



Best practices guide

# Reducing costs in your Oracle database environment

Best practices for the modern data center



**Hewlett Packard  
Enterprise**



What if you could reduce the cost of running Oracle databases and improve database performance at the same time? What would it mean to your enterprise and your IT operations?

Oracle databases play a critical role in many enterprises. They're the engines that drive critical online transaction (OLTP) and online analytical (OLAP) processing applications, the lifeblood of the business. These databases also create a unique challenge for IT leaders charged with improving productivity and driving new revenue opportunities while simultaneously reducing costs.

## The challenge for IT leaders

Oracle databases are critical to your business. They are necessary, but they are also costly and can be challenging to manage. With the rapid growth in enterprise data, IT leaders are seeing several issues emerge in their Oracle environments:

- **Database growth:** As data volumes grow, the Oracle landscape becomes more complex and difficult to manage. Studies show that Oracle licenses are often oversubscribed.
- **High costs:** Oracle license and support costs consume too much of the IT budget, reducing the resources available to support innovation. Oracle licensing can be unnecessarily complex and license audits often cause disruption and reveal unexpected costs.
- **Deployment:** Workloads best suited for a scale-up design may be deployed with a scale-out architecture that limits performance while requiring the use of costly Oracle Real Application Clusters (RAC). The time and effort required by complex deployments reduces the net value of Oracle solutions.
- **Oracle "lock-in":** Locking in to the full Oracle ecosystem is costly and limits your ability to leverage other technologies and products.

Many organizations run Oracle on UNIX® or Oracle's proprietary Linux® distribution, adding to the cost of maintaining and growing those deployments. In many cases, the existing Oracle infrastructure is outdated, putting additional pressure on IT to find creative ways to manage growing data and transaction volumes and more demanding service-level agreements (SLAs). CIOs and VPs of infrastructure face some difficult questions:

- What can we do to reduce the cost of Oracle databases without reducing service levels?
- Can we enable growth without adding cost and complexity?
- How do we get the absolute best value from the dollars we spend on Oracle?
- How do we support enterprise efforts to move to industry-standard Linux platforms?
- What virtualization technology is best for our Oracle environment?

To get to the answers, let's begin by taking a look at the Oracle ecosystem.

## The Oracle Database Ecosystem

Oracle promotes a closed ecosystem including five key components: Oracle Database, Oracle RAC, Oracle Linux, Oracle VM Servers, and Oracle Exadata hardware. While Oracle can run on a wide variety of UNIX or Linux hardware, Oracle encourages this proprietary approach.

Oracle Database is the most-used relational database in the world and is supported on most major computing platforms. The latest version, Oracle Database 12c, includes options for a multitenant architecture, enabling database consolidation and in-memory data processing capabilities to support real-time analytics.

## Oracle RAC

Oracle RAC provides scale-out server clustering and near-real-time failover capabilities for Oracle Database. RAC adds complexity to an Oracle environment, and even with deep Oracle discounts, can increase the basic licensing cost by nearly 50 percent. In addition, RAC clustering adds overhead which requires additional CPU cores (and licenses) to process, driving up license and support fees.

To get an idea of RAC costs: At a 50 percent discount, the cost of RAC on a fully equipped Exadata X5-2 is \$1.6 million USD, plus 22 percent (\$350,000 USD) in annual support. In addition, for many workloads, customers will need to pay for additional processor licenses to support the overhead of RAC clustering.

Oracle built its business case for RAC based on scale-out capabilities and total cost of ownership (TCO). In the past, when large x86 servers were not available and the only scale-up option was expensive UNIX hardware, that business case made sense. The high cost of RAC was offset by the low cost of smaller x86 scale-out servers, providing a cost advantage over UNIX options. Using RAC also answered the RAS issues of the Linux OS. Today, with the availability of high-capacity and low-cost scale-up x86 servers and the maturity of the Linux ecosystem, the dynamics of the business case have changed. Now, RAC is a solution for a problem that no longer exists.

## Exadata

Exadata is a scale-out compute and storage platform designed and sold by Oracle specifically to run Oracle databases. There are some advantages to this type of synergy, but there are also some distinct disadvantages, including:

- The cost of Exadata software and support, added to the Exadata hardware purchase, is enormous (including RAC, which is required for deployments larger than two sockets to support Exadata's scale-out architecture).
- Exadata uses an inflexible architecture with preset building blocks. This can result in excess core count and licensing costs as well as requiring expensive floor space. Although the most current Exadata models offer more configuration flexibility than earlier models, other platforms supply even greater flexibility, including filling a partial rack to add capacity for other workloads.
- This strategy does not leverage customers' investment in existing storage infrastructure as it requires you to use Oracle Exadata Storage Servers. In addition to the cost of storage hardware, Exadata Storage Server Software must be licensed on a very costly per-storage-device basis.
- Managing multiple Oracle RAC nodes adds complexity.
- There are memory limitations that inhibit implementation of Oracle Database in-memory and real-time analytics.



