



The Compelling Economic Value of Incomparable Database Performance

by, Marc Staimer

January 2022

Research Premise

When Oracle announced the latest X9M generation for Exadata Cloud@Customer, Exadata Database Machine and Zero Data Loss Recovery Appliance (Recovery Appliance), it radically pushed the database performance envelope to mindboggling heights. Oracle performance specifications are remarkable at $\leq 19 \mu\text{s}$ average read latency with more than 27.6 million Oracle Database 8K SQL read IOPS and 7.4 million write IOPS, and over 1 TB/s throughput per Exadata X9M full rack. Note that this is not the maximum performance per Exadata X9M system. Far from it. Every Exadata X9M on-premises can scale from what Oracle calls a base system of 2 database servers and 3 storage servers to a quarter rack, half rack, full rack, and up to 12 racks with linear scalability. In other words, adding more performance is as simple as expanding within a rack or adding more Exadata X9M racks to the same system. This kind of database performance would have been considered fantasy just a few short years ago.

Oracle did not limit these X9M extraordinary database performance gains to on-premises systems. They made them available via Oracle Cloud@Customer offerings that bring Oracle Cloud Infrastructure (OCI) to the customer premises¹—namely in the form of Exadata Cloud@Customer and Dedicated Region Cloud@Customer deployment options.

This currently unmatched database performance is available now. That means Oracle believes there is a market need or customer requirement for this exceptional database performance. Oracle skeptics and competitors on-the-other-hand, will ask what are the use cases requiring these very high levels of performance? They're likely to point to an automobile analogy such as the very high-performance Tesla

¹ Oracle Exadata X9M Cloud@Customer is a fully managed cloud service with a complete cloud management control plane. Performance comes in at a robust 22.4 million Oracle Database 8K SQL read IOPS and 540 GB/s throughput per Exadata X9M rack. The OCI control plane consumes some of the rack space decreasing the maximum Exadata X9M rack performance slightly. Oracle Dedicated Region Cloud@Customer is an entire OCI data center custom built on the customer premises with all of the OCI cloud services available.

Model S Plaid. It can go zero to 60 miles per hour (100 kilometers per hour) in less than 2 seconds while achieving a top speed of approximately 200 Mph or 322 Km/h. There are very few automobile situations that require those levels of acceleration or that top-end speed. This analogy falls completely apart when it comes to databases and the applications that use them. It assumes that a database is single purpose, i.e., only works with a single application. It also assumes the value of this incredible performance is nominal, solving the needs of a narrow few use cases. Research shows otherwise based on the many problems it solves.

This Wikibon research dives deep into high-performance database use cases and more specifically:

- What user problems does extreme database performance solve?
- How does the X9M platform for Exadata Cloud@Customer, Exadata Database Machine and Recovery Appliance solve those problems?
- How does it compare to alternative on-premises and public cloud solutions?
- Why do IT organizations that don't need that level of performance for any single application still need Exadata X9M?

In the Wikibon view, as Oracle continues to expand its impressive database performance advantage, it will continue to win new customers and expand its share of the large, medium, and small Enterprise database market.

Key Findings

This Wikibon research is designed to help CIOs understand how Oracle's Exadata X9M platform cost-effectively solves many of the multi-database, multi-workload performance and manageability problems and requirements. It examines how those problems and requirements are solved today and why legacy workarounds are neither sustainable nor cost-effective, especially with modern applications such as AI and machine learning.

Exadata is co-engineered with the Oracle Database. That means the Oracle Database provides features, functions, benefits, and performance only available on Exadata. It doesn't matter what other server or storage platforms offer, these cannot come close to the Oracle Database performance or functionality running on the latest Exadata. The cost per IOPS and/or throughput simply blows away every competitive system. This was true with Exadata X8M. Now Exadata X9M increased IOPS by up to 72% and throughput by up to 87%.

And yet, Oracle did not raise the price at all. Same cost as Exadata X8M with much greater performance. That reduces the IOPS cost by approximately 42% and the analytical scan (throughput) costs by as much as 47%. And it does this while increasing pooled resources by 33%. All with zero cost increase! That's pretty remarkable.

The Exadata X9M is so much more than a simple processor upgrade. It's a complete system upgrade including up to:

- 33% more CPU cores.
- 33% more DRAM.
- PCIe 4.0 with 2x bandwidth of PCIe 3.0.
- 28% more usable HDD capacity while keeping fast Flash SSD capacities unchanged.

Similar problem-solving cost/performance advantages are maintained with the Cloud@Customer offerings². Consider that other on-premises cloud systems such as Amazon Outposts have two cost components. The first is the monthly hardware rental component requiring a minimum annual commitment. The second is based on compute time utilized and the number of vCPUs³. Oracle Cloud@Customer services have similar components. Pricing is tied to the number of vCPUs per second when completing a database transaction, query, task, etc. Because Oracle Database and Oracle Autonomous Database running on Exadata are substantially faster than other public cloud database cloud services, the cost per query or job typically comes in well below **half** that of competitive database cloud services. Oracle's much lower cloud database cost is one of the reasons why many customers claim they

² Cloud@Customer CPU increase at 23% instead of 33% whereas DRAM stays the same. PCIe is Gen 4.0 and HDD capacity increases by 28%.

³ A vCPU is a ½ core.

switched to OCI or Exadata Cloud@Customer. The other reasons customers quote are performance, the way it solves significant and often intractable database problems, decreases costs, and increases revenues. This is explored more deeply in this research.

The Oracle X9M platform's unparalleled performance leads directly to dramatic database cost reduction through extensive consolidation, appreciable and measurable productivity gains, faster-time-to-market, faster-time-to-actionable-insights, and more direct revenues. Unmatched automation frees up DBAs and backup administrators for more strategic work. No other database and platform combination, or cloud database service comes within sight of this latest Exadata X9M engineered system from Oracle.

Table of Contents

Research Premise 1

Key Findings 2

Table of Contents 4

Problems the Oracle Exadata X9M Platform is Architected to Solve..... 5

 Out-of-control database-sprawl 5

 Conspicuous decline in database performance as database data proliferates..... 6

 Mediocre to poor user productivity 6

 Excessive time-consuming database management 6

 Regulations, rules, and laws inhibiting many organizations from on-premises cloud services..... 7

 Sub-optimal database availability from both planned and unplanned disruptions/outages 7

 Prohibitive database costs 8

How Oracle Solves These Problems 8

 Exadata X9M baseline performance 8

 Solving out-of-control database sprawl 9

 Solving the conspicuous decline in database performance as database data proliferates – i.e. scaling
 10

 Solving mediocre-to-poor user productivity 11

 Exadata Cloud@Customer X9M solves specific on-premises problems no one else solves 11

 Solving extensive time-consuming database management 11

 Solving regulations, rules, and laws preventing too many organizations from on-premises cloud
 services 12

 Recovery Appliance X9M also solves sub-optimal database availability from planned and unplanned
 disruptions-outages 12

 Improving database availability 13

 Solving Prohibitive Database Costs 14

**How Oracle’s X9M platforms—Exadata Cloud@Customer, Exadata and Recovery Appliance—
Match up against Major Competitors15**

 Exadata X9M Database Machine and the Competition 15

 Exadata X9M Database Machine vs HCI-Nutanix 15

 Exadata X9M vs External Storage: Pure Storage FlashArray//X90R2 16

 Exadata X9M Database Machine vs External Storage: Dell EMC PowerMax 16

 Exadata X9M Database Machine vs IBM Power Systems E1080 17

 Exadata Cloud@Customer X9M and the Competition 17

 Exadata Cloud@Customer X9M vs AWS Outposts 17

 Recovery Appliance X9M and Competition 18

Conclusion.....19

More Information.....20

Problems the Oracle Exadata X9M Platform is Architected to Solve

Databases of all kinds are essential to IT organizations. Commercial transactions, analytics, big data, graphics, document management, blockchain, AI and machine learning all depend on databases. Databases are ubiquitous, mission-critical, and incredibly valuable to everyone.

