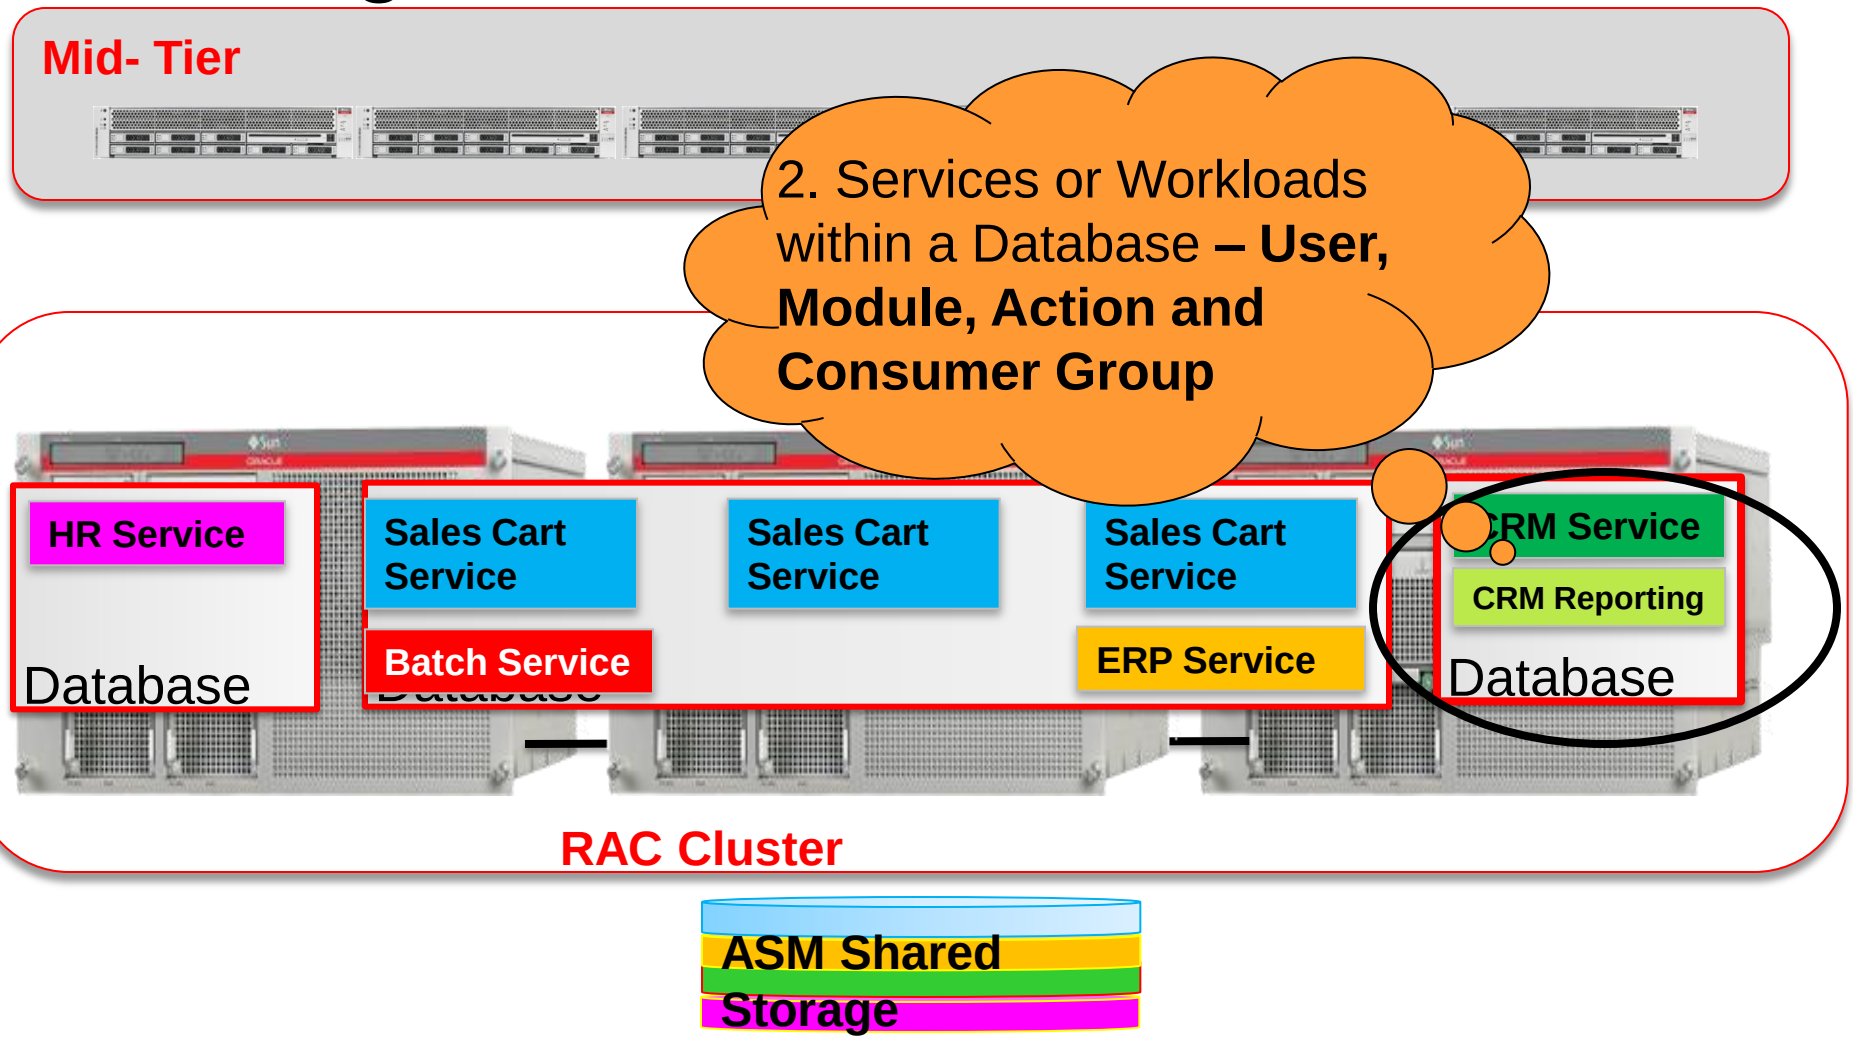

Resource Management- Four Ways to Manage

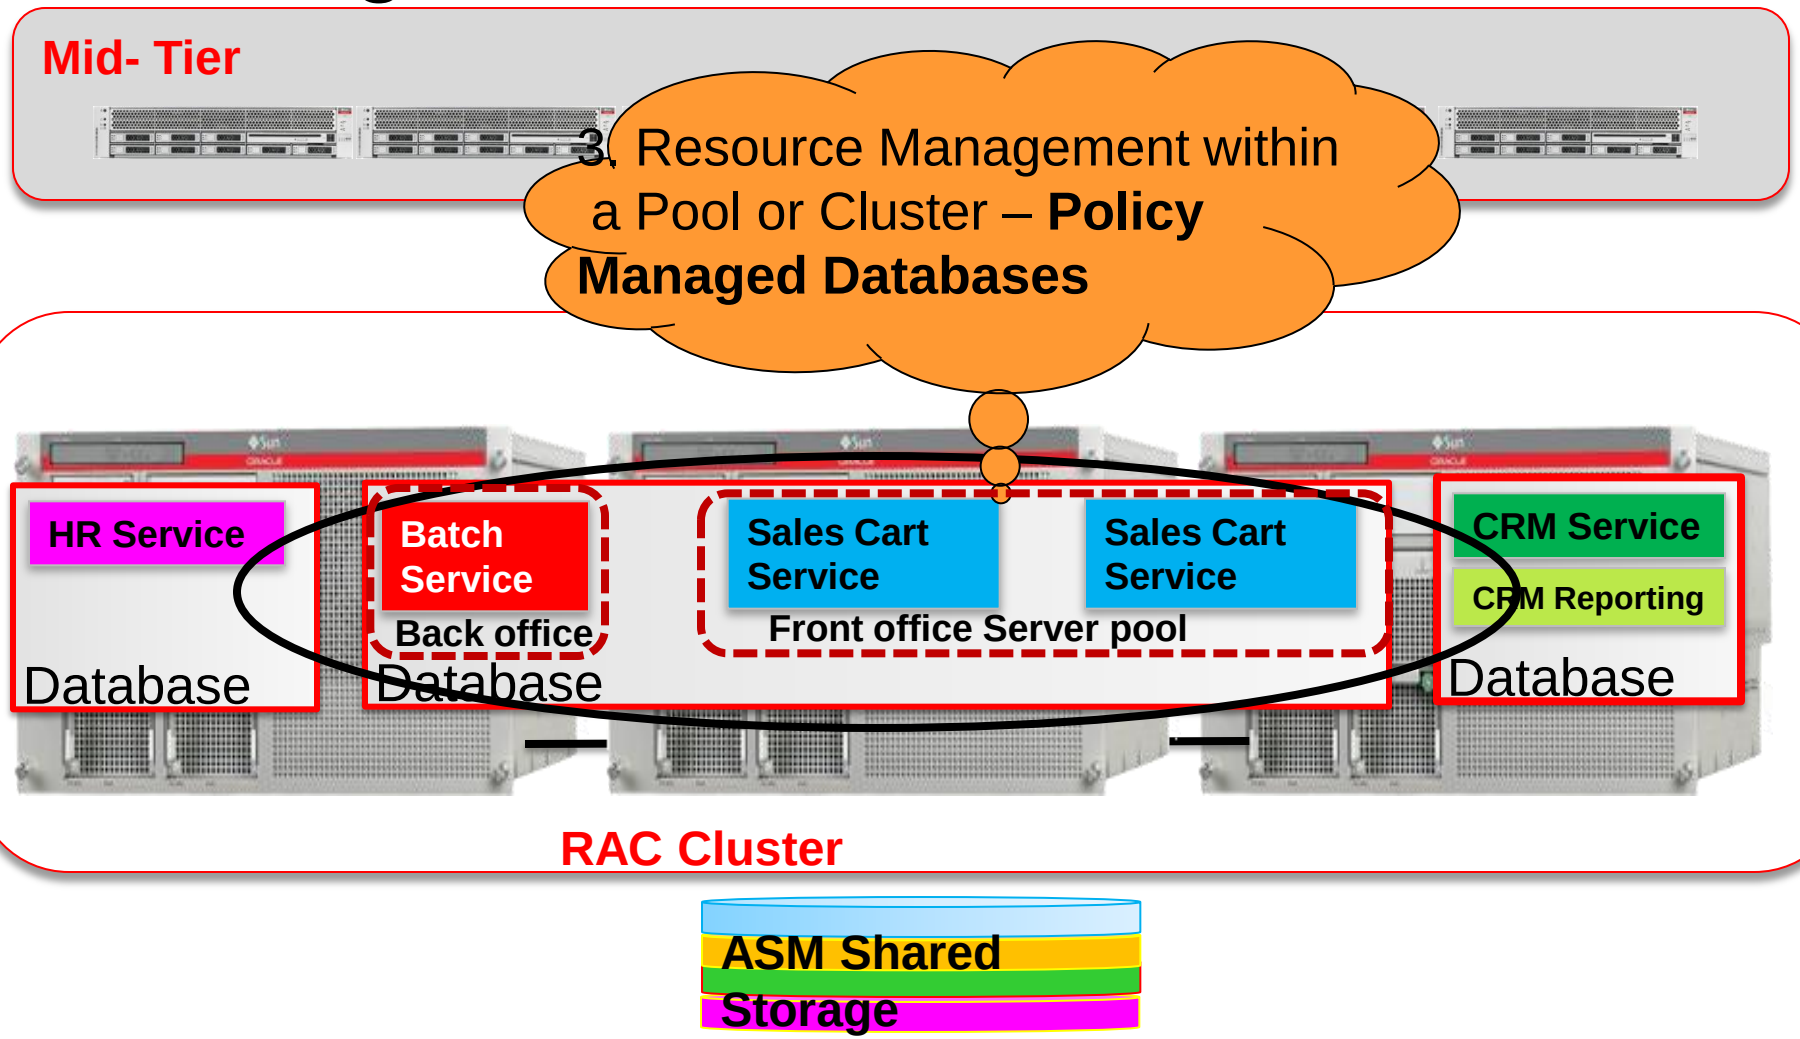

