Cloud Database Battle: AWS vs. DIY vs. Oracle   by, David Floyer

January 18th, 2021

The first premise of this research is that architecting the Oracle Cloud Database service to run on specialized hardware and software, either on-premises or in Oracle Cloud Infrastructure (OCI), allows the cloud database vendor to reduce costs significantly. This approach also allows the vendor to provide autonomous services based on economies of scale that further reduce the operational support costs. The combination of the two methods leads to halving the cost of running today's cloud database application suites.

The second premise is that future synchronous automation of business processes will require real-time integration between systems-of-record, advanced analytic/AI inference systems, and other data and cloud database types. This integration can only be achieved by sharing data between database types. Also, the operation of synchronous applications is too complex for traditional operational processes. Therefore, high levels of cloud database and application automation, and machine learning are imperatives for synchronous application deployment.

Oracle Cloud Database is Tier-1 and in a class of its own. Wikibon recommends that larger enterprises with mission-critical workloads should not convert from Oracle to other databases. Instead, Wikibon recommends migrating to Autonomous Cloud Database on Oracle Exadata Cloud@Customer X8M, Oracle Exadata Cloud Service on OCI, or other Oracle Database cloud services. Wikibon recommends that enterprises minimize the number of separate databases and data types and use the converged Oracle Cloud Database instead.

Cloud Database Battle: AWS vs. DIY vs. Oracle

At this time, Wikibon cannot recommend running large-scale Oracle Database mission-critical workloads and the surrounding portfolio of applications on AWS. The cost of running Oracle databases in AWS is prohibitive.

Wikibon recommends Microsoft as the best multi-cloud alternative for Oracle mission-critical workloads because of its adjacent Microsoft Azure strategy combined with Oracle Exadata Cloud technology.

Senior executives should press Oracle and AWS to bury the hatchet and develop a win-win-win cost-effective multi-cloud database services strategy for their joint customers.

Cloud Database Premises

The first premise of this research is that architecting the Oracle Cloud Database service to run on specialized hardware and software, either on-premises or in Oracle Cloud Infrastructure (OCI), allows the cloud database vendor to reduce costs significantly. This approach also allows the vendor to provide autonomous services based on economies of scale that further reduce the operational support costs. Combining the two methods leads to halving the cost of running today's cloud database application suites.

The second premise is that future synchronous* automation of business processes will require real-time integration between systems-of-record, advanced analytic/AI inference systems, and other data and cloud database types. This integration can only be achieved by sharing data between database types. Also, the operation of synchronous applications is too complex for traditional operational processes. Therefore, high cloud database and application automation and machine learning are imperatives for synchronous application deployment.

* *For a more in-depth discussion, see the "Application Support for Synchronous & Asynchronous Business Processes" section in the Footnote at the end of this research.*

Executive Summary

Cloud Database for Existing Systems

A cloud database resides on a private cloud, public cloud, hybrid cloud, and multi-cloud environments. From an application perspective, the database services are identical. The only difference lies in where the database resides. There is an option for databases to be fully managed from a public cloud as an Autonomous Cloud Database from an operational perspective.

This Wikibon research focuses on the IT Operational costs for running large mission-critical applications on Oracle Cloud Database Enterprise Edition in an on-premises cloud environment. It compares the operational IT budget costs for a traditional IT datacenter with a multi-vendor Do-It-Yourself (DIY) infrastructure, AWS running Oracle RDS on Outposts**, and Oracle Autonomous Database running on Oracle Exadata Cloud@Customer X8M.

Cloud Database Battle: AWS vs. DIY vs. Oracle

Traditional Datacenter 96% Higher & AWS Outposts 90% Higher than Autonomous Exadata Cloud@Customer for Large-scale Mission-critical Oracle Workloads

Y-axis: 4-year Infrastructure Costs, including Maintenance and Operational Costs for System & Database for Typical $2B Enterprise

- Traditional IT Datacenter: $40.1
- AWS RDS on Outposts (Projected): $38.7
- Autonomous Database on Exadata Cloud@Customer X8M: $20.4