Notwithstanding the importance and omnipresence of databases, there are several intractable and pervasive problems. Problems such as:

- Out-of-control database and infrastructure sprawl
 - Data silos across centralized IT and business units
 - Servers, NICs, switch ports, storage, cables, cable management, transceivers, conduit, rack space, floor space, and more.
- Conspicuous decline in database performance as data proliferates
 - Exacerbating database sprawl
- Mediocre to poor user productivity
 - Caused by slow and variable response times
- Excessive time-consuming database management
 - Limited automation
 - Demanding much of the DBA's time instead of focusing on more strategic pursuits
 - Such as SQL query optimization or application development
- Regulations, rules, and laws preventing too many organizations from moving to public cloud services
- Sub-optimal database availability from planned and unplanned disruptions/outages
 - Especially patching and upgrading
 - Further ruining user and customer productivity
- Prohibitive database costs
 - License, subscription, and/or support and maintenance fees

Out-of-control database-sprawls

Database sprawl is the out-of-control proliferation of databases. They can be the same type of database and/or different types. Database sprawl across central IT and business units leads to data silos which make it difficult to get an overall view of the business and generate insights, impeding information exchange between business units. Each one of these database instances creates additional work for the IT organization. Every database instance must run somewhere. It can be in a container, a virtual machine, or a bare metal system, on an on-premises platform, or in a public cloud. Databases don't live in a vacuum. They need infrastructure support including networking, storage, conduit, transceivers, cables, rack space, power, cooling, plus people support including DBAs, system administrators, data architects, data scientists, and more.

A key contributing factor to database sprawl is the rapid spread of specialized databases and public database cloud services. Many databases deal with a single type of database structure such as transaction management (relational), transactional analysis (data warehouse), documents (JSON), object (XML), financial transaction analysis (time series), interconnects between events or people (graph and spatial), deeply secure transactions (blockchain), AI analytics (machine learning or ML), IoT and more.

Database sprawl becomes more complicated and seriously problematic when the many proliferated and vastly different databases need to communicate with one another. Whether it's ETLs (extract transform and load) between different databases or joins between similar databases, the processes are complex. Increased complexity invites risk as "Mr. Murphy" – as in Murphy's Law - rears his ugly head. As the number of database connections escalates so does the likelihood of minor and major problems and errors. That becomes a serious issue of time. Human beings have a hard limit on the amount of work they can produce in a unit of time. As problems and errors swell so does the time and effort required to troubleshoot and perform root cause analysis. That doesn't include the massive time sink in managing each of these separate database instances. Patching alone becomes a nightmare. Support becomes unsustainable.

Another significant cause of database sprawl is the preponderance of various different developer tools and programs. Several of the database cloud services, proprietary databases, and open-source databases

enable developers to develop at no cost. They make it incredibly simple and easy to spin up new and unique or specialized databases to handle different data types. But once in production, those specialty databases are no longer “free.” Even the open-source databases have support costs. Make no mistake, few DBAs are willing to put mission-critical or simply important database applications into production without support.

Every single database adds cost. Some more than others. And the costs are not linear. Communication between databases adds more cost. Infrastructure growth is also non-linear, increasing more rapidly as incremental demands force extensive investment on-premises or in public clouds. Database sprawl devolves into an IT mess of technical complexity and exponentially rising costs. More on management and costs later.

Conspicuous decline in database performance as database data proliferates

Database performance generally declines as the amount of data in the database grows. Reasons that performance drops include insufficient compute resources available for the load; limited memory resources forcing frequent page swaps; poor storage latencies, IOPS, and throughput; and inadequate bandwidth and performance in the storage network. As an example, steep decreases in storage system performance as capacity utilization grows above 50% have been observed—and it drops off a cliff at 80% or more of capacity utilization for most storage systems.

Other data growth factors negatively impacting database performance include database backups taking too many server resources for too long, complicated joins and upgrades with application regression testing consuming too many resources and time.

The most common workaround for poor database performance is to add more database instances and limit the number of applications utilizing a given database. Another workaround is to shard the database (break it into parts) with one primary database and many subordinates. Sharding is incredibly complex to set up and maintain. While these workarounds may improve performance, the by-product is generally greater database-sprawl. With one exception, the [Oracle Database](#).

Mediocre to poor user productivity

As database performance declines so do the applications’ response times. Multiple studies demonstrate a direct correlation between application response time and user productivity. IBM’s groundbreaking 1982 Redbook on “[The Economic Value of Rapid Response Time](#),” demonstrated that as response time decreased productivity increased. In contrast, response times greater than 3 seconds resulted in plunging productivity and sagging revenues and profits. Application response times of 3s was the threshold in which productivity falls off a cliff. As application response times fall below 3s, productivity increases. When application response times decreased to 300ms (.3s), productivity soared to more than 2x that at 3s. One unexpected side effect was the increased quality of work. Another key finding from this research was user addictiveness to sub-second application response times. In lay persons’ terms, they love it.

Translating this research to applications such as AI neural networks and machine learning shows similar productivity results. Faster application response times are highly impactful for analytical applications, such as machine learning, that derive insight from large amounts of ingested (read) data, i.e., throughput. Faster insights translate into faster-time-to-action that in turn leads to increased revenues and/or decreased costs.

When database performance declines, increasing application response times, productivity tumbles, time-to-actionable-insight slows, revenues decline, and market share is lost. The workarounds most commonly deployed include upgrading to faster servers, storage, and networks and/or the implementation of more database instances on VMs, containers, or bare metal machines. I.e., more database-sprawl.

Excessive time-consuming database management

Managing any database requires skill, knowledge, and experience. Most databases generally lack much in the way of database automation. This is true both for cloud database services and on-premises offerings. Database performance tuning is a common DBA task that’s highly manual and labor-intensive. It can take a significant chunk of time for the DBA to manage this single and quite important process. And that’s just one database process requiring management.

Each additional database instance adds more than 100% additional effort. Adjusting the network and storage resources for one database instance may negatively impact the performance of another database instance. That requires the DBA to consider how their tuning efforts of one database is affecting others and adjusting appropriately. That takes more time, trial and error, and effort. Technology advancements in high-performance computing, networking and storage have somewhat simplified database performance tuning. But not as much as lay people think. The DBA is still instrumental when it comes to tuning database performance.

Consider patching, especially security or vulnerability patching. Not just for the database itself but for everything supporting each database—OS, server, hypervisor, network, storage, microcode on everything, etc. For those patches requiring a reboot or disruption, they have to be scheduled. The schedule for all application-disruptive processes are typically one weekend per quarter. The first 24 hours are for implementing the disruptive processes. The second 24 hours are to back them out should they cause problems. That process is workable for a handful of database instances. But for hundreds or thousands of instances, it is far too labor-intensive to be workable or sustainable.

The workarounds rapidly gaining popularity are cloud database services, sometimes referred to as database-as-a-service. The other is migrating the organization's database licenses and data to a public cloud. These are different. A cloud database service provider licenses, manages, and supports the database with the customer paying for the service and supporting cloud infrastructure based on the resources (computing, storage, etc.) used. That means the provider handles the patching and upgrading. It doesn't mean it's non-disruptive. Customers still have to plan appropriately.

