

Exadata Deployment Bare Metal vs Virtualized

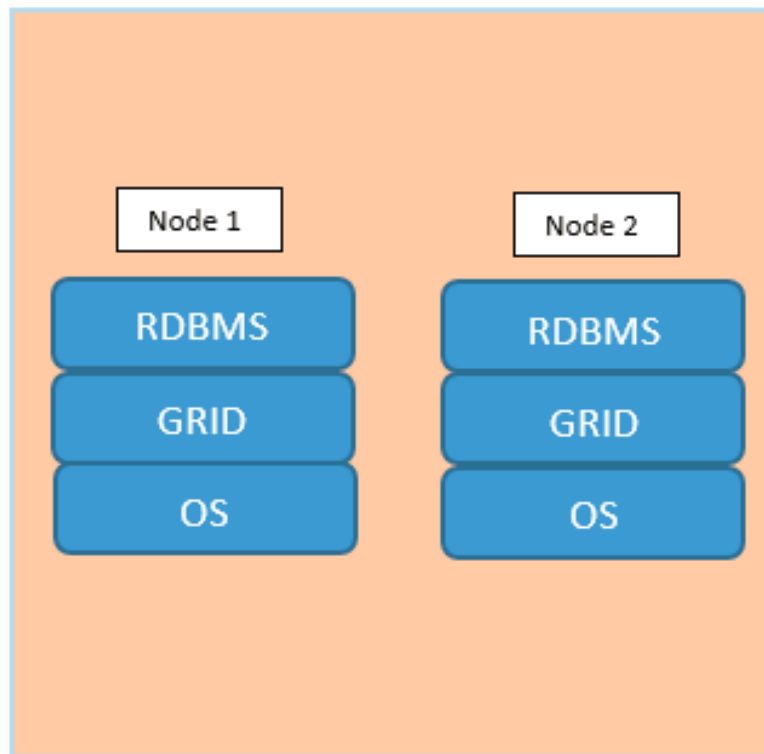
By
Umair Mansoob

Topics Covered

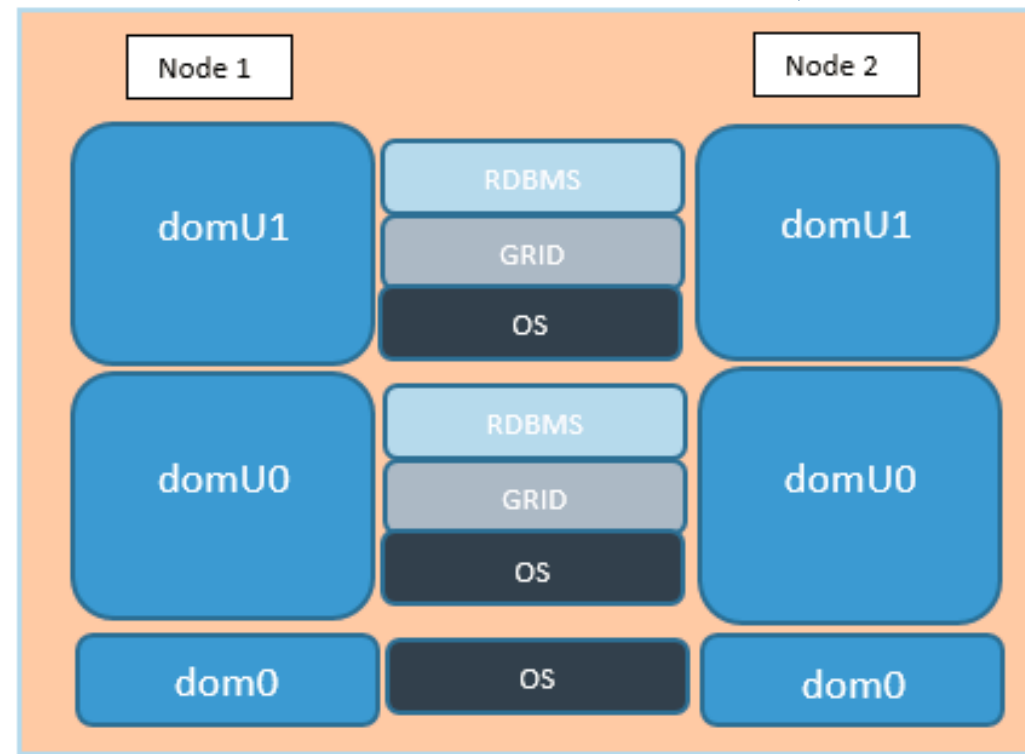
- Bare Metal vs Virtualized
- Bare Metal Layout
- OVM Layout
- OVM Considerations
- Exadata virtualization Used Cases
 - Oracle Licensing
 - Security & Compliance
 - GI Home / Database Home Maintenance
 - Workload Isolation
 - Database Consolidation
- Isolation vs Efficiently
- Migrating from physical to virtual
- Virtualization Pros / Cons