Source: © Wikibon 2021 Note: Full Stack Systems assumes AWS Outposts RDS Oracle (Projected), Traditional IT Datacenter, Autonomous Exadata Cloud@Customer XLM
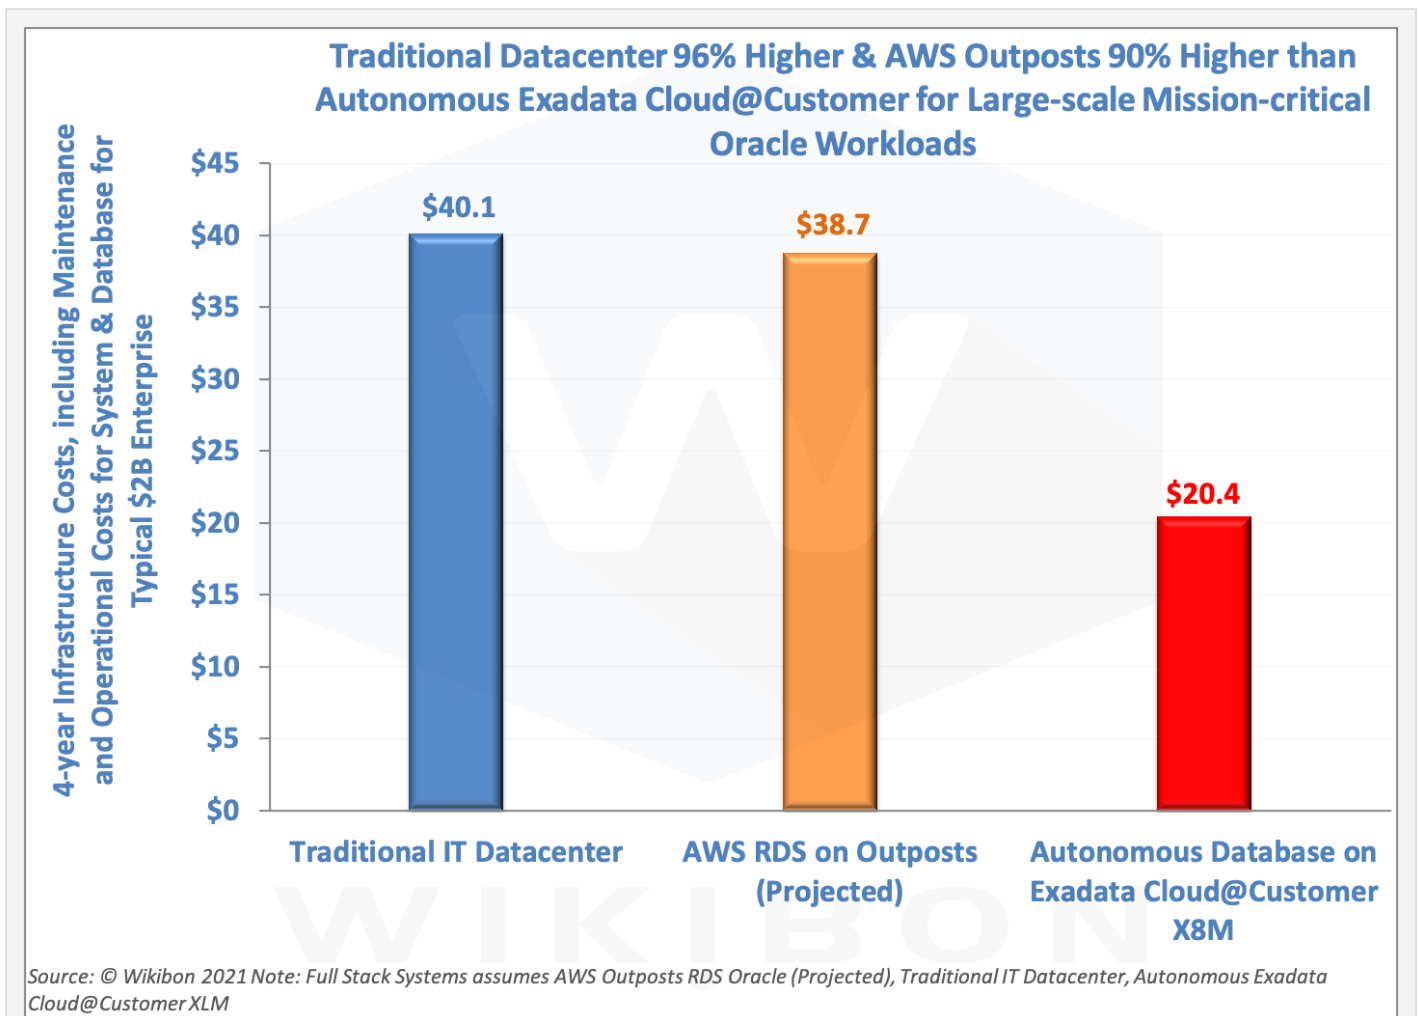
Figure 1: Executive Summary of IT Cost Comparison between Traditional IT Datacenter, AWS Outposts (Projected), and Autonomous Cloud Database running on Oracle Exadata Cloud@Customer X8M for Large-scale Mission-critical Applications for a typical US$2 billion enterprise.
Source © Wikibon 2021

Figure 1 shows the results of this research. The y-axis shows the four-year infrastructure, operational support, Oracle Database licensing, and additional datacenter costs. The blue column shows 4-year costs of $40.1 million for running mission-critical applications in a traditional on-premises datacenter with a traditional on-premises Oracle database. The orange column shows a slightly lower figure of $38.7 million for running Oracle as a cloud database in the RDS service on AWS. The third column in red shows a much lower figure of $20.4 million, which is the costs of running the Oracle cloud database on-premises on Exadata Cloud@Customer X8M.

The conclusion from Figure 1 shows that compared to Oracle Autonomous Cloud Database on Exadata Cloud@Customer X8M, the cost of running the same large mission-critical Oracle-based applications in a Traditional IT datacenter is 96% higher and on AWS RDS on Outposts is 90% higher.

