



Real-time Analytics Revolutionize Data Warehousing by, David Floyer

May 3rd, 2022

The premise of this research is that modern integrated hardware and converged database software will enable Real-time Analytics. Technologies such as Oracle Autonomous Data Warehouse with Exadata will revolutionize data warehousing by progressively eliminating complex ETL processes and providing the business with analytic queries in real-time. As a result, enterprises can radically improve the quality of data, simplify data management, and move towards the goal of a single version of the truth. No significant increase in IT budget is required to meet this goal.

Real-time Analytics is also a prerequisite for Transactional Analytics, where the transaction applications issue Real-time Analytic queries. Wikibon believes that Transactional Analytics will profoundly simplify and automate business processes and dramatically improve costs, quality, and customer satisfaction.

Wikibon strongly advises Executive Management to embrace Real-time Analytics, set a 5-year goal of eliminating at least 90% of all ETL operations, and embrace the vision of a single version of the truth with provenance and compliance. In addition, Wikibon recommends IT development start integrating Real-time Analytics with existing mission-critical Systems of Record.

Wikibon recommends implementing Real-time Analytics on Oracle Autonomous Data Warehouse on Exadata Cloud@Customer or Cloud Infrastructure X9M because only these platforms have the converged database architecture, performance, price performance, and level of investment to be effective converged real-time database platforms.

Premise

The premise of this research is that modern integrated hardware and converged database software will enable Real-time Analytics. Technologies such as Oracle Autonomous Data Warehouse with Exadata will revolutionize data warehousing by progressively eliminating complex ETL (Extract, Transform, Load) processes and providing the business with analytic queries in real-time. As a result, enterprises can radically improve the quality of data, simplify data management, and move towards the goal of a single version of the truth. No significant increase in IT budget is required to meet this goal.

Real-time Analytics is also a prerequisite for Transactional Analytics, where the transaction applications issue Real-time Analytic queries. Wikibon believes that Transactional Analytics will profoundly simplify and automate business processes and dramatically improve costs, quality, and customer satisfaction. Wikibon will publish follow-on research about the value of Transactional Analytics and the prerequisite technologies.

Research Scope

This Wikibon research paper defines Real-time Analytics and the crucial benefits of reducing and eventually eliminating complex, time-consuming, and costly ETL processes to populate data warehouses. In addition, we look at the integrated infrastructure and software technology requirements to achieve Real-time Analytics.

Wikibon then evaluates a case study where an enterprise runs a data warehouse based on Oracle Database Enterprise Edition on-premises in a traditional datacenter with “Best-of-Breed” infrastructure components. We assume that the organization has annual revenues between two and three billion dollars. Next, Wikibon evaluates the performance, IT costs, and business case to migrate and deploy two different cloud-based data warehouse solutions that can support Real-time Analytics. These alternatives are Oracle Autonomous Data Cloud Warehouse with Exadata Cloud Infrastructure X9M or Exadata Cloud@Customer X9M. Wikibon also evaluates AWS RDS for Oracle solutions for performance and IT costs and the ability to deliver Real-time Analytics.

We also evaluate additional major alternative cloud database vendors, including AWS, Couchbase, IBM, Microsoft, SAP, and Snowflake, and their ability to support Real-time Analytics.

A follow-on Wikibon paper will look at the practical steps in designing and deploying Transactional Analytics systems on Oracle Database and the challenges of achieving full automation of business processes. Finally, we will examine the potential business case to survive and thrive and why we expect successful implementation to have staggering returns.