Bare Metal vs Virtualized

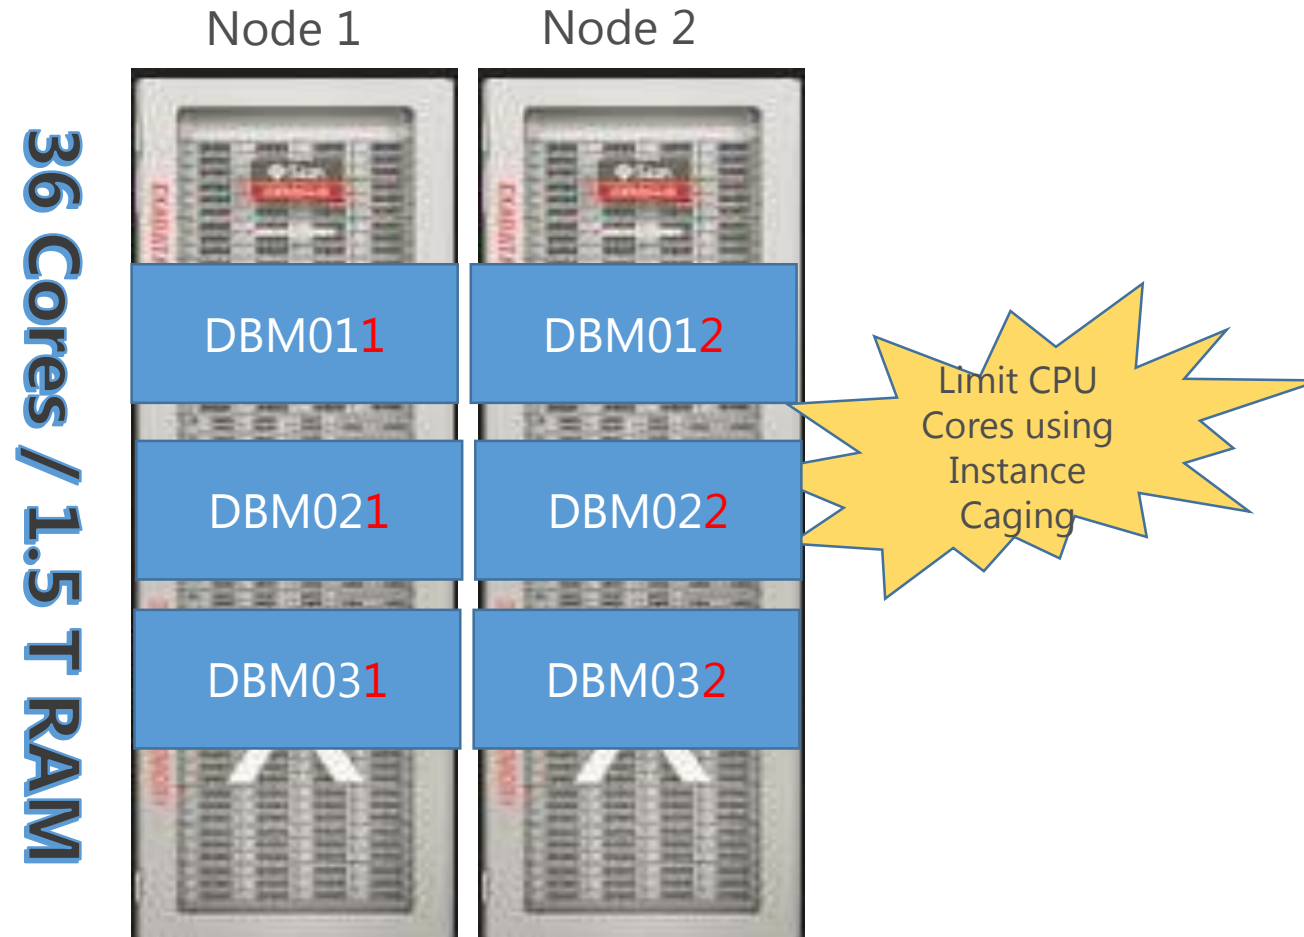
Bare Metal



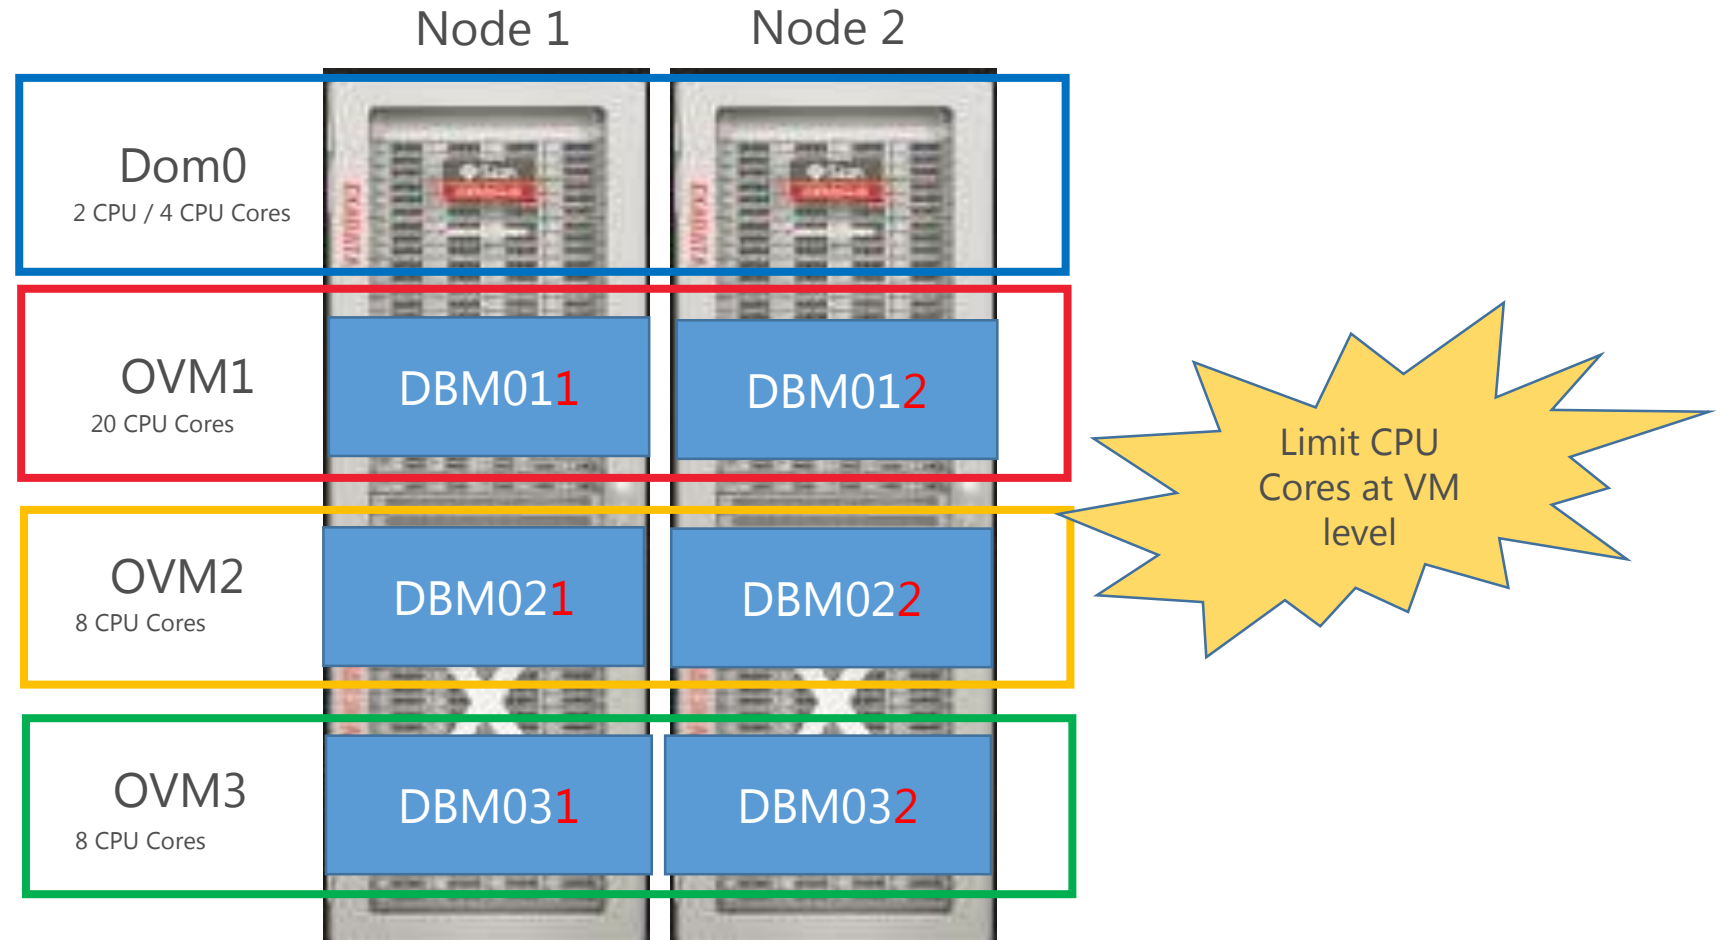
Virtualized



Bare Metal Layout



OVM Layout



Networking Considerations

Number of Ethernet Network Addresses Required

	X7-2				VM Impact
	Full	Half	Qtr	Eighth	
ILOM for Database Servers	8	4	2	2	-
ILOM for Exadata Storage Servers	14	7	3	3	-
Eth0 (mgmt) for DB Servers	8	4	2	2	+ 1x Num VMs
Eth0 (mgmt) for Exadata Storage	14	7	3	3	-
Mgmt port for IB switches	2	2	2	2	-
IP address for Ethernet Switch	1	1	1	1	-
IP address for PDUs	2	2	2	2	-
Mgmt Subnet Total	49	27	15	15	Mgmt + 1x VMs
Client Access Database Servers	8	4	2	2	= Num VMs
VIPs for Database Servers	8	4	2	2	= Num VMs
SCAN Addresses (per Cluster)	3	3	3	3	= 3x Num Clusters
Client Access Total	19	11	7	7	= 2 x Num VMs + 3x Num Clusters
Total	68	38	22	22	Mgmt + 3x Num VMs + 3x Num Clusters

Major Differences

Topic	How OVM differs from Physical
Hardware support	2-socket only
Cluster config	System has one or more VM clusters, each with own GI/RAC/DB install
Exadata storage config	Separate griddisks/DATA/RECO for each VM cluster; No DBFS disk group
Dbnode disk config	VM filesystem sizes are small; GI/DB separate filesystems
Software updates	Dbnodes require separate dom0 (Linux+fw) and domU (Linux) patchmgr updates
Exachk	Run once for dom0/cells/ibswitches, run once for <u>each</u> VM cluster
Enterprise Manager	EM + Exadata plugin + Virtualization Infrastructure plugin