The migration of the organization's database licenses and data to a public cloud, a.k.a. bring-your-own-license (BYOL), can reduce some costs. However, some of the benefits and value of a cloud database service is often lost. Benefits such as free no touch upgrades. Infrastructure granularity tends to be much coarser as well. Take the case where one more virtual CPU (vCPU) is required. It may require adding many more of them and memory because that's the next available increment. That's potentially a major price jump for a small demand. This makes it imperative for every IT organization planning on moving their database licenses to a public cloud, completely understand that cloud's infrastructure granularity for their database. When that granularity is too coarse and costly, they should look at a public cloud that can more closely meet their needs.

Both the cloud database service and the BYOL issues vary by service provider. However, neither reduces database sprawl. They in fact may make it worse. Take Amazon Web Services (AWS), as an example. The maximum IOPS AWS allows per database instance today is 260,000. When a database needs more there are no other AWS options. There are greater third-party storage options, but then there's the EC2 instance performance limitations that top out at 64,000 IOPS and latencies at around 1 millisecond (1m). That means breaking up the database and... more database sprawl. And to make matters just a little bit worse, this is limited to no automation, and DBAs will still have to perform ETLs between different database types.

[Regulations, rules, and laws inhibiting many organizations from on-premises cloud services](#)

Many regulations require tight control and auditing for all system and database management access regardless of where they're located. On-premises cloud deployments require vendor access to manage infrastructure. The problem is that the vast majority of on-premises cloud offerings, such as Amazon Outposts, simply do not allow the customer to fully control access with the cloud provider's staff.

Complicated workarounds limit and frustrate both customer and public cloud service providers. It either puts the customer in a non-compliance position with regards to the regulations, rules, and laws, or it makes it too difficult for the public cloud service provider to effectively manage their service SLAs. It's an exasperating situation.

[Sub-optimal database availability from both planned and unplanned disruptions/outages](#)

Database availability is a paramount requirement for all databases, but especially mission-critical ones. For many databases of all types, there are planned and unplanned outages or performance slowdowns that reduce availability. Something as simple as applying vulnerability patches or hot fixes may require a reboot of the system.

Backups consume database and database server resources and generally require the database to quiesce briefly, causing a relatively short outage. Database upgrades on the other hand are generally considerably worse due to application regression testing—validating that database-reliant applications can run without issues—that often consumes significant resources and time. Without application regression testing, those same upgrades can cause unplanned application outages.

One popular workaround to reduce or limit the impact of backups, recoveries, patches, hot fixes, and upgrades is to increase the number of database instances while constraining the dependent applications on any given database. Also reducing the total amount of data per database reduces backup time, recovery time, and impact of any given outage.

Database recovery time varies by the outage's severity, data amount, and technology used to backup and recover the database. Database data growth causes both backups and recoveries to take longer. Once again, the most common workaround is to increase the number of database instances. Database-sprawl is the obvious go-to answer for too many DBAs, until it becomes unmanageable.

Prohibitive database costs

Database sprawl is insidious. It sneaks up on organizations. It's analogous to the boiling frog fable⁴. Database sprawl is similar. At first it seems manageable and appears to solve performance and backup/recovery problems. But over a relatively short period of time, the management and especially the costs, spiral out of control. At which point DBAs feel trapped by their own processes.

The costs become prohibitive whether it's on-premises capital and operating expenditures or cloud operating ones. Trepid costs at first, boiling costs before too long.

On-premises or traditional costs include the physical servers, VMs, containers, switches, storage, cables, transceivers, conduit, OS licenses or subscriptions, database licenses or subscriptions, ETL tool licenses or subscriptions, rack space (allocated overhead), power, cooling, UPS, database protection, training, maintenance, spares, support, management costs and more. Every database instance adds to these costs. Of these costs, management and support grow more quickly than any other as database ecosystem complexity increases.

There is no easy workaround for prohibitive database sprawl costs, regardless of where it's located. Those costs will continue to spiral until they become unsustainable.

How Oracle Solves These Problems

Each of the Oracle Exadata, Exadata Cloud@Customer, and Recovery Appliance X9M products and services play a unique role in solving all of the aforementioned problems. They all leverage the incredible X9M performance to do so. A deeper dive reveals more.

Exadata X9M baseline performance

Starting with previously discussed extraordinary performance⁵ per **Exadata X9M Database Machine** rack.

- Average read latency at $\leq 19 \mu\text{s}$.
- ≥ 27.6 million Oracle Database 8K SQL read IOPS.
- ≥ 7.4 million write IOPS.
- 1 TB/s throughput.

While this performance is per **Exadata X9M Database Machine** rack, **Exadata X9M Database Machine** is not limited to one rack—it can scale out to 12 racks. And the performance gains are linear. Four Exadata X9M racks delivers 4x the performance of a single rack. Every **Exadata X9M Database Machine** cluster can be as a single database system or logically partitioned into multiple discrete databases. Oracle production-hardened Real Application Clusters (RAC) greatly simplifies scaling **Exadata X9M Database Machines**

⁴ This is the story of a frog is put in tepid water while gradually raising the temperature until the frog is boiled alive. If the frog were put directly into boiling waters first it would leap out immediately. But by increasing the temperature gradually, it does not become aware of the danger until it's too late.

⁵ Performance results are based on 8K SQL IOPS, not 4K Iometer unrelated IOPS used by most storage vendors, in an elastic configuration with 10x Exadata X9M-2 Database Servers and 12x Exadata X9M-2 Extreme Flash Storage Servers in a single rack.

enabling dynamic processing power additions. Storage additions can also be added dynamically with Oracle Automatic Storage Management (ASM).

For those requiring less performance, the **Exadata X9M Database Machine** can start out at a much smaller configuration. The entry-level system is a 1/8th rack with up to 2.8 million read IOPS. It can be upgraded to a quarter, half, or full rack thereafter, plus up to 11 more racks.

Exadata Cloud@Customer X9M provides the exact same per server performance as **Exadata X9M Database Machine**, but packs fewer servers into a single full rack by virtue of its OCI control plane. **Exadata Cloud@Customer's** full rack performance is as follows:

- Average read latency still at $\leq 19 \mu\text{s}$.
- ≥ 22.4 million Oracle Database 8K SQL read IOPS.
- ≥ 5.9 million write IOPS.
- 540 GB/s throughput.
- Currently limited to one rack per Exadata system.

However, only **Exadata Cloud@Customer X9M** and **Dedicated Region Cloud@Customer** support Oracle Autonomous Database on-premises, which more than makes up for the performance delta.

This unprecedented performance solves the following previously mentioned problems

- Out-of-control database sprawl.
- Conspicuous decline in database performance as database data proliferates
- Mediocre-to-poor user productivity

Peeling back the curtain a bit explains how.

[Solving out-of-control database sprawl](#)

The **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** are uniquely co-engineered with the Oracle Database, Intel CPUs, Linux OS, Intel® Optane™ Persistent Memory (PMem), RoCE switches, and RDMA storage networking. It can run all supported Oracle Database versions. **Exadata Cloud@Customer X9M** goes a step further in that it's the only platform on-premises that also supports Oracle Autonomous Database – for transaction processing and mixed workloads (ATP) or Autonomous Database for analytics and data warehousing (ADW).