Traditional Analytics

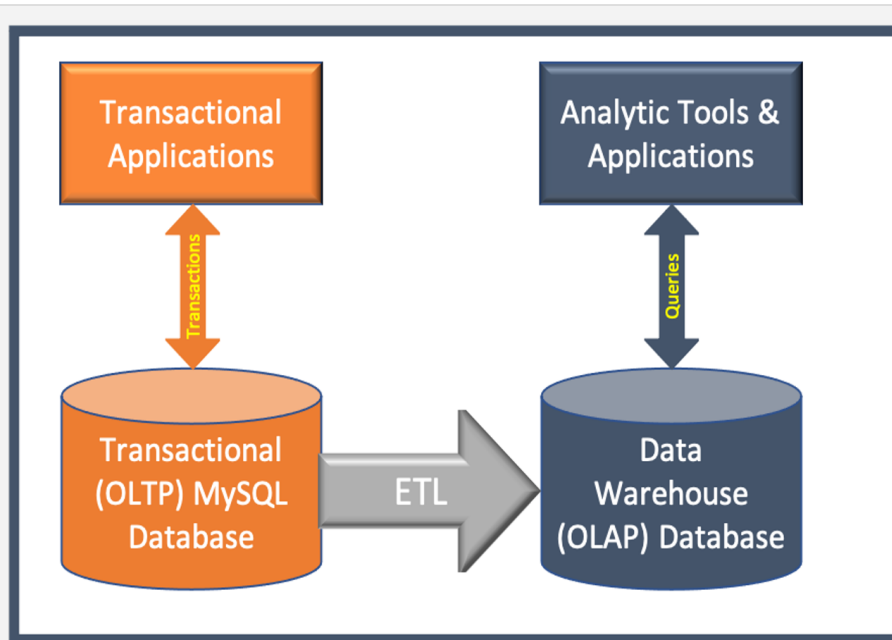


Figure 1 - Extract Transform Load (ETL) is a Time-consuming Circuitous Process that Separates Transactional & Data Warehouse Data and Databases. ETL is not compatible with Real-time Analytics

Source: © Wikibon 2022

Figure 1 to the left illustrates the separation between transactional and analytic databases. The transactional database supports the creation of data, which must move to the data warehouse through an ETL process (Extract, Transform, Load) to the separate data warehouse database. The process is much more circuitous and time-consuming than Figure 1 suggests. Data flows into data warehouses from many sources. Specialized data warehousing software ETLs this data into data warehouse databases, with the same data appearing hours, days, or weeks later. The multiple copies of data and variable timing make data warehouses extremely complex. The bottom line is that executives and operational staff cannot rely on data accuracy. As a result, enterprises have had to design complex business processes to deal with data arriving inconsistently late.

Automation of business processes is complicated and challenging in this environment. Although companies like UiPath have improved productivity sharply with RPA (Robotic Process Automation) tools, the final steps to fully automate and simplify business processes are elusive. For example, if the organization is not sure a transaction is fraudulent, it must be able to undo the transaction later, requiring advanced people skills when the data is inaccurate.

The author has yet to meet a senior executive who is happy with their traditional data warehouse.

Real-time Analytics

Real-time analytics allows enterprises to analyze data at the time of creation and extract the value from the data in real-time. As a result, all data in the data warehouse has an accurate timestamp and provenance, which dramatically simplifies storing, understanding, and retrieving this data. This capability means enterprises can streamline and automate their analytic processes, dramatically improve data quality, and streamline data management and compliance.

The Technology Requirements of Real-time Analytics

Transactional vs. Analytic Databases

Designers of transactional systems usually organize data in database rows — for example, the application stores the customer, product, and delivery information in a row. Their transactional databases typically use 3NF (Third Normal Form) data models.

However, analytics typically comprise columns of data and analyze customer, product, and financial data.

Real-time Analytics Revolutionize Data Warehousing

Analytic (Data Warehouse) databases often use a dimensional data model (facts and dimensions).

Enterprise IT often feeds their data warehouses from multiple source databases, which is another reason why they have separate OLTP and data warehouse databases. Some enterprises have also grown through acquisition, so they might have different OLTP systems that perform the same business function but for formerly separate parts of the business. In those cases, the users often need a consolidated view of the business, so the data is loaded into a single, common data warehouse.

As a result, enterprises usually use separate databases for transactional and analytic systems. Data is generally transferred from transactional databases to data warehouses using ETL, as in Figure 1 above. This process is a source of much unhappiness in IT.

Converged transactional and analytic databases are a potential solution, as discussed in the next section.

Converged Databases

The last decade has seen many additional applications use different data types (e.g., document, location, graph data, and many more). As a result, vendors introduced many specialized databases (e.g., Mongo for document databases). Unfortunately, document databases are not well suited to analytics, as the basic structure of the data (how it's modeled) doesn't work well for performing analytics against that data. A separate document database means you need to ETL that data somewhere else and transform it for querying and analytics.

In addition, blockchain and Machine Learning (ML) are recent innovations that require specialized software and databases. All these separate databases make the ETL processes even more complicated, increase the time-complexity of the data warehouse, and decrease the accuracy and value of the data.