Resource Management- Four Ways to Manage

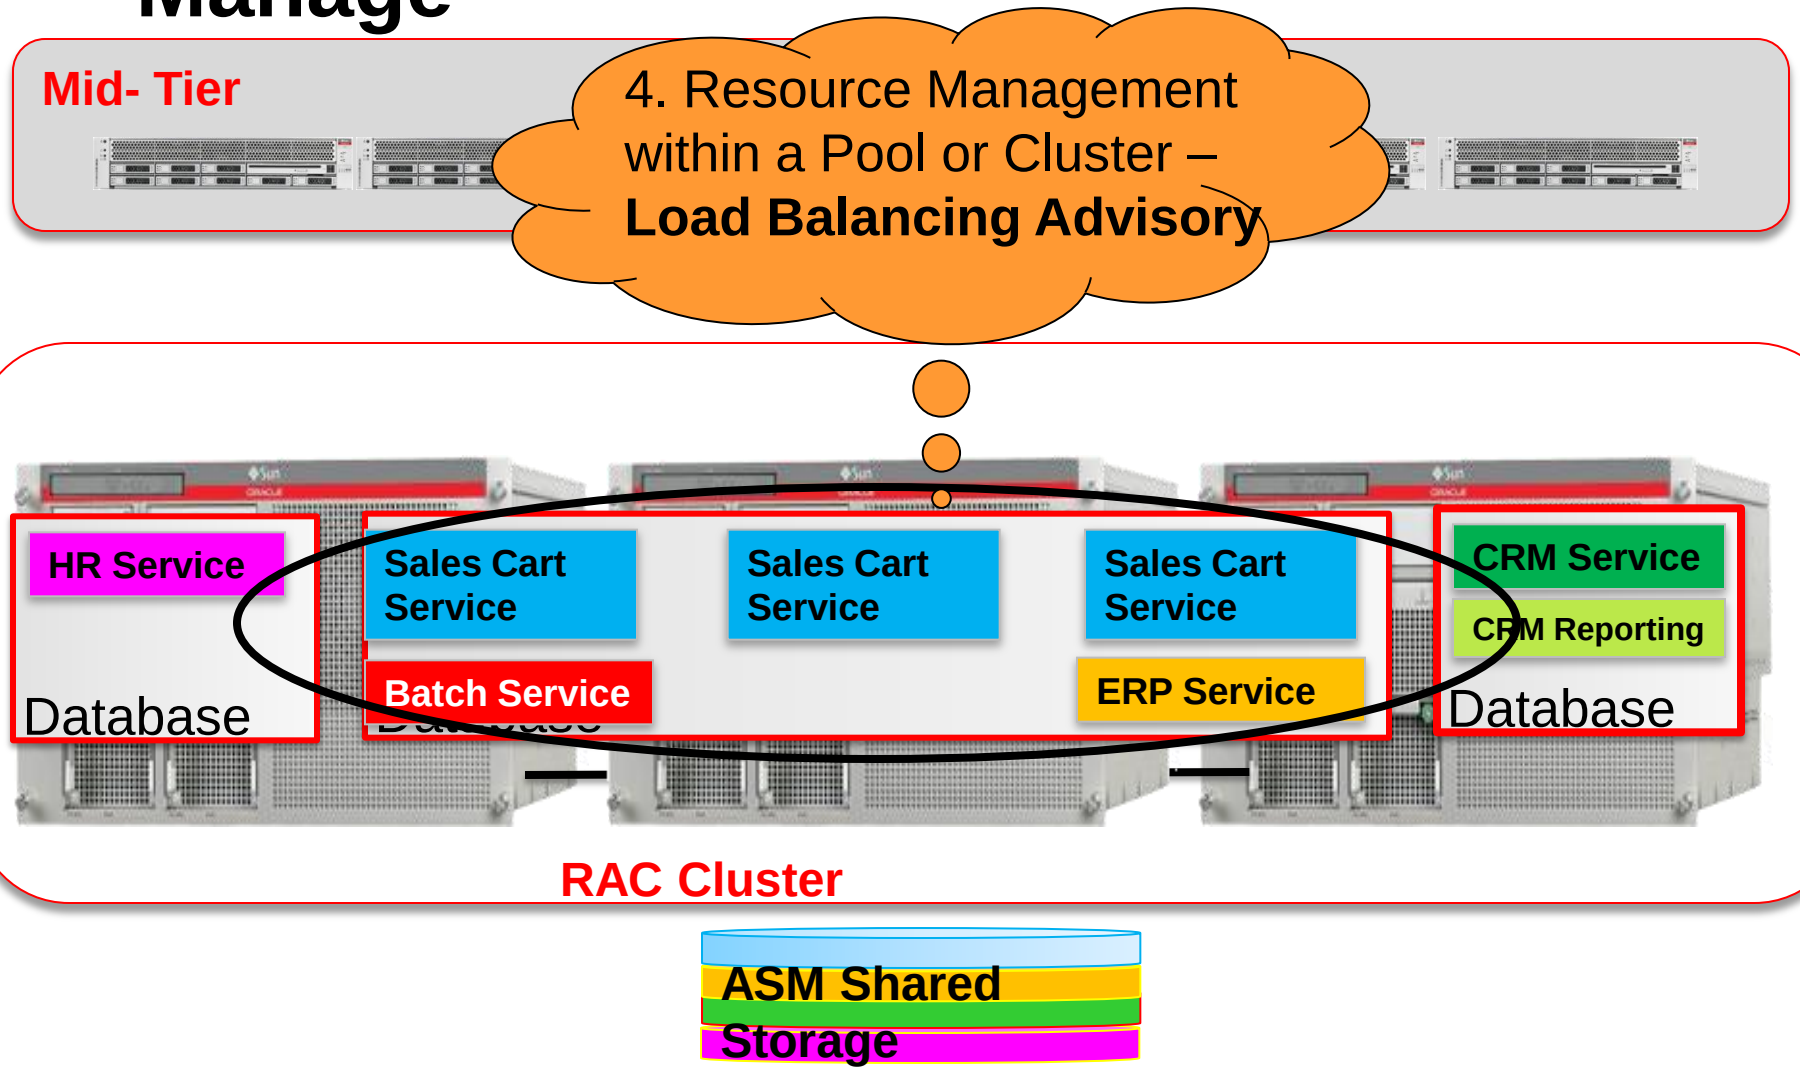