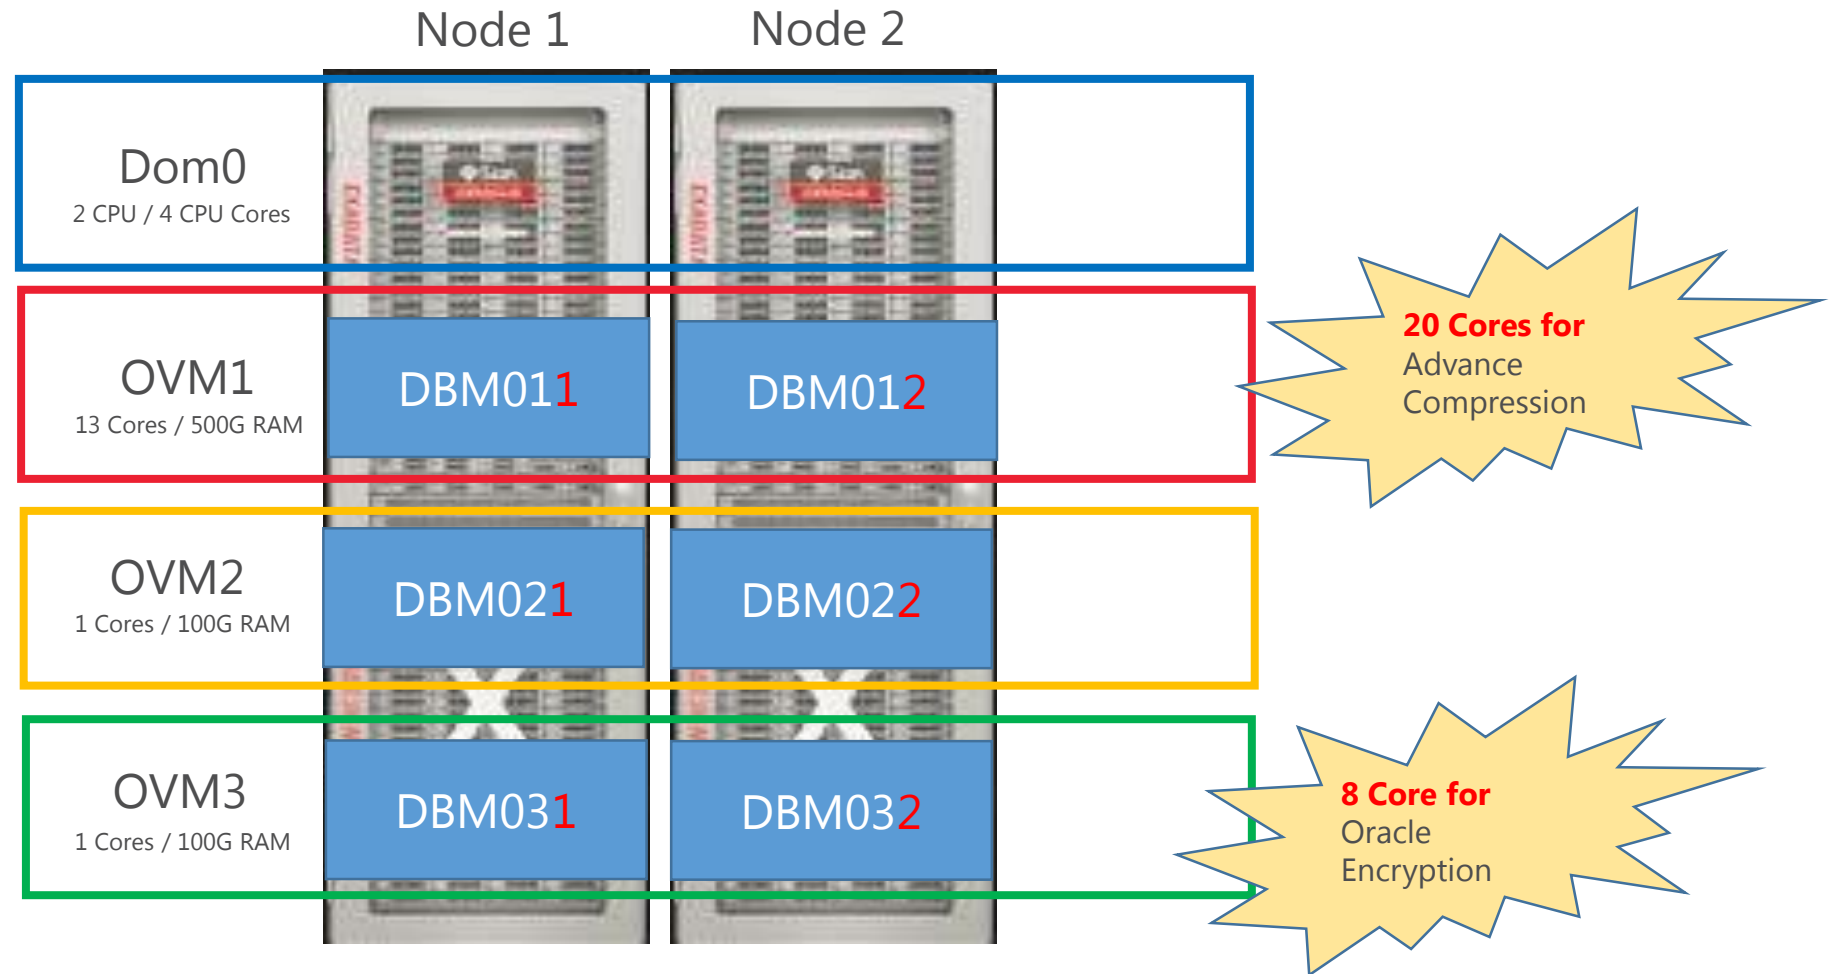
Maintenance Considerations

- Memory
 - VM memory can **not** be changed online
 - Do not over-provision physical memory
 - Minimum 16 GB per VM to support starter database, plus OS, Java, GI/ASM, etc.
- CPU
 - CPU over-provisioning is possible, but not recommended
 - Dom0 allocated 2 cores (4 vCPUs)
 - Maximum per VM is number of cores minus 2 for dom0
 - Number of vCPUs assigned to a VM can be changed online

Exadata Virtualization Used Cases

1. Oracle Licensing (Cost Saving)
2. Security & Compliance (Data Classification)
3. GI Home / Database Home Maintenance (Patching)
4. Workload Isolation (Test / QA / PROD)
5. Database Consolidation (Gold / Silver / Bronze)