Converged databases are the basis for solving this problem. For example, the Oracle Database allows developers to use all the data types discussed above in a single converged database. In addition, Oracle GoldenGate software allows real-time sharing between databases. Finally, some real-time transformations may be necessary for databases such as the document databases mentioned above, and having a converged database simplifies these operations.

Converged databases and good data modeling are essential to achieving the benefits of Real-time Analytics and Transactional Analytic applications. However, specialized hardware and software are also required to perform real-time analytics at scale, which is the subject of the next section.

Specialized Database Hardware & Software

Databases make application development much faster, more reliable, and safer. Converged databases also dramatically improve application development speed and quality. Commodity hardware and software are fine for running smaller databases and lower transaction rates. Microservices can link processes together if the databases are small, and real-time is not required.

However, integrated applications with more extensive data volumes, higher transactional rates, and very high availability need specialized hardware and software. For example, the ACID (Atomicity, Consistency, Isolation, Durability) properties of transactional databases require highly specialized software and low-latency hardware at scale. In addition, converged databases need to combine access and management of rows and columns and need specialized memory hardware and software to achieve this. Furthermore, large-scale mission-critical databases require scalability and recoverability built-in. Finally, cloud databases need even higher scalability and performance.

Amazon and other cloud vendors advocate small databases linked by microservice links because their databases and infrastructure cannot handle scale. The same goes for legacy on-premises best-of-breed hardware vendors. However, Wikibon believes that almost all mid-sized and large customers will need largescale databases to achieve the transformational simplicities of Real-time Analytics and Transactional

Real-time Analytics Revolutionize Data Warehousing

Analytics.

Oracle is the largest database vendor and has developed highly specialized hardware called Exadata to run it efficiently. The list below picks out elements of the Exadata architecture that contribute to performance and latency improvements to enable Real-time Analytics. Readers who do not require the detail below can skip to the next section.

- Exadata X9M includes high-bandwidth, low-latency active-active 100 Gb/sec network fabric connections, which use a low latency protocol called RDMA over Converged Ethernet (RoCE).
- With Exadata Persistent Memory Data Accelerator, Oracle Database accesses data cached in Persistent Memory using Remote Direct Memory Access (RDMA) for the lowest possible latency of less than 19µsec.
- Exadata Persistent Memory Commit Accelerator automatically enables the database to issue a one-way RDMA log write to Persistent Memory. RDMA and Persistent Memory technologies allow log writes to occur without the typical request/response acknowledgment, and intelligent software distributes the write across multiple servers for resilience. This feature leads to a significant performance increase for log write operations.
- Each storage server has 1.5 terabytes of Optane Persistent Memory as a high-speed caching tier between DRAM and flash. This Persistent Memory improves performance.
- Exadata Persistent Memory Data and Commit Accelerators bypass the network and I/O stack and eliminate expensive CPU interrupts and context switches. As a result, latency reduces by 10x, from 200µs to less than 19µs.
- For transactional environments, a single rack configuration of Exadata X9M-8 can achieve up to 15 million random 8K database reads and nearly 7 million random 8K flash write I/O operations per second (IOPS).

The Exadata X9M flash storage servers are capable of scanning one terabyte per second. Data warehouse environments require high bandwidth performance.




Figure 2 - The Exadata X9M-8 has State-of-the-Art Communication enabling High Bandwidth and Low-latency with Racks and between Racks. Exadata is a key enabler for Real-time Analytics