Most large enterprises run their mission-critical systems on Oracle databases. The business conclusions are that there is no business case for migrating large mission-critical Oracle applications running on traditional on-premises ("DIY") environments to an AWS Oracle RDS cloud database. There is a strong business case for enterprise IT management to migrate Oracle Database large mission-critical applications, running on

either a "DIY" or on AWS Oracle RDS environments, to an Oracle Autonomous Cloud database running on Exadata X8M.

The "Autonomous Database on Exadata Cloud@Customer X8M Architecture" section below gives a more detailed break-out of costings and all the assumptions.

Converged Cloud Database for Future Systems

Synchronous automation of business processes is essential for improving business efficiency and cycle time. The ability to process enormous amounts of data in real-time or near real-time is at the heart of developing synchronous systems. The compute power to process that data must be adjacent to the data because the costs and elapsed time to move large amounts of data to a compute resource are too high to achieve synchronicity.

Multiple database types are increasingly required to process this enormous amount of data. There are transactional databases for systems-of-record, advanced analytic and graph databases for analysis systems, AI databases for inference, document databases, time-series databases, blockchain databases, and many more.

The traditional approach of using separate "best-of-breed" database types for each application means complex data transfers and transformations are required between databases. Wikibon concludes that this architecture will make future synchronous applications needing to blend different data types from various database types much more expensive and challenging to implement, if not impossible. Synchronous applications require that the databases share the same data in real-time because it simply takes too long to move data from one database to another.

Autonomous Cloud Database for Future Systems

Due to the complexity and real-time requirements of synchronous systems, traditional system operators, storage admins, and database administrators (DBAs) cannot react fast enough or accurately enough to move data between isolated databases and manage availability levels and recoverability required. Instead, the hardware and software vendor that offers these systems must provide full automation of the stack, including hardware, operating systems, and databases. Also, the vendor must run this automation for many customers to achieve economies of scale.

Summary of Conclusions

This analysis means that deploying a single autonomous converged or universal database that supports different data and database types is the correct long-term strategy for most large organizations. These database architecture benefits are much greater than the operational savings examined in this study, and Wikibon plans to address this in future studies.

At the moment, Wikibon assesses that Oracle is the leading provider of a converged database, and Oracle has clearly stated it intends to continue investing in ensuring performance, automation, and integration of different databases and data types within the Oracle Database.

Wikibon recommends that IT executives deploy their Oracle mission-critical workloads on Oracle Autonomous Database on Exadata Cloud@Customer X8M to reduce costs today and prepare for the synchronous applications soon.

Architecture Requirements for Modern Systems of Record

Synchronous Business Processes

All enterprises have a mixture of synchronous and asynchronous business processes. Using data to

increase the percentage of synchronous business processes is the most effective way to improve employee and partner productivity, improve customer satisfaction, reduce the cost of doing business, and reduce business cycle time.

Data-driven enterprises can accelerate productivity from synchronous automation by marrying OLTP systems of record and other real-time analytic/AI database systems. However, the systems architecture and services must integrate and optimize the different components to allow organizations to meet real-time response-time requirements. There is a significant simplification of the architecture when a single database can share separate databases and data types.

Cloud Database Architecture

The traditional systems of record are synchronous, and online transactional processing (OLTP) systems are installed in almost every enterprise.

Traditionally the analytic systems were batch systems, and the technology was not available to run analytics systems in real-time. However, infrastructure technology has improved dramatically to process far more significant amounts of data in real-time. Also, many additional data and database types have been created, which can operate much more efficiently to different analytics and support business requirements in real-time. These databases often start as standalone databases to solve specific problems. The successful ones (like MongoDB for document handling) create their own ecosystem.

Moving data from one database system for use in another requires data transformation and data transfer, both of which can take significant elapsed time. The more databases and data types that exist, the more specialized transfer systems are required. Sixteen databases would require 120 different transformation/transport systems. Fifty databases would require 1,225. This approach is not sustainable because it leads to data fragmentation, data inconsistency, security holes, and massive data management costs. Also, it does not scale and is not a suitable platform for creating synchronous solutions.

The ideal architecture implements an autonomous unified or universal cloud database, with all the data directly available to all applications. This architecture makes possible far richer application systems. For example, a system of record can ask for real-time analytics to help automate a business process. This analytics process requires shared databases running on a high-performance, purpose-built, and optimized stack. The cloud database and data types need to include relational operational, advanced analytic, AI learning and inference, document, graph, key-value, in-memory, blockchain, etc. Wikibon expects new databases and data types to evolve.

This approach provides a high-performance and flexible cloud database system for synchronous application systems. It eliminates the data transfer and transformation processes necessary to move data across multiple specialized databases. It also avoids data fragmentation and data consistency issues, enhances security, and reduces data management costs.

Impact on Cloud Database Vendors

Oracle is the leading cloud database vendor for systems of record and analytics. Oracle has invested heavily in creating a hybrid-cloud database system that allows the same database and hardware to be available on-premises, in Oracle Cloud Infrastructure (OCI), and other cloud datacenters such as Equinix. Their database architecture approach provides all database, and data types within a single converged or universal database.