A full rack of Exadata storage can add \$2.24 million USD in license fees, and nearly \$500,000 USD in annual support.



## Oracle licensing

In most enterprise environments, Oracle Database Enterprise Edition and Oracle RAC are licensed on a per-processor-core basis, with different processor cores licensed on a weighted basis (called a Core Factor). For example, an IBM POWER6 or later processor has a Core Factor of 1.0, while an Intel® Xeon® processor has a Core Factor of 0.5. The licensing formula is:

**Number of processors \* Core Factor \* price per core = license costs**

What this means in practice is that the license cost for an Intel Xeon processor core is half the price of the same license for an IBM POWER6, 7, or 8 core. This is important to understand when looking at growing or modernizing an Oracle environment. The choice of processor directly and significantly affects license and maintenance costs. Annual maintenance costs are based on a percentage of the licensing fee, generally 22 percent.

**Scenario A:** An Oracle Database requiring 20 processors on an IBM POWER8 platform (Core Factor = 1.0) and a heavily discounted price per core of \$10,000 USD would cost  $20 * 1.0 * 10,000 = \$200,000$  USD. Annual maintenance at 22 percent = \$44,000 USD for a five-year TCO of \$420,000 USD.

**Scenario B:** An Oracle Database requiring 20 processors on an Intel Xeon platform (Core Factor = 0.5) would cost  $20 * 0.5 * 10,000 = \$100,000$  USD with annual maintenance of \$22,000 and a five-year TCO of \$210,000 USD.

## Improving return on Oracle investments

It should be clear from the previous discussion that there are at least four ways to reduce Oracle licensing costs:

- Reduce the number of processing cores
- Migrate to processors with lower core-licensing factors like the Intel Xeon
- Review the need for and use of costly RAC licenses
- Replace Exadata deployments to eliminate storage software licensing costs

For IT leaders needing to improve Oracle TCO and return on investment, infrastructure modernization that eliminates the need for RAC wherever possible can be an attractive option.

## Modernize Oracle Database infrastructure

Many companies have aging Oracle Database architecture that needs modernization. Whether running on UNIX or Linux x86 scale-out solutions, modernizing infrastructure to industry standard x86 scale-up architecture with the latest processors and a fast-failover solution can provide several key business benefits including a much lower TCO.

New servers with more efficient processors can provide cost reductions in several areas:

- Using fewer processors to handle existing Oracle workloads equals lower license costs
- Switching to industry-standard Intel Xeon processors lowers the software license cost per processor
- Eliminating scale-out clusters can eliminate the need for RAC licenses
- Replacing proprietary infrastructure reduces maintenance and support costs
- Eliminating the hidden costs of learning proprietary solutions
- Reducing footprint resulting in lower network, power, and cooling costs

The savings in Oracle licensing and maintenance alone can be more than enough to cover hardware acquisition costs.

New scale-up infrastructure can support consolidation of medium and large-sized databases and servers to simplify management and tuning for optimum performance. Scale-up architectures can provide better performance and scale to run or virtualize complex workloads. They also have the flexibility to mix OLAP and OLTP workloads and different IT landscapes on the same frame. Fast-failover solutions can provide reliability and availability equal to or better than enterprise mission-critical UNIX servers without the cost and complexity of RAC.

IDC estimates that moving from a scale-out to scale-up architecture can reduce solution costs by 35 percent or more.<sup>1</sup>

## Eliminate Oracle RAC

**Moving to scale-up architecture** eliminates the need for RAC clustering. The other critical capability that RAC provides is near-real-time failover. Scale-up x86 solutions like HPE Integrity Superdome X can provide five-nines availability (99.999%), and when paired with HPE Serviceguard, can provide immediate automated failover and disaster recovery functionality, and eliminate planned downtime. Which option makes sense for your business?

If you measure the cost of downtime in hundreds of thousands or millions of dollars per minute, then you may need the near-real-time failover that Oracle RAC provides. If availability is a matter of life and death (as in some medical applications), then you may look to RAC.

It's a question of balancing risk aversion and the cost of downtime. Is it worth hundreds of thousands or millions of dollars in licensing, maintenance, and operations costs to ensure an extra three to five minutes of availability each year (the difference between 100 percent and > 99.999 percent availability)? For some enterprises, the answer is "yes," but for many enterprises, RAC is an expensive luxury that is no longer needed.

What does five-nines reliability really mean? Technically, it means that the maximum downtime for a system will be no more than 5.36 minutes per year. In practice, it means that the server hardware is so reliable that the likelihood of a hardware failure is next to zero.