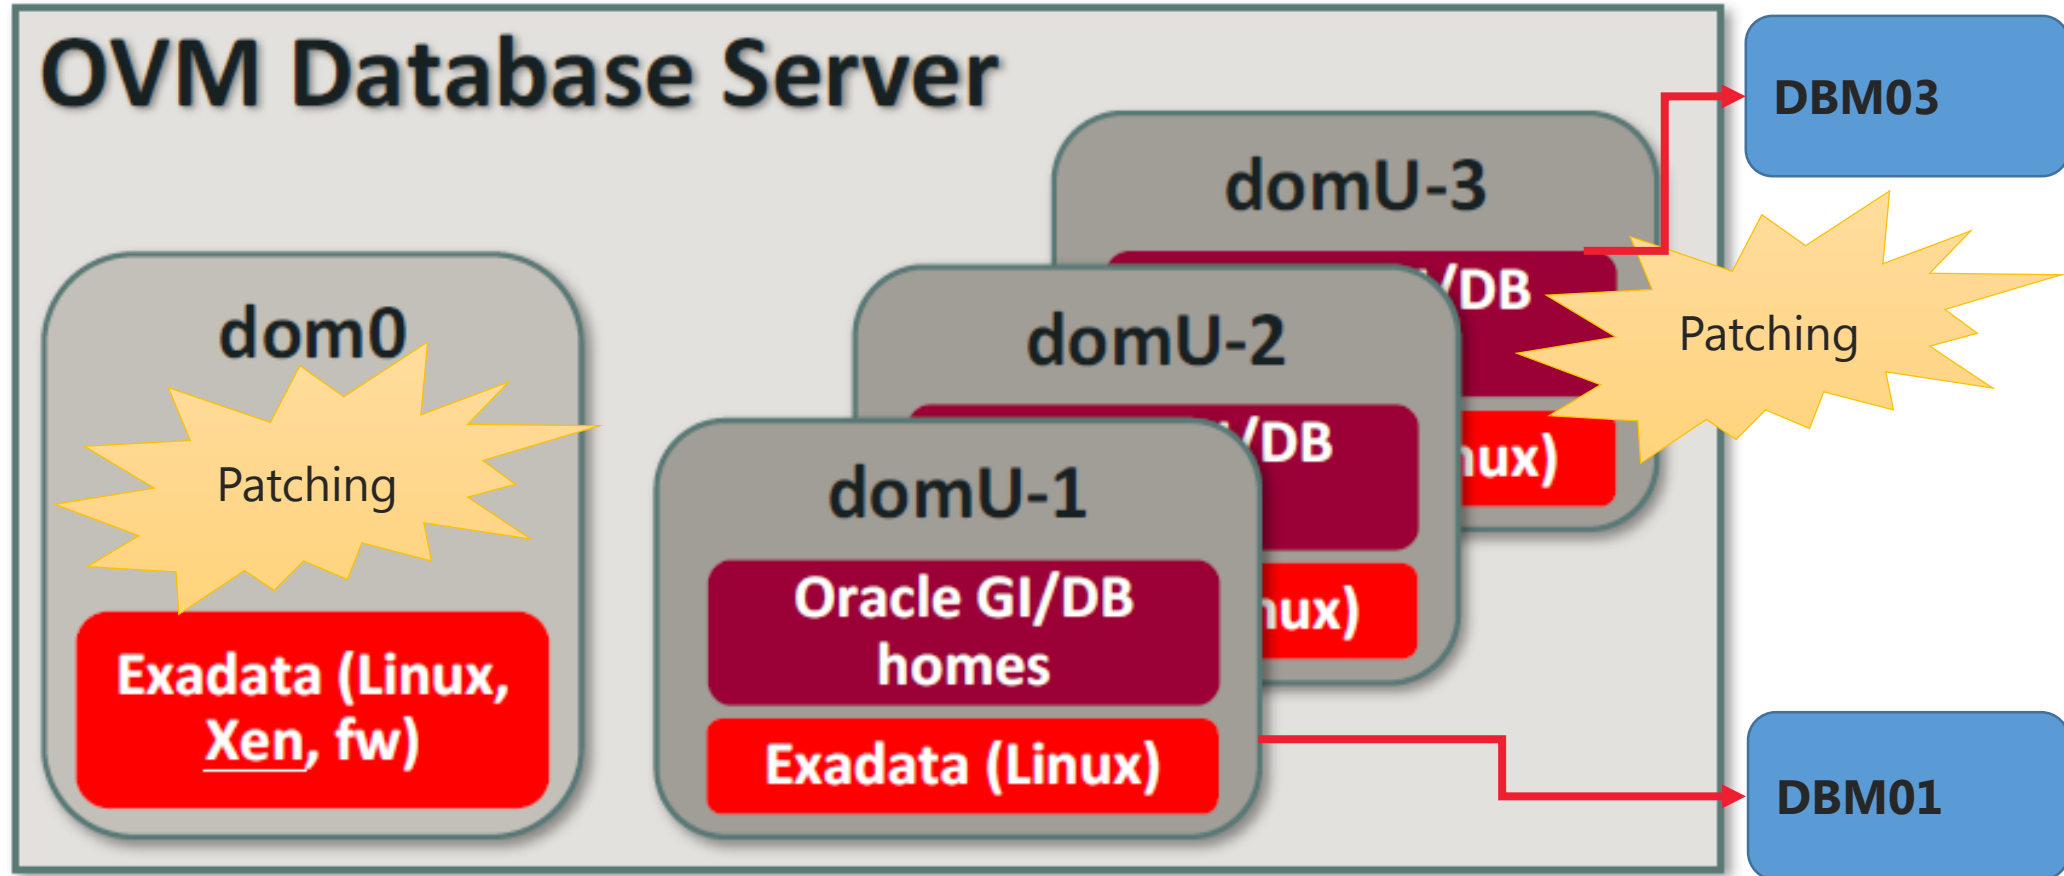
Oracle Licensing (Cost Saving)



Security & Compliance (Data Classification)

- There are different types of compliance requirements HIPPA, PCI DSS and Certifications, none of them will require you to virtualize Exadata machine
- You might be required to isolate your workload at database level or cluster level or operating system level by customer or client.
- if your databases contain sensitive client data from different business partners you might be required to isolate data at operating system level or even physical level