Oracle has designed an Autonomous Cloud Database architecture to automate the management of the most mission-critical workloads, including provisioning, tuning, clustering, disaster protection, elastic scaling, securing, and patching. As a result, many of the traditional manual processes and human errors are eliminated. This autonomous cloud database solution delivers far lower costs and significantly lower risk for the most demanding synchronous workloads relative to alternative approaches.

Wikibon believes that this type of database architecture is critical for implementing data transformation initiatives that most enterprises have embarked on. Gartner has issued a report on Cloud Database Critical Capabilities, which defines some of the above requirements. Not surprisingly, Oracle Autonomous Database is the leader in most of the categories.

This should be a wake-up call for AWS, which has taken the approach of offering 16 different mainly open-source databases and supporting them well in the AWS PaaS and AWS infrastructure. There is support, including provisioning, tuning, disaster protection, elastic scaling, securing, and patching. AWS has provided some interconnects between some databases. However, AWS will need to invest far more heavily in integrating these databases to provide an effective cloud database system for mission-critical synchronous applications, which have the highest value for enterprises.

Autonomous Database on Exadata Cloud@Customer X8M Architecture

Oracle Exadata Compute

Oracle Exadata Cloud@Customer X8M's scale-out architecture separates storage from compute and provides high-performance storage nodes and compute nodes. It also enables scaling up and down without interrupting database operations while customers pay only for the resources used. The billing is in 1-second increments, with a 1-minute minimum.

Oracle Exadata Internal Networking

The RDMA over Converged Ethernet (RoCE) network architecture is the foundation for Exadata X8M's OLTP and analytics performance. RoCE provides fast 100 Gbits/Sec, low-cost compute, and storage interconnectivity based on industry-standard Ethernet. Wikibon believes that 400 Gbit bandwidth will be available soon. Also, NVMe storage allows an any-to-any low-overhead connection between servers and storage. The levels of performance achieved by Exadata X8M also power automated database management, including Autonomous RAC clusters and high availability capabilities such as Autonomous Data Guard.

Oracle Exadata Storage & IO Latency

Operational databases are usually stateful. Fast locking and logging dramatically improve throughput and performance consistency to stateful and complex databases.
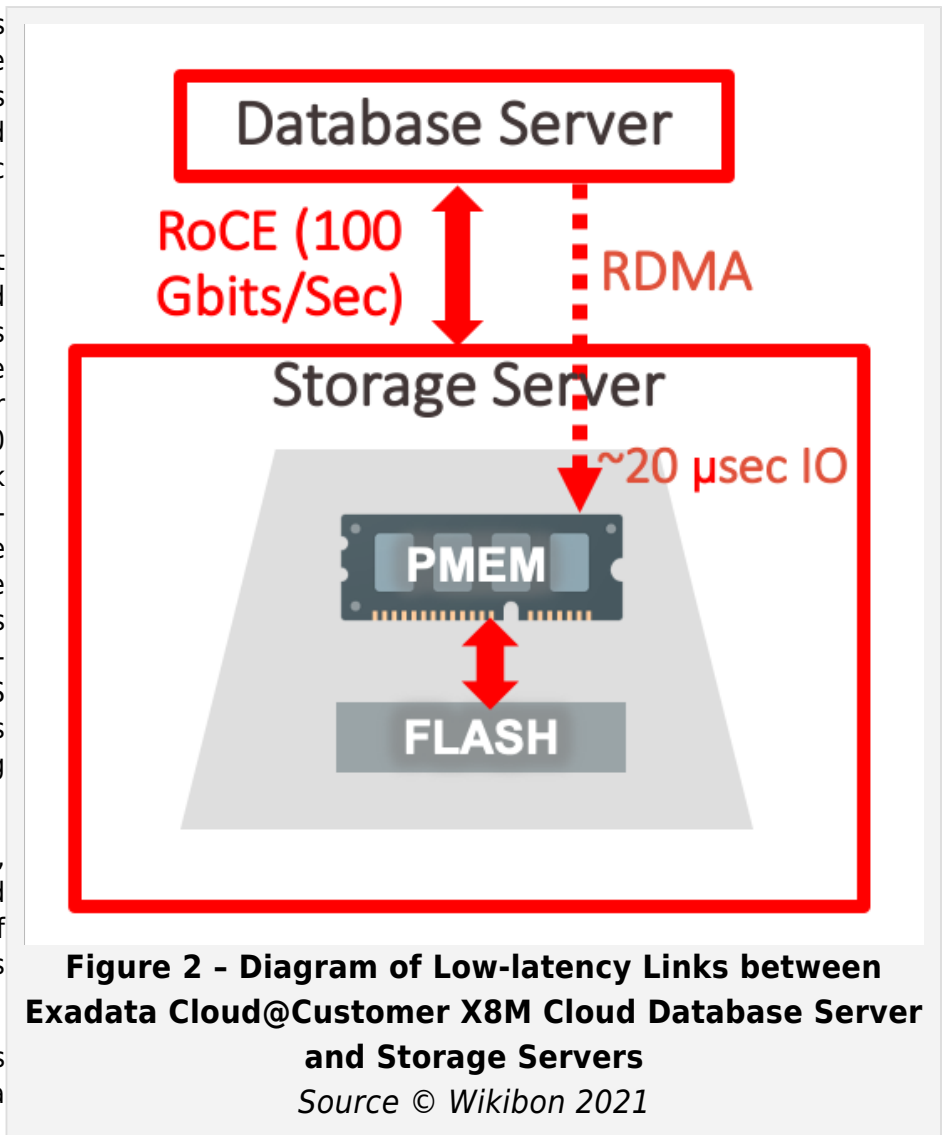
Exadata's smart storage servers process queries offloaded from the compute servers. The data is columnized to speed up finding and analyzing the data for analytic applications.