Source: Oracle, 2022

Real-time Analytics Revolutionize Data Warehousing

- The Exadata SmartScan technology offloads data-intensive SQL operations from the database servers to all the storage servers. As a result, SQL data filtering and processing occur in parallel while data is read from disk and flash. Only rows and columns relevant to a query transfer to the database servers.
- The Exafusion Direct-to-Wire Protocol allows database processes to read and send Oracle Real Applications Cluster (RAC) messages directly using RDMA and bypass the OS kernel and networking software overheads. This feature improves the response time and scalability of RAC OLTP configurations on Oracle Exadata Database Machine, especially for workloads with high-contention updates.
- Exadata uses Machine Learning (ML) to implement automatic indexing. The ML continually analyzes executing SQL and creates new indexes to accelerate performance. In addition, ML continuously learns and tunes the database as the underlying data model or usage patterns change.
- Exadata implements real-time statistics gathering as DML operations insert, update, or delete data. This function ensures that the SQL optimizer continuously updates statistics to make more intelligent decisions on SQL execution paths. In addition, real-time statistics allows the SQL optimizer to adapt plans dynamically as the distribution of data changes.
- Exadata implements the same pure columnar format in FlashCache used by Oracle Database In-Memory. This feature extends the Exadata Columnar Flash Cache by automatically transforming data into In-Memory pure columnar formats when loading into FlashCache. Smart Scans also process multiple column values with a single instruction by leveraging Single Instruction Multiple Data (SIMD) vector instructions. Smart Scan results are passed back to the database server in Oracle Database In-Memory formats, further reducing the load on the database servers. The effect is to seamlessly extend the In-Memory columnar store size from DRAM in the database server to flash in storage servers. An Exadata X9M-8 Full Rack HC5 has 360 TB of Flash Cache, capable of servicing the largest in-memory workloads.
- Exadata In-Memory Fault Tolerance duplicates in-memory data across the clustered database servers to provide higher availability. Queries will transparently access the duplicate copy on a surviving database server if a database server fails and continue without interruption.
- Exadata integrates with Active Data Guard to allow In-Memory analytics on a standby database.
- Exadata uses Transparent Data Encryption (TDE) to move decryption processing from database servers into the Exadata Storage Server hardware. Exadata storage leverages hardware decryption and compression to provide high-performance, secure databases. Encryption occurs after the data compression so that the cost of decryption decreases by the degree of compression. As a result, Exadata can query fully encrypted and compressed databases with minimal overhead at hundreds of gigabytes of original user data per second by leveraging both technologies. Oracle Transparent Data Encryption provides a key management solution to keep all data encrypted and secure.
- Exadata is preconfigured and pre-tested and is identical whether on-site or in the cloud. The volume of identical Exadata systems installed improves the overall reliability of the hardware and the support services.
- The Exadata system optimizes performance for all database applications and enables real-time analytics in the real world.

Oracle Converged Database and Exadata Performance Reduces or Eliminates Need for ETL

Real-time Analytics Revolutionize Data Warehousing

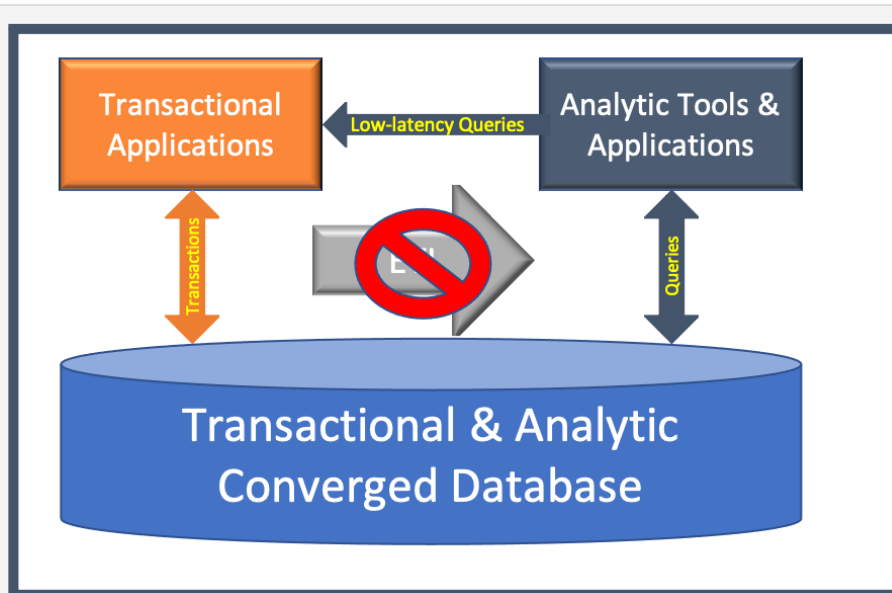


Figure 3 - The Oracle Converged Enterprise Database, ADW in-memory Technology, and Exadata database performance mean that IT can drastically reduce or eliminate ETL & enable Real-time Analytics.

Source: © Wikibon, 2022

Figure 3 to the left shows the essential business conclusion from this section. The speed of the Exadata, which can accelerate both transactional and data warehouse workloads together with an integrated database and in-memory features, allows transactional workloads to run about 3-4 times faster.