Competitors frequently state that Oracle's co-engineering claim is marketing fluff. It's not. The databases and hardware are specifically optimized for one another down to the microcode, enabling performance levels simply not obtainable any other way. In fact, there are more than 60 net incremental database capabilities only available on Exadata.

Exadata X9M solves out-of-control database sprawl via consolidation utilizing Oracle Database multi-tenant container databases (CDBs) and pluggable databases (PDBs)⁶. For more detail about CDBs and PDBs—go to the [Oracle's Administrator's Guide](#).

The key here is the marriage of Oracle Database multi-tenant CDBs and PDBs with **Exadata X9M** makes a serious dent in the out-of-control database sprawl problem facing companies of all sizes today. Users can plug multiple PDBs into a single CDB or they can run multiple CDBs and only utilize PDBs with the same service level agreements (SLA) in a given CDB. Most IT organizations really don't want to run their mission-critical applications and databases alongside non-mission-critical or low-level applications.

CDBs and PDBs additionally leverage the converged database capabilities of the Oracle Database. Too many IT pros still see the Oracle Database as only an Enterprise-grade transactional relational database (RDBMS). That is a grossly out-of-date perception. The Oracle Database today is considerably more than that. It's a highly converged database. In addition to its traditional transactional RDBMS, it is also a data warehouse, object/XML, document/JSON, time series, spatial, graph, block-chain, and machine-learning database. All

⁶ Each CDB contains one or more user-created PDBs and application containers, supporting —depending on the license—up to 3 (SE2), 252 (EE), or 4,096 (EE-ES) PDBs. Each **Exadata X9M** can support thousands of CDBs. CDB at the physical layer, is essentially a set of files including the control file, online redo log files, and data files. The database instance manages the files that make up the CDB. A PDB is the portable collection of schemas, schema objects, and nonschema objects that appears to an application as a separate database. PDB at the physical layer, is its own set of data files that store the data for the PDB. Whereas the CDB includes all the data files for all the PDBs contained within it, and a set of system data files that store metadata for the CDB itself.

of these database types come with the same Oracle Database license. Fewer licenses also reduce database sprawl and cost. Running 15 different databases only increases surface area exposures, requiring dozens of different APIs, new training and processes, different management skills and learning curves and always results in much higher costs.

CDBs and PDBs enable users to leverage all of those database types and workloads into a single system by radically optimizing hardware utilization. Less hardware further reduces database sprawl. The amount of database sprawl decrease is only as steep as the supporting hardware infrastructure can support. And that's where the **Exadata X9M's** performance and scalability come very much into play. The hardware is all integrated, optimized and tuned for the best possible performance without any effort required by the user.

What about Oracle competitors' claims of market fluff? Can't an IT organization or system integrator achieve the same levels of performance with the latest hardware and infrastructure available off-the-shelf? Actually, no. And it's not close. Adding the latest storage with NVMe, fastest Ethernet, Infiniband, or Fibre Channel networking, the fastest, most powerful servers, and more memory to a database will very likely improve any database's performance. However, it still will not come close to the performance of Oracle **Exadata X9M** with the Oracle Database. A big reason for this is that co-engineering can't exist currently in any other platform without access to database source code, non-public hardware APIs, and more.

To put that in perspective, the more than five dozen Oracle Database on Exadata capabilities **not available on any other platform**, are primarily targeting performance. Here are just some of the highlights:

- **Smart Scan** that automatically offloads data and low-level SQL queries, JSON, analytics, machine learning (ML) algorithms, and data decrypt operations to the Exadata storage servers. This delivers much faster analytics by enabling more cores for database operations. Sharing Smart Scan metadata on storage servers across parallel query processes improves scalability. NOTE: Oracle Database licensing, subscription and Cloud@Customer vCPU fees are tied only to the database server cores, not the storage ones. This enables more database processing at no additional license cost. There are no other databases that can offload database processing to storage processors. Oracle Database can only do that when running on Exadata. Smart Scan basically knows exactly where to look to find a specific set of data rapidly in Oracle Database—everything else in the market spends excess CPU cycles trying to figure out where to look first. It's like the difference between knowing where to go in your head or relying on a 1982 paper map.
- **Smart Flash Cache** and **Storage Index** that automatically accelerate database I/O.
- Automatic conversion of data to fast **In-Memory Columnar** format in flash that accelerates analytics.
- Exclusive **RDMA algorithms** for inter-node scale-out cluster coordination that accelerates OLTP I/O
- **Faster decryption** algorithms for all storage server generations, accelerating processing on storage servers.
- **Smarter Management** providing enhanced database alerting and faster software updates increasing security.
- **Enhanced algorithms** for Oracle ML, Graph, Spatial, In-Memory, and Blockchain Tables that accelerates queries and security.

Exadata Cloud@Customer X9M and Exadata Database Machine X9M deliver performance and scalability that support unparalleled massive database consolidation for multiple database types and workloads not available in any other on-premises platforms. Consolidation for dozens, hundreds, thousands, even tens of thousands Oracle databases of all kinds.

Consolidation shatters database sprawl.

[Solving the conspicuous decline in database performance as database data proliferates – i.e. scaling](#)

Exadata X9M's scalability in compute, performance, and capacity eradicates this issue. **Exadata X9M** is far more than a hyperconvergence, a reference architecture, or validated design for do-it-yourself architects consisting of servers, networking, and storage. It is carefully thought-out, precision engineered, tight integration of processing power, RAM, 100Gb/s RoCE fabric, switches, PMem, high performance Flash SSDs,

high-capacity HDDs, Linux, and the Oracle Database. The result is unprecedented scaled up performance of **Exadata Cloud@Customer X9M** or **Exadata Database Machine X9M** system. There's just nothing else that comes close in the market today. That's just the beginning. The on-premises **Exadata X9M** version also scales-out providing mind blowing performance. A scaled-out system of 12 racks configured with 8 database servers and 14 storage servers per rack, provides **up to**:

- ≥ 331.2 million Oracle Database 8K SQL read IOPS.
- ≥ 88.4 million write IOPS.
- 12 TB/s throughput.
- 6,144 cores.
- 196.61 TB RAM.
- 252 TB PMem.
- Incredible total drive capacities:
 - 45.6 PB raw disk storage.
 - 11.04 PB raw flash storage.
- Hybrid Columnar Compression (HCC) also increases the effective storage and memory capacity by a proven factor of 10x.

Because **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** can independently scale database servers and storage servers as well as storage server drives non-disruptively, the performance and/or capacity can be added as required as the database data increases. Scalability costs are greatly limited by adjusting only those parts of the infrastructure required to maintain performance as data scales. And by cleverly moving data across active memory, flash, and disk tiers, Oracle **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** currently deliver unmatched scale-up and scale-out database performance, for the most concurrent applications and mixed workloads, at the lowest cost in the industry⁷.

Solving mediocre-to-poor user productivity

Aiming at the [Doherty Threshold](#) users prefer detailed in "[The Economic Value of Rapid Response Time](#)," Oracle focused **Exadata X9M** performance on ensuring exceptional user productivity. At ≤ 19 μs latency, it delivers in spades by proving the industry's lowest and best database response times to applications. That enables the applications to provide sub-second response times to the users, resulting in much improved productivity. Saving time for applications provides faster response times and results.

As good as **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** are in dramatically improving user productivity, that's only part of the story. For analytic applications such as machine learning, data warehouses, blockchain, and more, they deliver faster actionable insights, which leads to faster-time-to-market. Faster-time-to-market means realizing significant revenues and/or cost savings previously unobtainable.