Exadata X8M delivers the lowest latency storage IO for cloud providers of about 20µsecs, as shown in Figure 2. The database server software uses RDMA over Converged Ethernet (RoCE), a 100 Gbits/Sec internal network, to talk directly to a 1.5 Terabyte non-volatile Persistent Memory Module DIMM (PMEM) slot in each storage server. This ~20µsec IO latency is faster than any other public cloud-native IO performance (e.g., AWS RDS or Microsoft Azure). PMEM is also used extensively for writing redo logs and caching data.

The combination of performance, elasticity, scalability, and automated patching and indexing of Exadata Cloud@Customer X8M is unrivaled in the industry today.

Previous Wikibon research provides additional details of the Exadata X8M architecture.

Figure 2 – Diagram of Low-latency Links between Exadata Cloud@Customer X8M Cloud Database Server and Storage Servers
Source © Wikibon 2021

Business Case for Autonomous Cloud Database on Exadata Cloud@Customer X8M

Wikibon compared Oracle Exadata Cloud@Customer with a Traditional Datacenter and a Microsoft Azure Stack approach. This new research took the latest Oracle Cloud@Customer offering based on Exadata X8M with Autonomous Cloud Database and updated the comparison with the Traditional IT Datacenter and AWS RDS on Outposts. AWS has not announced official support for Oracle Database on RDS. However, Wikibon believes it will be announced shortly.

Figure 3 below shows the results of the comparison. The y-axis shows the 4-year IT budget costs, including a detailed breakdown of datacenter infrastructure costs, maintenance and operational costs for the system, and Oracle software licenses. Figure 3 shows that compared to Oracle Autonomous Database on Exadata Cloud@Customer X8M ($20.4 million), the cost of running the same large mission-critical Oracle-based workloads in a traditional IT datacenter (US$40.1 million) is 96% higher and on AWS RDS on Outposts (US$38.7 million) is 90% higher. In other words, the costs for DIY and RDS on AWS Outposts are nearly 2X higher than the Oracle solution.

Cloud Database Battle: AWS vs. DIY vs. Oracle

Traditional Datacenter 96% Higher & AWS Outposts 90% Higher than Autonomous Exadata Cloud@Customer for Large-scale Mission-critical Oracle Workloads

	Traditional IT Datacenter	AWS RDS on Outposts (Projected)	Autonomous Database on Exadata Cloud@Customer X8M
Datacenter Infrastructure	$17.3	$14.0	$9.9
Operational Support	$8.6	$7.9	$0.8
Oracle Database	$8.6	$12.3	$6.7
Additional Datacenter Costs	$5.6	$4.5	$3.0
Total	$40.1	$38.7	$20.4