The converged database allows the analytics to run on the same database as the transactions. As a result, the data is always consistent and up to date. Users do not have to wait for ETL processes to complete. Queries are simpler to code and debug.

Like everything in life, compute resources are finite, and IT will need to shield sufficient resources for transaction processing. However, new ways to query data in real-time or near real-time will drive improvements in business processes.

Oracle Autonomous Data Warehouse for Analytics and Data Warehousing (ADW)

Oracle Autonomous Data Warehouse (ADW) is an automated cloud database service optimized for analytic workloads. It runs on Exadata Cloud Infrastructure and Exadata Cloud@Customer and contains preconfigured Oracle Database analytical features such as columnar format, partitioning, and large joins.

ADW moves the responsibilities for automating operational procedures from the enterprise data center to Oracle OCI operators. The theoretical result is that higher implementation volumes should make Oracle more efficient. Users have reported more efficient database administrators, infrastructure management, and lower infrastructure costs. As expected, users find radical improvements in planned and unplanned downtime. In addition, the application development and analytic data groups have a better system experience and spend less time in systems administration.

The Cost Analysis Section ahead includes the benefits of ADW and Exadata listed above.

Performance & Cost Analysis of Oracle Autonomous Data Warehouse

Performance Analysis Oracle Database-only Cloud Platforms

Figure 4 shows the compute power required to drive different Data Warehouse Infrastructure Platforms. The left-hand axis is the performance measured in virtual CPUs (vCPUs). The blue column on the left represents


Real-time Analytics Revolutionize Data Warehousing

the starting point for the number of vCPUs in an on-premises "Best-of-Breed" DIY configuration of enterprise servers, storage, and networking running Oracle Data Warehouse software in a virtualized environment. The data in Figure 4 below come from Table 2 in the *Footnotes* section, and the model assumes a 25% saving from virtual sharing of vCPUs for the blue column.

The second orange column from the left is the number of vCPUs in the AWS RDS for Oracle service in a Single AZ (Availability Zone) environment. Table 2 also shows the selection of larger instances to meet the vCPU requirement.


The third column from the left is the AWS RDS for Oracle service in a Multi-AZ environment. The number of vCPUs is the same as the Single AZ, but the cost is higher.

The red column on the right shows the number of vCPUs required in the Oracle Autonomous Data Warehouse on Exadata Cloud Infrastructure X9M or Cloud@Customer with Exadata X9M. Table 2 shows that the specialized hardware and software detailed above only needed 65% of the vCPUs to do the same work. In addition, the base number of vCPUs that must be allocated is set at 45%, and the average burst vCPUs is 28%. As a result of these settings, the net number of vCPUs required 316, 25% of the AWS requirements because Oracle built its cloud for database efficiency.



Performance Analysis of Oracle Databases, MySQL, & Snowflake Cloud Platforms

Figure 5 below is Figure 4 with an extra green column, which represents an initial conservative estimate of the performance of other data warehouse offerings. These offerings include converged MySQL from AWS Aurora and Couchbase, and Snowflake running on AWS. Snowflake is a Data Warehouse-only platform that cannot share a common transaction database and, therefore, cannot support Real-time Analytics. Wikibon has compared Heatwave performance with converged MySQL & Snowflake Data Warehouse benchmarks and other Oracle Database Benchmarks. As a result, Wikibon is comfortable that the performance requirements of these workloads are likely to exceed 1,500 vCPUs.




IT Budget Cost Analysis

Figure 6 below shows the Wikibon Research findings of detailed 5-year IT Budget estimations for migration from an on-premises data center to a cloud database approach. The starting point is an on-premises "Best-of-Breed" datacenter, the light-blue column on the left, and the 5-year IT budget for staying put is \$48 million. The origin of the data is Wikibon's IT and Budget model for an enterprise with revenue or budgets of \$2-3 billion/year.