## Oracle VM Server

Oracle documentation seems to imply that customers planning to run Oracle Database in a virtualized x86 environment must license every core within the physical server if any core runs Oracle.<sup>2</sup> Obviously, this can be quite expensive. A 44-core server running Oracle on a small virtual instance would require 22 licenses by this policy. HPE Integrity Superdome X nPars can help address the uncertainty of managing licenses in a virtualized environment by creating hard partitions as small as eight cores (or four Oracle licenses).

<sup>1</sup> HPE's Superdome X: "The Mission-Critical Scale-Up Platform for SAP®, Oracle, and SQL Server," IDC, June 2016

<sup>2</sup> Oracle customers can refer to their individual contracts for clarifying language regarding virtualization and where Oracle must be licensed. Oracle's Partitioning Policy document does note that it is for educational purposes only and does not constitute a contract.

As an alternative to VMware®, Oracle has introduced Oracle VM Server (commonly referred to as OVM). However, there are some considerations in adopting a new virtualization technology:

1. Many businesses have standardized on VMware, and it is frequently key to their cloud strategy. Introducing a second virtualization technology is not helpful and adds complexity.
2. Gartner's Magic Quadrant for x86 Server Virtualization Infrastructure ranked OVM in the lowest quadrant, while VMware ranked as a leader in vision and ability to execute in virtualization technology.<sup>3</sup> Given the importance of virtualization for cloud computing, compromising in this area represents an unwarranted risk.
3. Oracle will point to things like the OVM templates that are available to customers deploying Oracle applications. The real value of OVM is that it is free and only requires support costs. From an enterprise perspective, if reduced functionality for your virtualization engine is worth the reduced acquisition costs, then it may be right for you.

## A case study

A major manufacturer with a large catalog of configurable and custom-built products runs one of the world's largest Oracle databases to track the massive volume of stock-keeping units (SKUs) it manages. The organization needed to improve the performance and scalability of its enterprise resource planning (ERP) solution along with other applications supported by multiple Oracle databases. They also needed to ensure redundancy and disaster recovery for all of their mission-critical databases while reducing licensing costs.

The company chose to migrate from UNIX to Linux, using a scale-up server solution that allowed them to deploy without Oracle RAC. They ported 26 Oracle databases to the new environment and achieved several important objectives:

- Improved performance and scalability
- Avoided incremental Oracle licensing costs (including Oracle RAC)
- Met high-availability and fast-failover requirements for mission-critical databases
- Reduced annual operating costs by over \$200,000 USD

## HPE Server Portfolio Solutions

Hewlett Packard Enterprise provides a full portfolio of right-sized server solutions allowing IT organizations to match processing power and scale with current and future needs, from small to large enterprise deployments, at price points that fit within almost any IT budget.

- **HPE Integrity Superdome X** servers provide the largest capacity, fastest processing, and maximum availability across Linux, Windows®, and VMware to support any size deployment.
- **HPE Integrity MC900 X** servers provide a highly scalable, high-performance, Linux-only alternative for price-sensitive organizations that may not need all of the capabilities available in the Superdome X.
- **HPE ProLiant DL580** servers provide a Linux, Windows, and VMware alternative for mid-size databases or small and medium enterprises.

From mid-size to large global enterprises, **HPE Integrity Superdome X scale-up servers** provide capabilities uniquely suited for Oracle Database deployments of any scale.

- **Scalability**—Scale-up capacity supporting two to 16 sockets, up to 384 CPU cores, and up to 24 TB of main memory on a single node
- **Speed**—World's fastest x86 scale-up platform, with **15 published world-record benchmarks**
- **Reliability**—UNIX-like, end-to-end RAS features with five-nines availability and up to twenty-times more reliability than other x86 servers; even greater business continuity when paired with HPE Serviceguard<sup>4</sup>

<sup>3</sup> Gartner Magic Quadrant for x86 Server Virtualization Infrastructure, July 2015

<sup>4</sup> HPE nPars hard partitions have ~5 percent the number of Single Point Of Failure (SPOF) of software-only partitions. Based on HA and field data modeling, HP Labs (Now Hewlett Packard Labs), August 2013



Where near-real-time failover is a requirement, HPE Integrity Superdome X supports Oracle RAC. However, Superdome X simplifies RAC deployment because HPE Integrity Superdome X can support larger node sizes. Unique to the industry, Superdome X offers nPar hard partitioning that creates separate compute blocks that can be configured to specific workloads to reduce core count and limit Oracle license costs.

**The HPE Integrity MC990 X server** provides a highly scalable Oracle database solution. While it does not include all of the features available in Superdome X (RAS, nPars), it does provide a high-performance 8-socket alternative at a more budget-friendly price point and is certified for Oracle Linux and OVM.