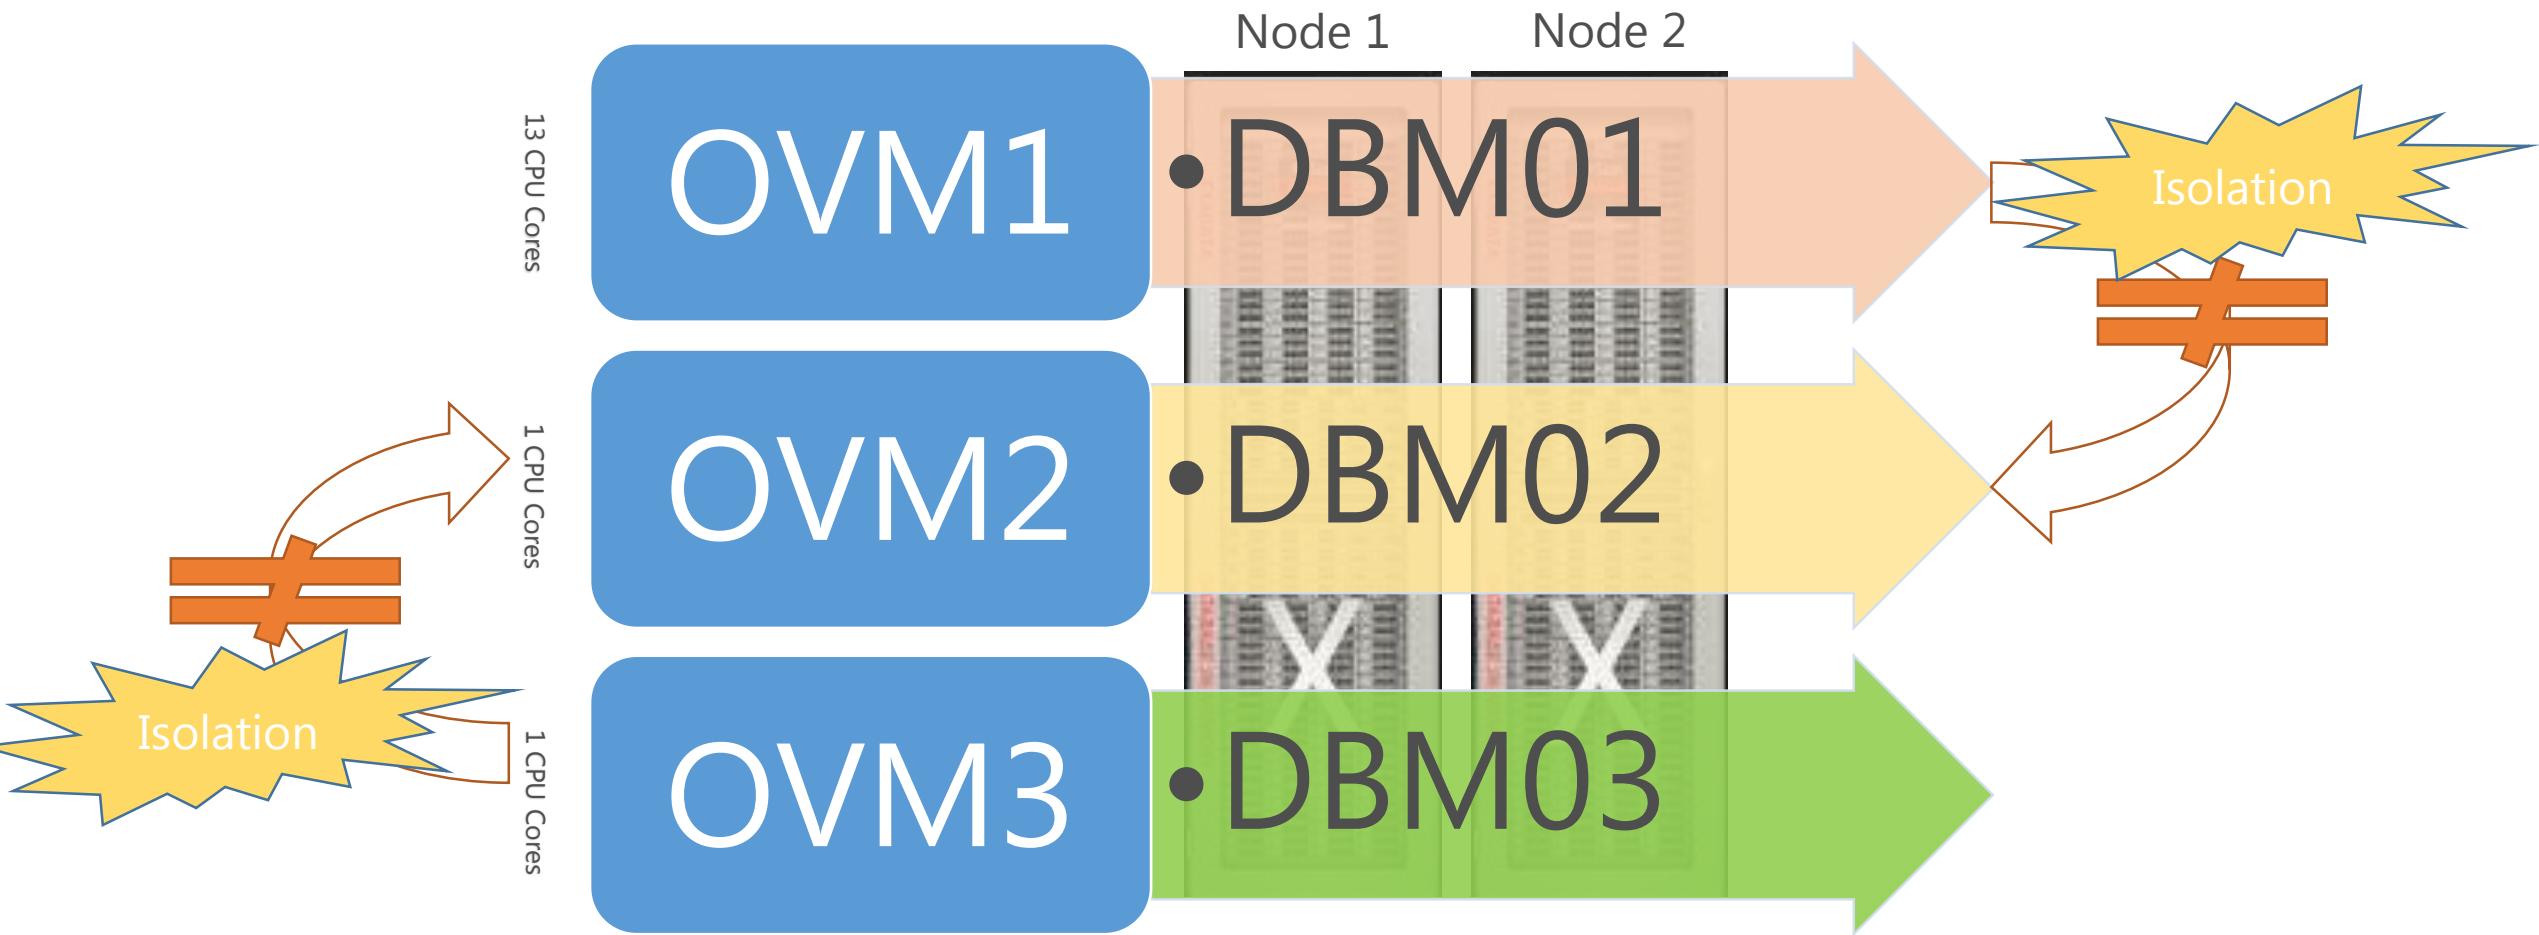
GI Home / Database Home Maintenance (Patching)