Source: © Wikibon 2021 Note: Full Stack Systems assumes AWS Outposts RDS Oracle (Projected), Traditional IT Datacenter, Autonomous Exadata Cloud@Customer XLM
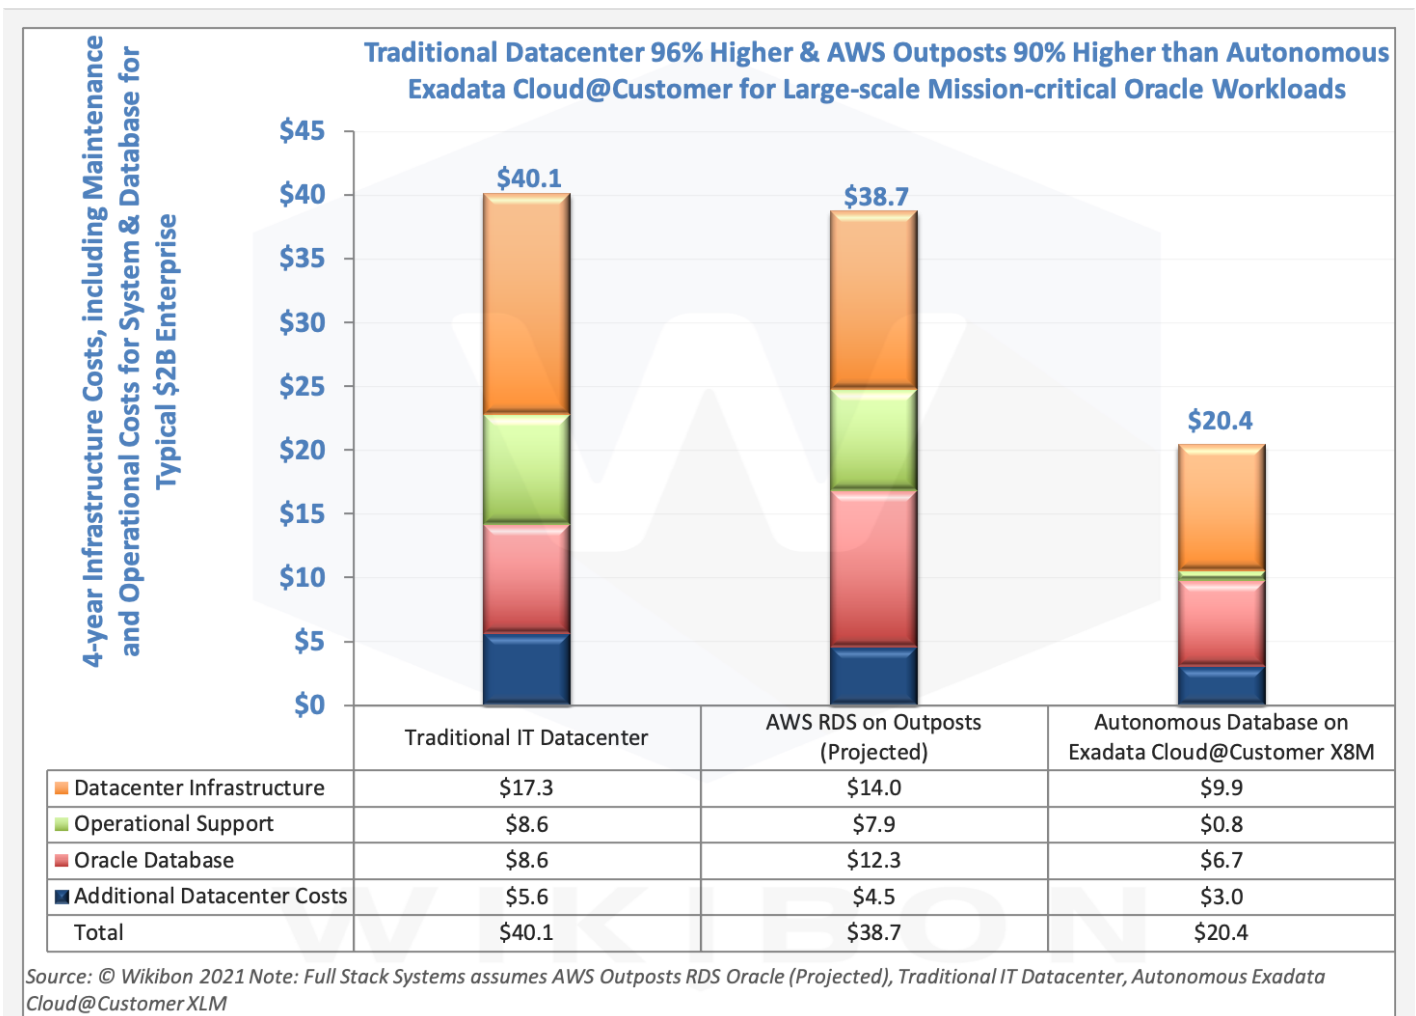
Figure 3 – Detailed IT Cost Comparison between Traditional IT Datacenter, AWS Outposts, and Oracle Cloud@Customer X8M when running Large-scale Mission-critical Applications on Oracle Automated Cloud Database.
Source: © Wikibon, 2021.

The components of the 4-year costs in Figure 3 include:

Infrastructure Costs
Server Costs

Traditional IT datacenters need to have the capacity to meet the highest peak-performance requirements. There is also no offload to the storage or network systems.

The AWS Outposts system offloads some IO, network, and security processing to a lower-cost ARM Nitro card.

The Oracle solution minimizes IO and system wait times to access data using 100-Gbits/Sec RDMA over Converged Ethernet directly to the storage server described in the "Oracle Exadata Cloud@Customer Architecture" section above. Exadata's Smart Scan technology also offloads SQL queries, analytics, and ML algorithms to storage servers, further increasing the amount of work that each database server CPU can process. This architecture reduces the CPU cores required for OLTP and data warehousing, allowing customers to enable fewer server cores for their given workloads and reduce Oracle Database license costs.

Infrastructure Software Costs

Infrastructure software includes operating systems, infrastructure platform software such as VMware, and software services to manage assets, resources, and data center processes. DCIM (Data Center Infrastructure Management) provides the ability to run efficient data center operations and improve data center infrastructure planning and design. Physical & logical security is also addressed.

This software is at full cost for Traditional IT.

AWS Outposts can reduce some of the local infrastructure software costs but still uses some infrastructure services in the AWS cloud.

The Oracle Autonomous Database solution includes the Oracle infrastructure software services available in Exadata Cloud@Customer X8M managed by Oracle, further reducing costs.

Storage Costs

The traditional IT datacenter costs assume high-performance flash storage (e.g., Dell PowerMax) and a storage area network (SAN).

The AWS costs assume its highest performance cloud storage and storage network at a lower cost than the traditional datacenter.