Exadata Cloud@Customer X9M solves specific on-premises problems no one else solves

When it comes to resolving the excessive time-consuming database management problem and regulations, rules, or laws that prevent IT organizations from on-premises cloud services, **Exadata Cloud@Customer X9M** is currently matchless. That may sound like hyperbole but looking at **Exadata Cloud@Customer X9M** shows it's more of an understatement.

Solving extensive time-consuming database management

First and foremost, there is extensive automation built into all **X9M** platform iterations. This engineered system is architected to be simple with the bare minimum of user administration. And it succeeds brilliantly. Perhaps more importantly, Exadata Cloud@Customer runs not only the Oracle Database, but also Oracle Autonomous Database (ADB).

Oracle ADB running on **Exadata Cloud@Customer X9M** nearly eliminates all database management tasks because ADB is "self-driving, self-securing and self-repairing." Self-driving means all of the database and infrastructure management, monitoring, indexing, and tuning processes are automated.

⁷ Cost will be explored in more detail later in this research.

Self-securing means inherent fundamental capabilities that protect against both external and internal (malicious insiders) attacks are also automated. Security is a major emphasis for ADB. Self-securing additionally means timely vulnerability patching is automated online without scheduled outages. It also means every configuration is secure with full end-to-end encryption. More importantly, the data is hidden from the cloud administrators in the cloud customer's Database Vault. And Oracle goes significantly further beyond these important measures with "Data Safe."

Data Safe adds an unprecedented level of security to ADB with full database risk assessment, database user risk assessment, database auditing, database sensitive data discovery, and data sensitive masking. Oracle does not charge for Data Safe. The self-securing Oracle Autonomous Database reduces cyberattack concerns without DBA expertise or intervention.

Self-repairing means automated database healing, again without time-consuming DBA expertise or intervention. This greatly reduces any outage downtime from unplanned maintenance. Oracle ADB customers' production experience has proven to be zero downtime for normal database operational lifecycle. That lifecycle includes patching of all types and averages less than 150 seconds of downtime in any month experiencing unplanned outages. What's remarkable is that Oracle does not put caveats around those numbers in its SLAs. There are none of the exceptions prevalent from other cloud service providers. This translates into Oracle ADB delivering the highest database uptime in the industry.

By enabling Oracle ADB in their **Exadata Cloud@Customer X9M** services, Oracle has made extensive database management a "non-issue." The **Exadata Cloud@Customer** combination of the **Exadata X9M platform** and Oracle ADB makes the management autonomy so powerful that DBAs will take it for granted. DBAs can now spend much more of their time on productive strategic initiatives or application development instead of time-consuming manual labor aka "grunt work."

Those are not the only automation gains with the **X9M** platform. Extensive **Recovery Appliance X9M** automation appreciably reduces backup admin tasks on average by $\geq 80\%$ —based on previous [Wikibon research](#). This frees up the backup admin for more strategic work such as disaster recovery and cyber resilience testing on an ongoing basis.

[Solving regulations, rules, and laws preventing too many organizations from on-premises cloud services](#)

Enabling on-premises cloud services to meet regulatory requirements is conceptually simple. In practice it's not. Solving this problem requires the customer to control any and all access to the on-premises cloud services. That means they must have:

- Complete tight control and auditing of database cloud access.
- Equally tight control and auditing of the infrastructure supporting the database cloud services.
- Control that extends to physical and virtual access by the cloud service provider staff.

Exadata Cloud@Customer is the first on-premises cloud database service to solve these problems with a series of unique customer controls. **Preventative controls** accessible via Oracle Cloud Infrastructure (OCI) empower customers with no login shell/ssh permitted until the customer has authorized it, Oracle operator shell restricted by customer-authorized Access Control Profiles and named users only with no direct access to service accounts. **Detective controls** accessible via OCI Audit provide customers full Oracle operator command logging, optional keystroke logging and identification of individual Oracle staff. **Responsive controls** provide end-to-end termination of TCP connections, processes, and child processes.

These once again are unique **Exadata Cloud@Customer X9M** capabilities for open on-premises cloud services to IT organizations previously locked out of the cloud. No other on-premises cloud service provider—including Amazon, Google, Azure, and IBM—is currently doing this.

[Recovery Appliance X9M also solves sub-optimal database availability from planned and unplanned disruptions-outages](#)

Oracle databases running on **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** increase database availability and uptime from both planned and unplanned disruptions/outages. It does this via tight controls between the Exadata hardware infrastructure and Oracle Database software, non-disruptive patching, Real Application Clusters (RAC), recovery manager (RMAN), and more. As previously discussed,

Exadata Cloud@Customer X9M automates these processes. But what about disruptions/outages such as site failures, system failures, accidental deletions, data corruptions, or even malware/ransomware? That's where **Recovery Appliance X9M** comes into play.

Improving database availability

The **Recovery Appliance X9M** is the only modern data protection system specifically designed for Oracle databases running on any platform and, of course, Oracle Exadata. It uses the latest **X9M** platform optimized for lower cost data protection storage.

Unlike any other data protection system, **Recovery Appliance** is specifically architected with RMAN. **Recovery Appliance** delivers continuous protection for critical and non-critical databases while offloading all backup processing from production servers, minimizing overhead. It offloads backup CPU-intensive processes, freeing up as much as 25% of Oracle Database server resources. This means users can reduce their Oracle Database infrastructure by up to 25%. Or they can increase their available Oracle Database processing performance by approximately 33% without actually adding any resources.

What's more impressive is the fact that **Recovery Appliance** is specifically designed to eliminate the loss of critical database data while significantly improving Oracle Database availability and uptime. It includes Oracle Database data protection generally not available with any other products or services, most of which radically improves database availability:

- Real-time redo transport.
- Continuous transaction-level protection, minimizing data loss exposure.
- Recoveries to any point-in-time.
- Individual transaction recoveries.
- Column or row recoveries.
- Efficient replication and synchronization with other **Recovery Appliances** providing backup and recovery anywhere.
- Autonomous tape archival.
- End-to-end database recovery validation.
- Recoverability status dashboard.
- Incremental-forever backups, creating space-efficient virtual full backups.
- Cyber-Vault architecture, adding significantly in the battle against malware and ransomware.
- Separation of Duty between backup operator and appliance administrator to control access to the system.
- Enterprise-wide database-level protection policies.
- Database-aware space management in a cloud-scale architecture.
- Average of 10x database backup space reduction via Oracle block-aware incremental forever and compression versus typically 2-3x for deduplication data protection with PBBAs.
- Archive to OCI for long-term retention directly from **the Recovery Appliance** and restore to any location.

The latest **Recovery Appliance X9M** version additionally increases backup capacity by 30% and protects more databases with longer retention rates. A new entry level **Recovery Appliance X9M** base rack comes with 207 TB of physical capacity—2 PB effective backups, and a 15 TB/hour restore rate while lowering price by 50% and expanding the potential market to smaller organizations. A full **Recovery Appliance X9M** rack comes with more than 1 PB physical capacity—13 PB of effective backups, and a 24 TB/hour restore rate. These are very impressive database-specific restore rates compared purpose built backup appliance systems from Dell EMC Data Domain, Infinidat, Cohesity, and Rubrik.

What **Recovery Appliance X9M** means to Oracle Database users is more automation, more uptime, less database processing hardware for the same performance. Or more performance with the same hardware, with faster and more granular recoveries at the same price as the previous generation at a much lower entry-level price.