Resource Management – Within the Server using Instance Caging

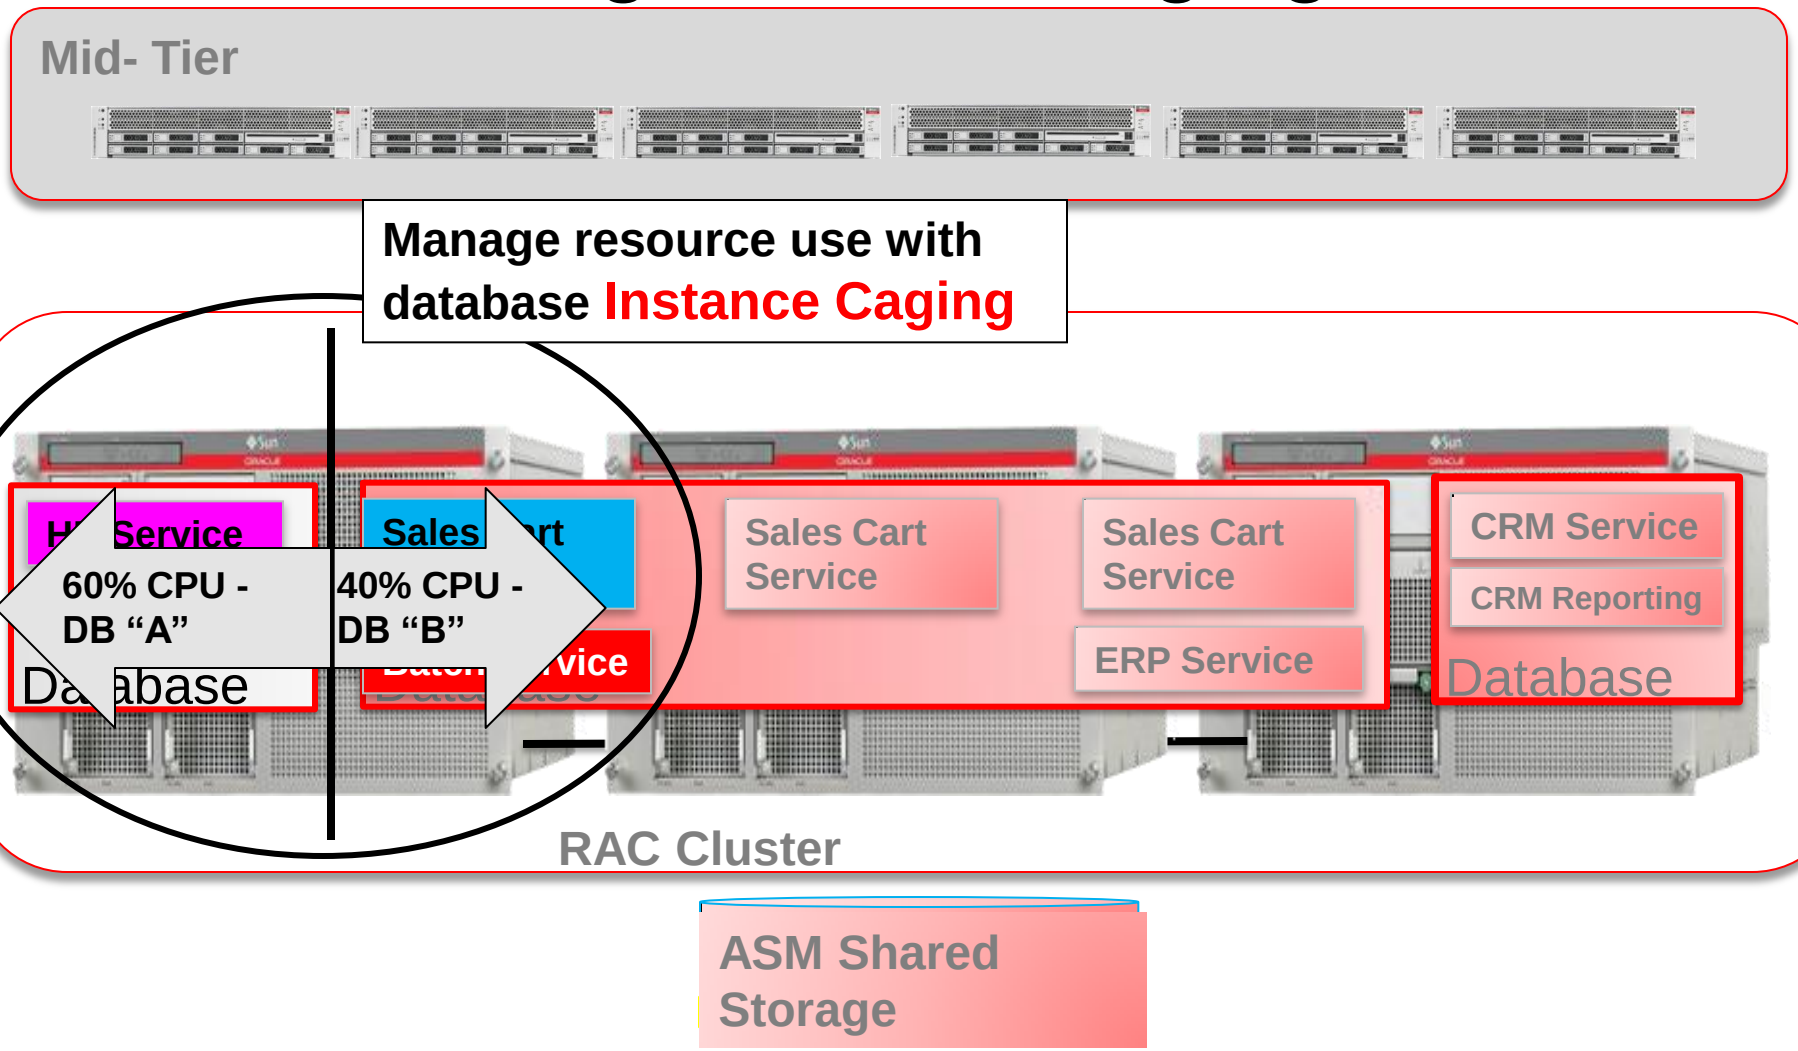

Resource Management Within a Database

Vice President of Sales
Quarterly report
(user,module, action)

Mid- Tier

In the resource plan, Tom Scott, running Quarterly Sales reports has a higher consumer group ranking than just “CRM Server” and per the plan directive, will get more cpu share when cpu is at 100% utilization