Workload Isolation (Isolation)

- Exadata VM combine with InfiniBand partitioning & VLAN tagging can provide great isolation
- Exadata VM provide maximum possible isolation that can be achieved within a single hardware
- VMs have good Isolation but poor efficiency and high management
- VMs have separate OS, memory, CPUs, and patching
- Alternate options are available to achieve workload isolation like instance caging , DBRM and IORM

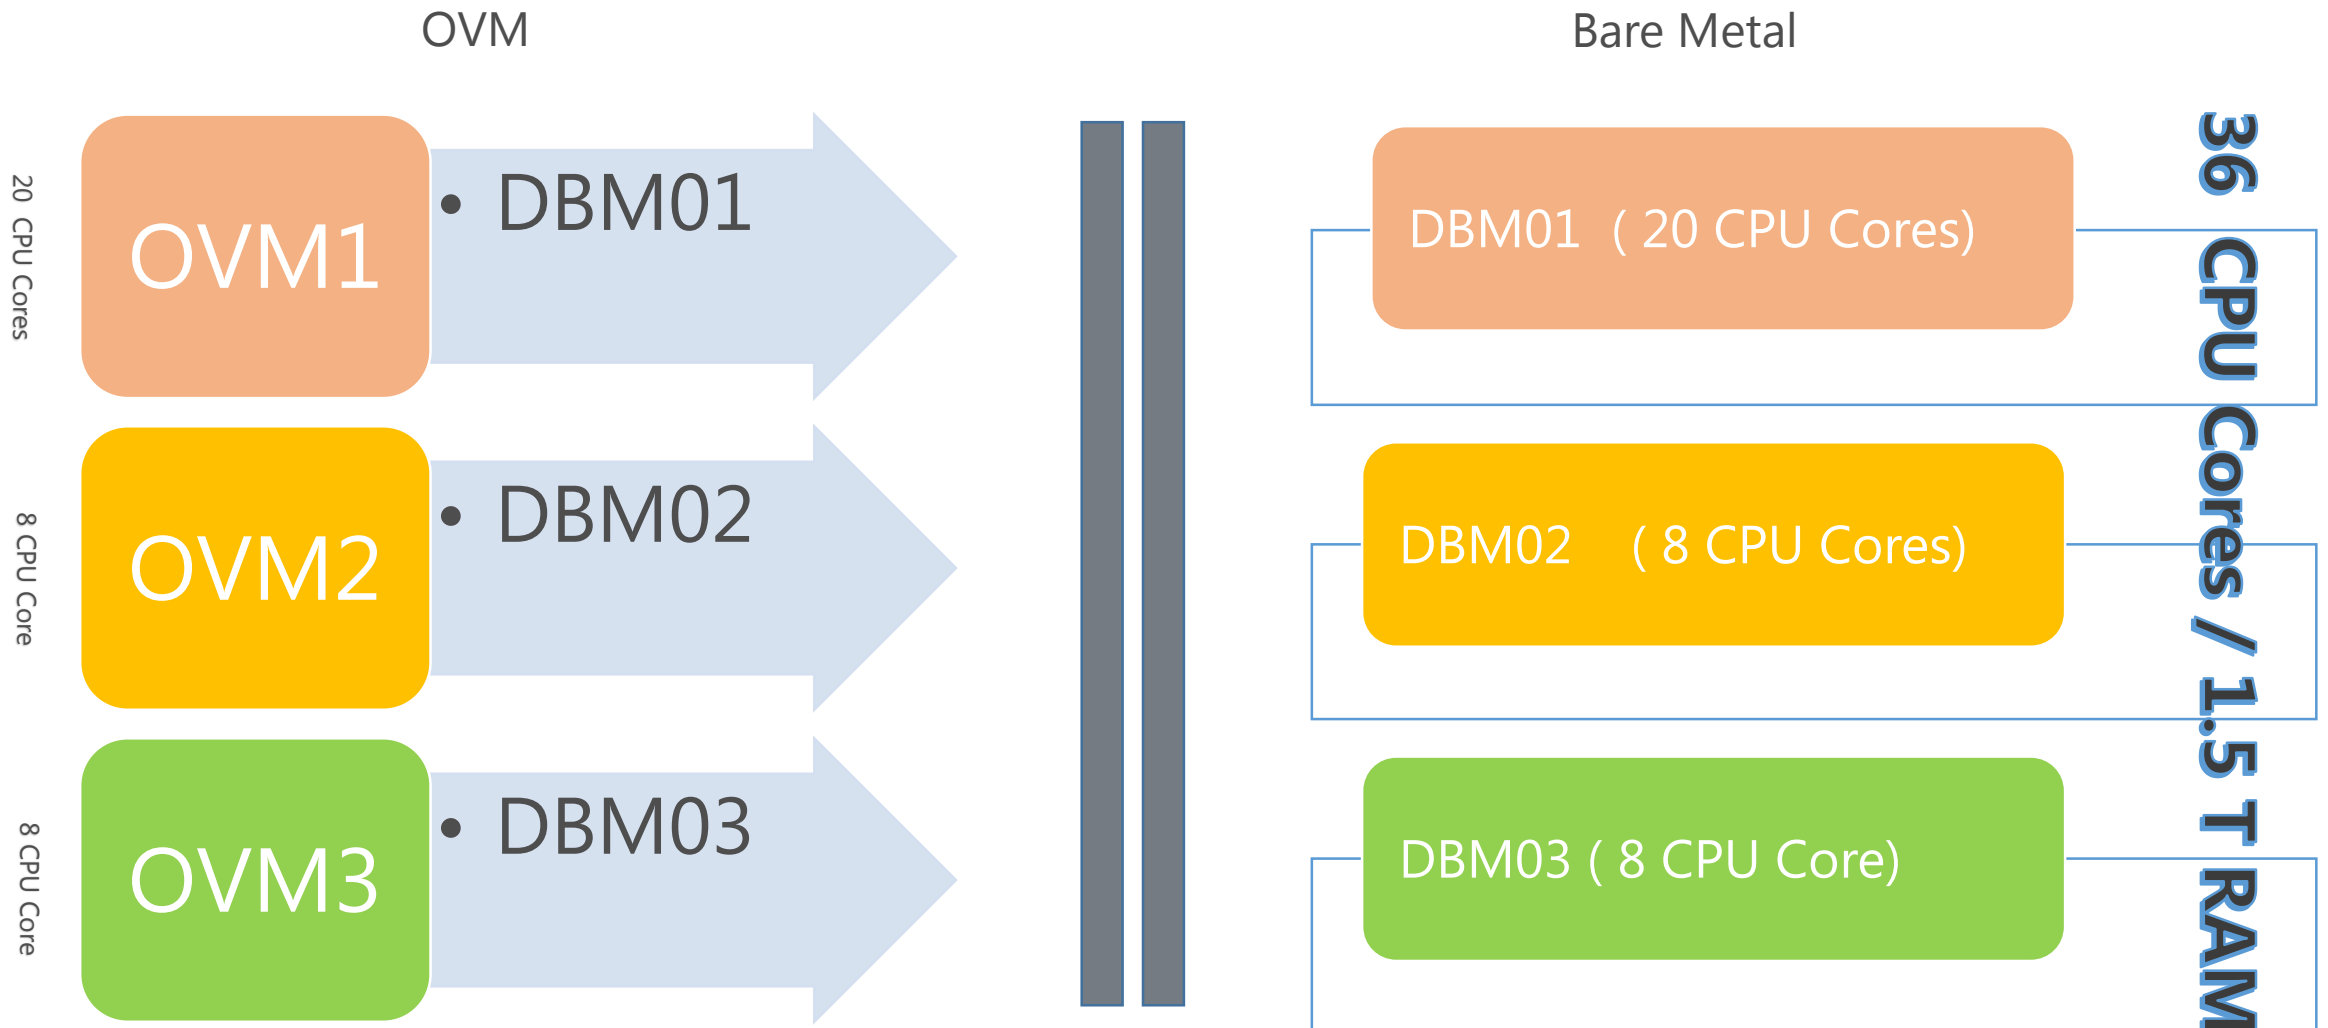
Workload Isolation with OVM



Workload Isolation without OVM

- The Oracle Database Resource Manager limits the amount of CPU that the database instance consumes.
- The `cpu_count` database parameter specifies the limit.
- Exadata I/O Resource Management (IORM) can be used to manages I/O resources
- Oracle Database Resource Manager has been enhanced to work with IORM to provide database resource management.
- Multiple Database home can be using for patching

Workload Isolation without OVM



Database Consolidation (Silver / Gold / Bronze)

- Database consolidation in a single OS is highly efficient but less isolated
- DB Resource manager isolation adds no overhead
- Resources can be shared much more dynamically
- Hosting, cloud, cross department consolidation, test/dev, non-database or third party applications
- Database consolidation based of SLA's and availability can be a good used case for virtualized Exadata Machine

Isolation vs Efficiently

- VMs have good Isolation but poor efficiency and high management
- Virtualization will introduce complexity and inefficiently but provide greater isolation
- Backup/Restore of Virtualized Environment