Solving Prohibitive Database Costs

Oracle **Exadata Cloud@Customer X9M**, **Exadata Database Machine X9M** and **Recovery Appliance X9M** all reduce TCO. Some reading this will immediately reject that assertion. That would be a mistake. Here's why.

As previously discussed, **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M** deliver blazing fast Oracle Database performance. That performance translates into much faster application response times. Those faster response times convert into significantly greater productivity, much faster time-to-actionable-insights, and much faster-time-to-market. Depending on the industry, accelerating time-to-market can generate a substantial impact on revenues and increased market share—just as being late-to-market can cost substantial revenues and market share. Consider the example of a major semiconductor manufacturer. Each week they can move up one of their major chips' time-to-market will frequently increase revenues far more than the potential cost of doing so.

Exadata X9M's consolidation of dozens, hundreds, thousands, and tens of thousands of databases delivers a profound reduction in cost. From physical and virtual servers, hypervisors, NICs, cables, transceivers, conduit, switches, storage, rack space, power, cooling, maintenance, allocated overhead, and more. Plus, much lower Oracle Database license costs and reduced operational costs. Remember that Oracle licenses are based on the number of database server cores. But Exadata uses storage server cores to offload a many of the CPU-intensive processes. Oracle does not count those storage cores for Oracle Database licensing and no other storage system can offload Oracle Database queries. That includes all storage systems from Dell EMC, Pure Storage, Nutanix, IBM, NetApp, or any other emerging or established vendor.

Multi-workloads are not the only consolidation saving tons of money. Multi-database types running on **Exadata X9M** also gain from that unprecedented performance. From traditional transactional databases to data warehouses, JSON, XML, time series, AI machine learning, block-chain, spatial, and graphic databases all running in the same system at the same time with the same storage and data. This eliminates costly database data storage silos. It also eliminates additional cost in ETL tools, services, license costs, supporting hardware infrastructure, personnel, training, management, troubleshooting, and time. Especially time.

Exadata Cloud@Customer X9M also appreciably reduces costs via consolidation and also, by the way the service fees are structured. The two principal components of monthly billing are a subscription fee for the **Exadata X9M** hardware and a separate compute charge per second for Oracle Database processing. Oracle charges per vCPU per second. The key to this charge is time. Query processing time affects cost. Shorter equals lower cost. The unparalleled performance of **Exadata Cloud@Customer X9M** means shorter processing times and a lower bill vs. AWS Outposts, Google Anthos, or Azure Stack.

The **Recovery Appliance X9M** further reduces costs as previously discussed. It offloads much of the Oracle Database backup processes, saving up to 25% of the database infrastructure. That's a huge savings. And because the **Recovery Appliance X9M** can protect many Oracle databases from dozens to hundreds, thousands, even tens of thousands, it saves on data protection systems. Savings that include the data protection licensing commonly based on protected TB, socket, core, or machine (physical or virtual); and the hardware infrastructure supporting that data protection including servers, storage, VMs, containers, NICs, cables, transceivers, conduit, switch ports, switches, rack space, allocated overhead, power, cooling, training, management, troubleshooting, and more. Once again, significant cost savings.

The **Recovery Appliance X9M** additionally brings an unprecedented level of automation as previously covered in extensive detail in Dave Vellante's Wikibon research "[Oracle's Recovery Appliance Reduces Complexity Through Automation](#)." That research revealed **Recovery Appliance's** extensive automation for \$5B or greater Enterprise organizations running Oracle mission-critical applications reduced their costs on average by 31% over purpose-built-backup-appliances (PBBA) such as Dell EMC Data Domain (since rebranded as "PowerProtect"). It reduced the average number of steps in backup and recovery by more than 80%, the probability of a recovery error by 96%, and the probability of a manual recovery or data loss by 98%. The reduction in downtime costs alone averaged \$370M for the global 2000 companies that are \$5B in size or greater organizations running mission-critical applications on Oracle Database. That's a massive savings for any organization.

These substantial TCO savings frequently cover the costs of **Exadata Cloud@Customer X9M**, **Exadata Database Machine X9M** and/or **Recovery Appliance X9M**. Yeah but, aren't all Exadata platforms

expensive? That’s what many competing vendors assert. That’s known as deception by omission. It’s the net total cost of ownership (TCO) that matters. Failure to look at that is analogous to comparing a new battery electric vehicle (BEV) versus a traditional internal combustion engine vehicle and not looking at the government incentives on BEVs and the much lower cost of charging a battery versus buying gas. It paints a distorted picture.

To evaluate the TCO for **Exadata Cloud@Customer X9M**, **Exadata Database Machine X9M** and/or **Recovery Appliance X9M** requires adding back in all of the savings into a TCO calculation, including dramatically reduced database sprawl; reduced database server hardware; reduced storage hardware; reduced networking hardware; reduced database licenses/subscription costs; reduced power and cooling; reduced maintenance; reduced rack space; reduced allocated overhead; reduced training; reduced troubleshooting; reduced maintenance; reduced management costs; increased user productivity; faster time-to-actionable-insights; faster time-to-market; increased revenues; increased profits; and higher uptime with reduced downtime costs.

Speaking of competitors, the next research section looks at how **Exadata Cloud@Customer X9M**, **Exadata Database Machine X9M** and/or **Recovery Appliance X9M** match up against major competitors.

How Oracle’s X9M platforms—Exadata Cloud@Customer, Exadata and Recovery Appliance—Match up against Major Competitors

Each of the Oracle’s **X9M platforms** compete with different vendor’s offerings. How competitors fall short in solving the problems discussed in this research is crucial to understanding why they are ineffective offerings.

Exadata X9M Database Machine and the Competition

In the case of the **Exadata X9M Database Machine**, it doesn’t have any actual direct equivalent competitive offerings. It has partial competitors in hyperconverged infrastructure (HCI), high-performance external storage with high bandwidth with low latency networks, and high-performance servers. Each is only a partial solution to the previously discussed costly problems. They cannot be considered in the same class as the **Exadata X9M Database Machine**.

Exadata X9M Database Machine vs HCI-Nutanix

Exadata X9M Database Machine technically competes with HCI systems such as Nutanix, by far the performance leader in the HCI space. A close look at Nutanix’s Oracle Database performance shows Nutanix has made major clever performance improvements over the years. They have implemented fast NVMe drives, RDMA NICs, and even a couple of storage-class memory (SCM) Optane drives per node. Their latest Oracle Database performance numbers⁸ are nearly double what they were previously. However, they’re still noticeably short of **Exadata X9M Database Machine**. Nutanix derived these numbers using a SLOB⁹ test in a 4-node (2U each node) cluster running 8 discrete Oracle Database instances, consuming 8U of rack space. The results are generally the best possible performance taken under ideal conditions. It utilizes a single IO size and pattern for each measurement. Real-world database processing tends to have a variety of IO sizes, patterns, and workloads. A look at the numbers shows that even ideal numbers come up way short. To provide an apples-to-apples rack comparison means multiplying Nutanix’s best published results by 5.

Per Rack	Oracle Exadata X9M	Nutanix	Oracle Advantage
Latency - <i>lower is better</i>	≤ 19µs	≤ 1ms	> 50x
IOPS - <i>higher is better</i>	27.6 M	2.68 M	> 10x
Throughput - <i>higher is better</i>	1 TB/s	Unpublished	N/A

⁸ [Nutanix performance for Database Workloads](#), 24 November 2021

⁹ [SLOB](#) is an Oracle IO workload generation tool kit.