- **Scalability**—Scale-up capacity supporting up to 192 CPU cores
- **Performance**—Up to 12 TB of main memory on a single node can provide a large performance boost when paired with Oracle's in-memory database technology
- **Reliability**—Fast recovery and enhanced downtime protection when paired with HPE Serviceguard

Where near-real-time failover is a requirement, the HPE Integrity MC990 X servers support Oracle RAC, with RAC deployment simplified because the HPE Integrity MC990 X server can support larger node sizes.

For small to midsize enterprises with light to medium workloads, **HPE ProLiant DL580 Gen9 scale-up servers** provide a cost effective solution for your Oracle database deployments and are certified for Oracle Linux and OVM.

- **Scalability**—Scale-up capacity supporting two to four processors, 96 cores, up to 6 TB of memory along with redundant power and cooling, storage, and network connectivity
- **Speed**—Winning big on eight world-record benchmarks<sup>5</sup> by delivering up to 30 percent<sup>6</sup> more performance
- **Reliability and availability**—Increased system availability and less need for service with advanced error recovery, error diagnosis, and built-in redundancy; all features of the HPE ProLiant DL580 Server Comprehensive Fault Management and Diagnostics and even greater business continuity when paired with HPE Serviceguard for Linux.<sup>7</sup>

**Note:** For high density or power constrained environments the HPE ProLiant DL560 and DL380 Gen9 Servers can be suitable compute alternatives.

Whether you are consolidating your database workloads from scale-out x86 servers, upgrading your aged x86 servers to meet performance and RAS expectations, or migrating from costly proprietary servers, HPE ProLiant DL580 Gen9 Scale-up Servers provide the foundation to enhance your data-driven infrastructure with greater performance, scalability, and reliability.

<sup>5</sup> The stated results are published as of 06-06-16; see [spec.org](http://spec.org). TPC-H results show the HPE ProLiant DL580 Gen9 with a result of 2,140,307.2 QphH @ 3000 GB and \$0.38 USD/QphH @ 3000 GB with system availability as of July 31, 2016; see [tpc.org/3323](http://tpc.org/3323). TPC-H results show the Cisco UCS C460 M4 with a result of 1,071,018 QphH @ 3000 GB and \$0.60 USD/QphH @ 3000 GB with system availability as of June 1, 2016; see [tpc.org/3322](http://tpc.org/3322). The TPC believes that comparisons of TPC-H results published with different scale factors are misleading and discourages such comparisons. Results as of June 6, 2016; see [tpc.org](http://tpc.org) for more information.

<sup>6</sup> Intel® measurements. Up to 1.3X average performance across key industry benchmarks (SPECjbb\*2015 Multi-JVM Critical and Max jOPS, SPECint\*\_rate\_base2006, SAP SD\* 2-tier, SPECvirt\_sc\*2013, and TPC-E\*) comparing best 4-socket OEM server publications on [spec.org](http://spec.org), [sap.com/benchmarks](http://sap.com/benchmarks) and [tpc.org](http://tpc.org). See [intel.com/content/www/us/en/benchmarks/server/xeon-e7-v4/xeon-e7-v4-world-record.html](http://intel.com/content/www/us/en/benchmarks/server/xeon-e7-v4/xeon-e7-v4-world-record.html) for more information.

<sup>7</sup> Recovery observed in HPE internal lab testing. System was based on HPE ProLiant DL380 Gen9 server with Red Hat® Enterprise Linux 6.1 running HPE Serviceguard 12.00.00. The failover time excludes any application-level operations.



## Conclusion

IT leaders looking to reduce the TCO for their Oracle Database deployment should consider modernizing Oracle infrastructure to a scale-up x86 platform. This can:

- Reduce the cost of processor licenses
- Eliminate the need for Oracle RAC, removing complexity and further reducing license costs
- Preserve or enhance existing levels of service
- Provide flexibility for deployment of in-memory database with Oracle 12c

HPE Integrity Superdome X, Integrity MC990 X, and ProLiant DL500 series servers provide a reliable, scalable, high-performance alternative to costly UNIX or x86 scale-out solutions, with features and sizing to fit nearly any Oracle database workload requirements or IT budget.

Hewlett Packard Enterprise offers an Oracle Database assessment that identifies ways to improve performance and reduce the high cost of Oracle licensing and support. In addition, HPE provides proven expertise in end-to-end planning, integration, and service of new and refreshed Oracle environments to achieve the best TCO for every customer. Check out how Hewlett Packard Enterprise can get you started on your journey to **optimizing Oracle costs by modernizing critical Oracle infrastructure**.

Learn more at

**[hpe.com/servers/superdomex](http://hpe.com/servers/superdomex)**

**[hpe.com/servers/mc990x](http://hpe.com/servers/mc990x)**

**[hpe.com/servers/dl580gen9](http://hpe.com/servers/dl580gen9)**

**[hpe.com/info/dcm](http://hpe.com/info/dcm)**



Sign up for updates