HR Service

Sales Cart Service

Sales Cart Service

Sales Cart Service

Batch Service

ERP Service

Database

CRM Service

CRM Reporting

Database

RAC Cluster

ASM Shared Storage

Resource Management – Load Balancing Advisory within the Pool

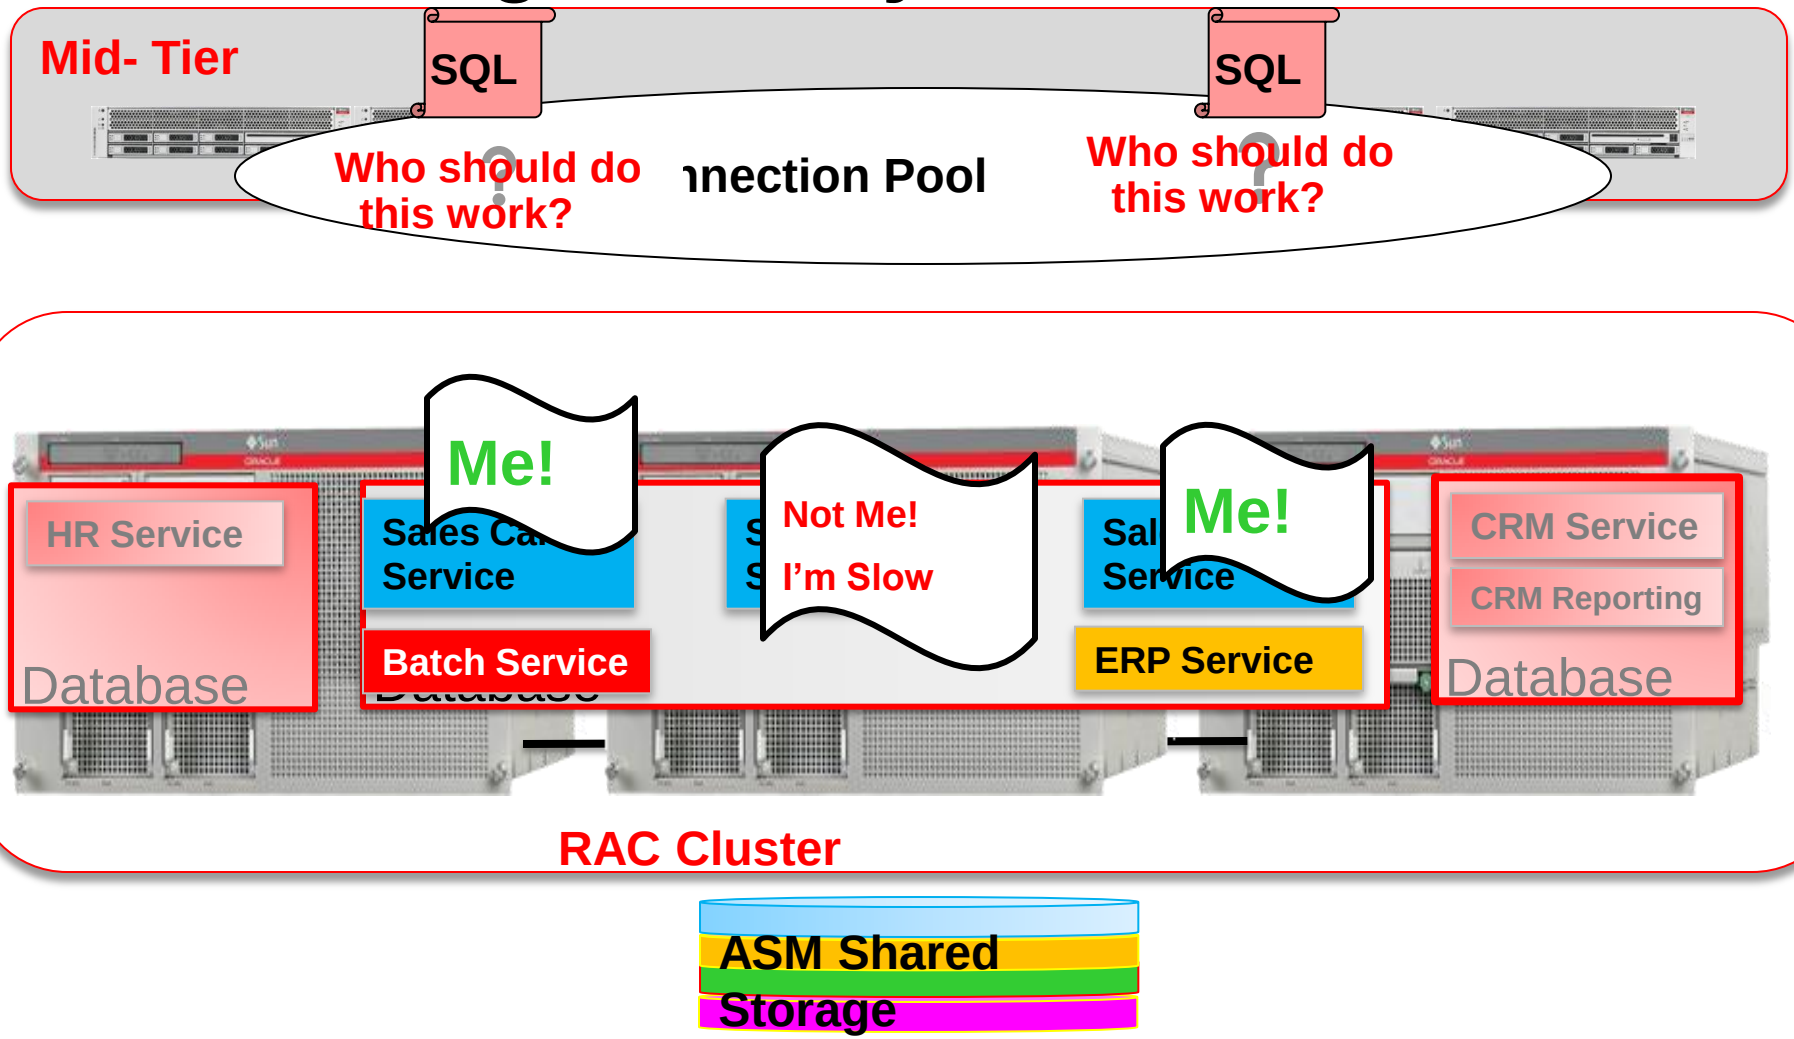

Resource Management – Policy Managed Databases