- OVM can be used to create virtual machines with specific number of cores to reduce licensing requirements
- Workload Isolation can be achieved using DBRM and IORM.

Migrating from physical to virtual

- Dynamic or online method to change physical to virtual
 - Data Guard or backups can be used to move databases – minimum downtime
 - Convert one node or subset of nodes to virtual at a time
- Dynamic or online method to change physical to virtual using any of the procedures below
 - Migrate to OVM RAC cluster using the existing bare metal Oracle RAC cluster with zero downtime
 - Migrate to OVM RAC cluster by creating a new OVM RAC cluster with minimal downtime
 - Migrate to OVM RAC cluster using Oracle Data Guard with minimal downtime
 - Migrate to OVM RAC cluster using RMAN backup and restore with complete downtime

Note : - Migrating from OVM to Bare Metal will require complete reimage / redeployment

Exadata Virtualization Pros/Cons Summary

Pros	Cons
Better Isolation (Resource are Hard partitioned)	Complexity (Network / Database / Storage layer)
Database Consolidation	Required more maintenance (Patching, etc.)
Oracle License Savings	Inefficient resource utilization
Oracle GI / DB Home Patching	Additional Layer of BUGs (OVM Layer)
Security & Compliance	Virtualization overhead

References

- <https://www.oracle.com/technetwork/database/availability/exadata-ovm-2795225.pdf>
- <https://www.oracle.com/technetwork/database/features/availability/exadata-consolidation-522500.pdf>
- https://docs.oracle.com/cd/E24628_01/doc.121/e27442/ch6_virtualization.htm

Thank you.

Umair Mansoob

773-297-2061

umairmansoob@gmail.com