Nutanix is nowhere close to **Exadata X9M Database Machine** in latency—more than 50 times slower. To close the gap in IOPS performance, Nutanix would need 10 racks instead of 1. Even then, it would come up short in performance. Usable capacity also falls well short of Exadata with Nutanix’s limited number of drives per node.

By definition, Nutanix does not solve any of the aforementioned database problems:

- Database sprawl.
- Conspicuous decline in database performance as its data grows.
- Mediocre-to-poor user productivity.
- No enhanced database availability and uptime.
- And definitely no reduction in prohibitive database costs.

[Exadata X9M vs External Storage: Pure Storage FlashArray//X90R2](#)

Storage vendors love to bash Oracle Database performance running on HCI and physical servers. They never compare their performance to Exadata for good reason. They come up very short. Pure Storage recently released their Oracle Database performance results in a [blog](#). They claim to set a new bar for Oracle Database performance with their FlashArray//X90R2. And yet, what they are comparing to is an older Pure Storage FlashArray//M70, which was hardly the old industry bar.

Nevertheless, Pure Storage tested with SLOB and again under ideal conditions. The FlashArray//X90R2 was configured with NVMe Flash SSDs. The company utilized an 8-node Oracle RAC environment hosting two database instances. Assuming they used small 1U servers and a typical 6U FlashArray//X90R2 system, the configuration consumed 14U. A full rack would house 3 such systems. To provide a relative apples-to-apples rack comparison means multiplying Pure Storage’ best published results by 3. This isn’t quite applies-to-apples as it does not account for the rack space required for storage network switches. However, the servers can be directly connected to the storage ports.

Per Rack	Oracle Exadata X9M	Pure Storage FlashArray//X90R2	Oracle Advantage
Latency - <i>lower is better</i>	≤ 19µs	≥ 250 µs	> 13.1x
IOPS - <i>higher is better</i>	27.6 M	~1.473 M	> 18.7x
Throughput - <i>higher is better</i>	1 TB/s	39 GB/s	> 25.6

The Pure Storage results, like Nutanix’s, are nowhere near Exadata X9M in latency, IOPS, or throughput and that’s with three complete systems.

Pure Storage announced on 8 December 2021 a new higher-end to their all-flash array line called [FlashArray//XL170](#). This is a bigger—5 to 11U—FlashArray, albeit denser leveraging quad-level cell (QLC) flash SSDs. They claim latencies as low as 150 µs with SCM drives, which is a 40% improvement. But still, **Exadata X9M Database Machine** is **7.9x** faster. They claim a significant improvement in throughput up to 36 GB/s. However, those aren’t SLOB results. But for the sake of argument, assume it is. Also assume they can still cram three systems at 6U each, i.e., not full systems that require as much as 11U. That would take the throughput up to approximately 108 GB/s per rack. **Exadata X9M Database Machine’s** throughput per rack is still **9.26x** greater. Assuming the Oracle Database IOPS increases the same as Pure Storage claims of up to 80% greater, the IOPS would grow to 3.271 M IOPS per rack. **Exadata X9M Database Machine’s** IOPS per rack is still **8.44x** better. A lot of assumptions in Pure Storage’s favor and still nowhere close. They fail to move the needle very much on the aforementioned database problems. And using simple deductive logic, their TCO are likely to be considerably higher because of it.

[Exadata X9M Database Machine vs External Storage: Dell EMC PowerMax](#)

Dell EMC is still the worldwide leader in external storage arrays. Their current most powerful and fastest array—per their data sheets—is the PowerMax 8000. Dell EMC has not published any Oracle Database results at the time of this research. For the sake of argument, Dell EMC’s purported performance claims will be assumed to be the performance for Oracle Database. As unlikely as that is, Oracle Database performance

is generally less than Dell EMC's data sheet claims, it doesn't matter for this comparison because it is still nowhere close to the **Exadata X9M Database Machine**.

Dell EMC claims as low as $\leq 100 \mu\text{s}$ latency, up to 15 M IOPS, and 350 GB/s throughput in two full racks not including the rack space required for the Oracle Database servers and storage network switches. Normalizing that to a single rack and comparing it to **Exadata X9M** demonstrates that PowerMax 8000 also comes up woefully short of the **Exadata X9M**.

Per Rack	Oracle Exadata X9M	Dell EMC PowerMax 8000	Oracle Advantage
Latency - <i>lower is better</i>	$\leq 19 \mu\text{s}$	$\leq 100 \mu\text{s}$	$> 13.8\text{x}$
IOPS - <i>higher is better</i>	27.6 M	7.5 M	$> 5.2\text{x}$
Throughput - <i>higher is better</i>	1 TB/s	175 GB/s	> 5.7

The Dell EMC PowerMax 8000 just does not measure up. The differences between Dell EMC PowerMax 8000 and Oracle **Exadata X9M Database Machine** are huge, and PowerMax clearly needs some more power to take it to the max. These are eye-popping material differences. And like Pure Storage, Dell EMC PowerMax 8000 makes an insignificant dent in the database problems previously discussed.

Exadata X9M Database Machine vs IBM Power Systems E1080

IBM has not published any Oracle Database performance specifications on its latest POWER10 systems. Nor did they publish and Oracle Database specific specifications in the previous POWER9 systems either. However, based on previous [Prowess](#) research, POWER9 was severely deficient compared to Exadata X8M in performance, TCO, availability, scalability, power, and more. IBM says POWER10 is 50% faster than POWER9. **Exadata X9M Database Machine** has 87% higher throughput and 72% more IOPS than Exadata X8M Database Machine. Based on those numbers, POWER10 is only falling much further behind Exadata.

It appears to Wikibon, the latest POWER10 is principally aimed at the scale-out supercomputer market with built-in chip accelerators for encryption and AI. It is definitely not aimed at the Oracle Database performance, scalability, simplicity, or automation market based on its non-competitive comparison to the **Exadata X9M Database Machine**.

Exadata Cloud@Customer X9M and the Competition

There are a few **Exadata Cloud@Customer X9M** competitors—AWS Outposts, Microsoft Azure Stack, and Google Cloud Platform (GCP) Anthos. Of these, AWS Outposts is the closest analog to **Exadata Cloud@Customer X9M** (compared below). And Outposts is not very close either. Azure Stack is a multi-vendor approach where the cloud stack is provided by Microsoft Azure and the generic hardware is built by their server partners. This is a multi-vendor, hybrid cloud on-premises cloud service and Microsoft requires customers to hire a minimum of one Azure Stack operator. No IT organization can afford to have only one cloud operator. What happens if that operator is sick or leaves? They typically require at least two. If they require 7 by 24 support, they'll have three or more. That adds \$400K to \$800K in cost before consuming a gigabyte of cloud services. Anthos is a different beast. It's much more hybrid cloud orchestration and management than an on-premises extension of the public cloud. Neither Azure Stack nor GCP Anthos are direct competitors to **Exadata Cloud@Customer X9M**.

Exadata Cloud@Customer X9M vs AWS Outposts

AWS Outposts, an on-premises cloud deployment, provides a neutered subset of the AWS cloud. Database cloud services are limited to MySQL, PostgreSQL, and Microsoft SQL Server. There is no current AWS RDS or Oracle Database support. There is no Redshift or Aurora support, Amazon's two flagship databases. In fact, the vast majority of the AWS cloud database services are not available on Outposts.

AWS Outposts comes in two flavors: Outposts 42U racks and Outposts 1U and 2U servers. Outposts' racks scale to as many as 96 racks at a location.