The Exadata X8M storage servers have Intel Optane persistent memory extension, providing a simple, fast, and cost-effective tier for writes and a cache for reads. This storage architecture provides the database and storage servers with extremely low overhead RDMA IO in about 20 microseconds. This design minimizes data movement from the storage back-end and reduces the number of SSDs and HDDs required, which reduces cost.

Network costs are similar for all the solutions.

Operational Support

Operational costs for system administration and database administration (DBA) are very high in the traditional IT datacenter. The system administration is usually split into a separate server, storage, and network groups. Many DBAs are required to manage the databases, including provisioning, patching, updating, cloning, and backing up/restoring the software.

The AWS RDS service automates updating the database and infrastructure services, saving people costs compared to traditional datacenters. However, DBAs are still required to manage the databases (e.g., ETL, cloning, scaling shapes, backup/restore, etc.) and interface with the developers.

Oracle Autonomous Database (for transactions and analytics, RAC, and Data Guard) automates many critical database management functions. For example, Autonomous Database's auto-tuning of database indexes are automatically managed and optimized, taking a manual process that can take years and turning it into an overnight job. Furthermore, Autonomous Database on Exadata Cloud@Customer provides auto-scaling, auto-patching, auto-securing, and other automated management capabilities while Oracle manages the infrastructure. This automation results in a dramatic lowering of operational support costs. DBAs can be redeployed to focus on innovation and creating business value. The feedback from initial users of Autonomous Database services has confirmed this level of reductions in operational costs. The users also reported that it takes time for the specialists to trust the automation!

Environmental Costs (On-premises Power and Space)

Traditional on-premises datacenters have the highest cost due to the requirement to meet capacity for the peaks, usually at month-end, quarter-end, and year-end.

The AWS Outposts and Exadata Cloud@Customer hybrid solutions allow some offloading of resources to the cloud and reduce on-premises environmental costs.

Oracle's approach of running the Autonomous Database on Exadata Cloud@Customer also allows customers to offload resources for data protection and disaster recovery to Oracle Cloud Infrastructure. In addition, Oracle Exadata Cloud@Customer users can also use local-only backups to meet strict data sovereignty and security requirements.

Oracle Cloud Database Costs

The database licensing costs are higher for the traditional IT datacenters as, again, the licenses have to be in place for the peak loads.

The database licensing costs for AWS are much higher due to the lack of agreement between AWS and Oracle on virtualized cores. This architecture leads to an effective doubling of Oracle Database licensing costs on AWS RDS.

The Oracle Database costs are much lower for Exadata Cloud@Customer. The much-improved system speed lowers the system wait times, which reduces the number of cores required and the amount of time needed to complete a computation, reducing overall costs. As a serverless architecture, the Oracle solution automatically scales to match changing workloads, providing true pay-per-use. Also, Exadata Cloud@Customer offers enterprises ways to increase and decrease the number of Oracle Database licenses they use without interrupting database operations. Running Autonomous Database on Exadata Cloud@Customer X8M improves this core capability by fully automating the increases, decreasing the licenses required, and doing so based on the queries currently being run instead of looking in a rear-view mirror. In doing so, Oracle allows enterprises to optimize performance and cost controls for their database infrastructure automatically.

Bottom Line: the cost of a traditional IT datacenter is 96% higher, and AWS Oracle RDS hybrid cloud solutions are 90% higher than Oracle Autonomous Database on Exadata Cloud@Customer X8M.

Adjacent Cloud Database Systems from Microsoft and Oracle

There is an alternative solution for enterprises running mission-critical Oracle Databases on Microsoft Azure. Enterprises can use database calls to an Oracle Exadata Cloud@Customer X8M deployment running adjacent to an Azure Stack in the same public cloud datacenter. Microsoft and Oracle have tightly integrated the two services and integrated support.

The advantage of this approach for enterprise customers is in avoiding a lengthy, costly, and risky conversion of the Oracle Database applications to another database platform. Wikibon has emphasized, on many occasions, the importance of avoiding database conversions, as this can lead to significant and costly delays in any digital transformation initiative. Wikibon believes that this Adjacent Database capability was a substantial reason for the JEDI contract award to Microsoft. The savings from avoiding the conversion costs from Oracle to AWS databases and the lack of AWS support for Oracle Database systems were significant factors.

This multi-cloud approach requires enterprise IT to integrate the two services, which can add cost. However, these costs are low compared with conversion costs.

Conclusions and Recommendations

Cloud Database Conclusions

The findings from this report are stark:

The costs for running mission-critical Oracle Database workloads in traditional datacenters and RDS on AWS Outposts are nearly 2X higher than Oracle Autonomous Database on Exadata Cloud@Customer X8M. Traditional data center architectures have a 96% higher TCO than the Oracle Autonomous Database on Exadata Cloud@Customer X8M solution evaluated in this study. Wikibon believes it is time to replace the best-of-breed DIY design approach for Oracle mission-critical infrastructure held together with skillful and expensive in-house expertise. Giving the design and automation responsibilities to the database vendor allows economies of scale in developing optimal hardware, software, and autonomous features. This capability cannot be replicated in-house, even by the largest Oracle shops.