A major AWS Outposts problem is that it does not provide any co-engineering synergy with the databases running on Outposts. None. It's quite similar to running those databases on any other on-premises server

in a VM or container with one exception. The hardware infrastructure, basic database implementations, management, and patching are done by AWS. There is no automation, enhanced performance, or enhanced database scalability. In fact, performance and scalability are severely constrained by AWS hardware limitations. For example, a database instance is constrained by EC2 limitations of approximately 80,000 IOPS, the same constraints as in the AWS public cloud.

Amazon does not publish Outposts' performance benchmarks, but strangely enough, some of their Outposts storage partners do. Pure Storage claims their FlashArray//X series are 4x faster in latency than native AWS Outposts Elastic Block Storage (EBS). This means that EBS is 4x worse than Pure Storage at approximately 1 ms latency. Pure Storage also claimed approximately 4x faster IOPS. Based on the previous comparisons that revealed Pure Storage FlashArrays//X or XL are not very competitive compared to **Exadata Database Machine X9M**, logical reasoning makes it clear that AWS Outposts at $\frac{1}{4}$ the performance of Pure Storage, is even less competitive.

It is evident that AWS Outposts was not architected for most Enterprise or even small-to-medium Enterprise requirements. More importantly, Outposts does nothing to address or solve any of the ongoing database problems discussed in this research. Outposts exacerbates several of them including out-of-control database sprawl, database performance degradation as database data grows, prohibitive database cost, user productivity, database management, and database availability. Most importantly, AWS Outposts does not currently have a workaround to the regulations, rules, and laws that inhibit too many organizations from implementing on-premises cloud services. For example, users cannot limit, restrict, or audit AWS personnel.

Because AWS Outposts does not address any of the problems discussed in this research, it actually has a higher total cost, and much lower performance than **Exadata Cloud@Customer X9M**.

Recovery Appliance X9M and Competition

There are a many data protection (DP) and disaster recovery (DR) products and services as well as target backup storage appliances on the market. Most work with Oracle Database's RMAN. However, **Recovery Appliance X9M** is a unique beast. It is the only Oracle Database data protection that specifically:

- Offloads RMAN processes from the Oracle Database saving approximately 25% of the database server resources.
- Continuously replicates the redo log.
- Enables finely granular recoveries.
- Provides end-to-end data validation.

Several of those backup appliances market themselves as high performance deduplication targets focusing on backup performance. Very few come close to matching **Recovery Appliance X9M** in raw performance. None match it in Oracle Database backup performance. Candidly, that only matters for the first full backup. After that, most DP products including **Recovery Appliance X9M** backup or replicate only the changes between backups, which is generally not that much data. **Recovery Appliance X9M** performs continuous backup, backing up every change as it happens.

The most important issue is when a major Oracle Database recovery is required. Most backup target appliances do not publish their Oracle Database restoration speeds. One can infer the reason they do not is because it is likely not very fast. If it were, they would publish, promote it on social media and offer temporary tattoos or T-shirts highlighting their accomplishment. But they don't, and the information omission speaks loudly.

Recovery Appliance X9M recovery capabilities for Oracle Database are in a word "matchless." Starting with the fine-grained recovery capabilities, 5x better compression compared to deduplication target storage and PBBAs, backup consolidation for up to tens of thousands of Oracle Databases, and the 24TB per hour recovery speeds, nothing else contests with the **Recovery Appliance X9M**.

When it comes to providing Oracle Database protection and high availability, the **Recovery Appliance X9M** has proven to be best in class.

Conclusion

Technology vendors generally promote their product or service's performance especially when they release a new version. But performance by itself in a vacuum means little. It only matters in how performance solves significant user problems.

The latest X9M generation of **Exadata Cloud@Customer**, **Exadata Database Machine**, and **Recovery Appliance** has radically upped the database performance ante with significant breakthrough specifications. And this is not just for relational databases, but also schema optional databases a.k.a. NoSQL in layperson's terms, XML-Object, JSON-document, data warehouses, graph, spatial, machine learning, blockchain, time-series, and more. These breakthrough capabilities are impressive for sure; however, it is how Oracle has used them to solve several pervasive and rapidly growing database problems that no one else is solving that's more notable. To summarize, below are the problems the new Oracle **X9M** platform solves and briefly how it solves them.

- Out-of-control database sprawl problem.
 - **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M's** performance and capacity scalability, facilitate extensive consolidation, with multi-tenancy, for multi-workload, multi-database type, concurrently¹⁰.
 - Radically reducing database sprawl.
 - **Recovery Appliance X9M** further enhances **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M's** performance, in addition to Oracle Databases running on other hardware, by offloading RMAN data protection from the database servers – giving back up to 25% of the database server processing resources.
- Conspicuous decline in database performance as database data proliferates problem.
 - **Oracle Exadata X9M Database Machine¹⁰** performance and capacity scalability from a base rack to 12 full racks eliminates this as a problem.
 - **ZDLRA X9M** offloads Oracle Database RMAN data protection processing saving up to 25% of the database server processing resources.
- Mediocre to poor user productivity problem.
 - **Exadata Cloud@Customer X9M** and **Exadata Database Machine X9M's** eye-popping performance by definition reduces database application response times, in turn increasing user productivity.
- Excessive time-consuming database management problem.
 - Extensive automation is built into the Oracle's **X9M** platform in general.
 - Oracle Autonomous Database (ADB) only runs on **Exadata** in Oracle Cloud Infrastructure (OCI) and on-premises in **Exadata Cloud@Customer**. ADB is the only level 5 autonomous database in the market today, eliminating the vast majority of database management, training, troubleshooting, etc.
 - **Recovery Appliance X9M** further reduces backup and recovery tasks by 80% on average for G2000 \$5 B Enterprise accounts.
- Regulations, rules, and laws are preventing too many organizations from on-premises cloud services.
 - **Exadata Cloud@Customer X9M** provides users with the means to regain security control for on-premises cloud services with full Operator Access Control. It enables users to meet regulation, rules, and law requirements for on-premises cloud services.
- Sub-optimal database availability from planned and unplanned disruptions/outages problem.
 - Oracle's **X9M** platform is extremely reliable with \geq six 9s of up time. However, there are always human errors, accidental deletions, site failures, natural disasters, etc. That's where the **Recovery Appliance X9M** changes the game. It is the only database data protection system specifically designed for Oracle Database. Highly automated, it offloads the data protection processes from the database servers. It has finely granular point-in-time and very fast recoveries with no data loss. The net of all this is a massive reduction in recovery errors and downtime costs and much greater database availability.
- Prohibitive database TCO problem.

¹⁰ Exadata Cloud@Customer X9M is currently limited to one rack, whereas Exadata X9M Database Machine scales to 12 racks.

- The **Exadata Cloud@Customer X9M**, **Exadata Database Machine X9M** and **Recovery Appliance X9M** each reduces so much cost from solving the aforementioned problems, that many times TCO calculations indicate that they more than pay for themselves.

Many vendors tout their performance. Performance is a key metric. But when compared to the Oracle **X9M** platform, it's the equivalent of holding a candle against a 20,000 lumens spotlight. There is no comparison. That **X9M** performance changes the game and solves several obstinate database user problems.

Wikibon's recommendation for medium to large to extra-large Oracle Database users, is to adopt this latest Oracle **X9M** platform. The decision is simple and the rewards ample.

More Information

[Oracle Exadata X9M](#)

[Oracle Exadata Cloud@Customer X9M](#)

[Oracle Dedicated Region Cloud@Customer](#)

[Oracle's Recovery Appliance X9M](#)

[Oracle Autonomous Database](#)