The hybrid AWS Outposts solution shows a 90% higher TCO than Oracle for the workloads analyzed. This analysis is a Wikibon projection of expected future announcements of RDS on Outposts by AWS and will be updated, if necessary, after any future AWS announcements.

Wikibon concludes that the current AWS database strategy of specialized, isolated databases with inadequate data sharing will inhibit the development of synchronous applications and the fundamental business improvements that can ensure enterprise survival and prosperity.

Microsoft Azure/Oracle Exadata Cloud Adjacent strategy is an alternative way of running Oracle Database applications in the Microsoft cloud. Although there is some support overhead, the joint approach avoids database conversion, saves significant time and expense, and lowers business risk.

Wikibon concludes that mission-critical, high-performance Oracle Database applications will benefit from the specialized hardware and software in Exadata X8M designed to support Autonomous Cloud Database. The benefits include lower costs, higher performance, availability, and radically improved serviceability, while Oracle manages an increasing level of infrastructure and database settings.

Wikibon concludes that future complex synchronous workloads will need database technology that allows sharing of different data types across transactional, analytic, AI, blockchain, and other database types and automation to meet too tight real-time response times. This converged approach is available now and is a critical strategic development tenet for future Oracle Autonomous Cloud Database development.

The Wikibon analysis shows that Oracle's superior cost performance comes from its fully integrated software stack and hardware specialized for Oracle Databases. The autonomous capabilities are built from decades of database experience. Wikibon concludes that it is now impossible to replicate these cost benefits with either a traditional datacenter approach or AWS services.

Cloud Database Recommendations

Wikibon recommends that enterprises running large-scale mission-critical and demanding database workloads consider architecting around the Oracle Exadata Cloud@Customer X8M for on-premises cloud services and Oracle Exadata Cloud Service on Oracle Cloud Infrastructure (OCI). As the analysis in this report and previous reports indicates, the result will likely lower IT budget costs relative to alternative solutions.

Wikibon believes that the AWS strategy of multiple independent database types will not allow enterprises the levels of data integration required to achieve high levels of synchronous business processes. Wikibon recommends that AWS must invest heavily to develop high-availability versions that are supported as an

integrated product by AWS, and not just provide the piece parts for enterprise IT to assemble, test, and maintain. Also, Wikibon recommends that AWS invest heavily in developing an integrated database strategy, with buying Couchbase as a possible first step.

Cloud Database Action Items

Oracle Cloud Database is Tier-1 and in a class of its own. Wikibon recommends that larger enterprises with mission-critical workloads should not convert from Oracle to other databases. Instead, Wikibon recommends migrating to Autonomous Cloud Database on Oracle Exadata Cloud@Customer X8M, Oracle Exadata Cloud Service on OCI, or other Oracle Database cloud services. Wikibon recommends that enterprises minimize the number of separate databases and data types and use the converged Oracle Cloud Database instead.

At this time, Wikibon cannot recommend running large-scale Oracle Database mission-critical workloads and the surrounding portfolio of applications on AWS. The cost of running Oracle databases in AWS is prohibitive.

Wikibon recommends Microsoft as the best multi-cloud alternative for Oracle mission-critical workloads because of its adjacent Microsoft Azure strategy combined with Oracle Exadata Cloud technology.

Senior executives should press Oracle and AWS to bury the hatchet and develop a win-win-win cost-effective multi-cloud database services strategy for their joint customers.

Footnotes

Application Support for Synchronous & Asynchronous Business Processes

Enterprises use synchronous transactional applications to support parts of their business processes. For example, ERP systems synchronously update order records, inventory levels, and many other fields. However, there are many parts of the business processes that are supported asynchronously. For example, executing price adjustments due to orders (or lack of orders) is usually (Uber & Lyft excepted) not synchronous. The result is sub-optimal revenue.

One primary IT reason for making the price update process asynchronous is that the analysis to determine the update requires a large amount of historical and other data, which historically took too long to find and process. Many enterprises could increase revenue significantly if the price updates were synchronous instead of asynchronous.

Today's integrated system architectures can now process both transactional data and analytic data to determine price-adjustments synchronously, that is, in real-time. In general, making asynchronous business synchronous can significantly reduce employee costs and improve enterprise responsiveness (i.e., improve business cycle times).

Technical requirements for the development of synchronous applications include integrated hardware and a complete software stack, ultra-low latency to data, in-memory databases, multiple databases sharing the same data, parallel processing of data, advanced machine learning inference functionality, and others.

*** Amazon Statement on Oracle Support for Outposts:*

"Amazon RDS on Outposts supports MySQL and PostgreSQL database engines, with support for additional database engines coming soon."

David Floyer spent more than 20 years at IBM, holding positions in research, sales, marketing, systems analysis and running IT operations for IBM France. He worked directly with IBM's largest European customers, including BMW, Credit Suisse, Deutsche Bank and Lloyd's Bank. Floyer was a Research Vice President at International Data Corporation (IDC) and is a recognized expert in IT strategy, economic value justification, systems architecture, performance, clustering and systems software.

David Floyer

@dfloyer
david.floyer@wikibon.org