Real-time Analytics Revolutionize Data Warehousing

AWS Single-AZ RDS for Oracle requires 2.9 times more IT Hardware Budget than Oracle ADW w Exadata
AWS Multi-AZ RDS for Oracle requires 3.4 times more IT Hardware Budget than Oracle ADW w Exadata
On-Premises "Best-of-Breed" requires 3.1 times more IT Hardware Budget than Oracle ADW w Exadata X9M
Much better Virtual Sharing, Ultra-low IO latency, IO Offloading, Reduced ETL, & Autonomous DW with Exadata X9M Architecture enable Dramatic Cost Reductions



Source: © Wikibon, 2022

Figure 6 - 5-year IT Budget for Migration from On-premises "Best-of-Breed" Oracle Data Warehouse to AWS RDS for Oracle Cloud Platforms & Oracle Cloud@Customer and Cloud Infrastructure Platforms (\$M)

Source: © Wikibon 2022

Figure 7 below shows a detailed breakdown of IT budget line items. The three main items are compute and storage for 933 vCPUs (See Figure 4 above, left-hand column), Oracle Database maintenance, and operational support and application maintenance. The Oracle Database features include the maintenance on Enterprise Edition, Active Data Guard, and RAC. This model assumes that IT has bought sufficient licenses in the past and only needs to pay maintenance.


The database maintenance value for Oracle Enterprise Edition increases with the inclusion of the in-memory feature, which is required to move towards real-time analytics. Both the AWS and Oracle clouds include the additional maintenance cost of the feature. The number of licenses is a function of the number of vCPU cores required to drive the workload.

The two orange columns reflect the IT Budgets for migrating to AWS RDS for Oracle cloud service offering in either a Multi-AZ configuration (compute and storage budgets are about double) or a Single-AZ configuration. In addition, the database budgets include the in-memory feature. Operational and application maintenance is the third large item, which reduces because of some automation of updates with the RDS service.

Real-time Analytics Revolutionize Data Warehousing

Total Cost of Running Oracle Data Warehousing
Databases On-premises Best-of-Breed,
AWS RDS Cloud, & Oracle Database Cloud (\$M)

Oracle Autonomous Data Warehouse on Exadata Cloud Infrastructure X9M has \$29.5M lower 5-year IT Budget than AWS RDS for Oracle Single-AZ, \$32.8M lower than On-premises, & \$36.5M lower than AWS Multi-AZ



Source: © Wikibon 2022, based on detailed expected IT budgets for 5-year Migration Projects (\$M)

Figure 7- 5-year IT Budget Line Items for Migration from On-premises “Best-of-Breed” Oracle Data Warehouse to AWS RDS for Oracle Cloud Platforms & Oracle Autonomous Data Warehouse Exadata Cloud@Customer and Cloud Infrastructure X9M Platforms (\$M)

Source: © Wikibon 2022

Oracle Cloud@Customer and Cloud Infrastructure enable a dramatic decrease in the number of vCPUs required. As a result, Figure 7 shows there is a dramatic reduction in the hardware budget. In addition, the Oracle Database license costs are cut for the same reason, and the automation of the administrative and technical tasks with the Autonomous Data Warehouse features accounts for the reduction in staffing costs.

Table 3 in the Footnotes section below details the build-up of charges shown in Figure 7 over five years.

IT Budget Financial Metrics

Table 1 below shows the detailed IT budget financial metric calculations between the current state and alternative migration paths. The top half of the table is a 5-year summary of the costs by IT Budget Item, derived from Table 3 in the Footnotes.

The bottom half of Table 1 shows the financial analysis metrics for the migration from the on-premises current environment for data warehousing to a cloud-based AWS RDS for Oracle or Oracle Cloud solution. The first observation is that AWS RDS for Oracle Multi-AZ solution does not show any return compared with the current on-premises environment, with the expected downtime costs used in the model. However, the

Real-time Analytics Revolutionize Data Warehousing

AWS Multi-AZ solution can break even in environments where expected downtime costs are significantly higher.

The first line shows the 5-year net benefits. Both Oracle solutions save more than ten times the cost of the AWS RDS for Oracle Single-AZ (\$3.3M vs. \$31M or \$32.8M). The NPV (Net Present Value) line, which considers the cost of money at 4%, also shows about ten times the benefit (\$2.7M vs. \$26.9M or 28.6M).

The breakeven line shows a breakeven for the AWS Multi-AZ line of 26 months vs. 6 or 7 for the Oracle Cloud solutions.

The IRR (Internal Rate of Return) shows a reasonable IRR of 78% for the AWS Single-AZ solution. The Oracle Cloud solutions are much higher (IRR 316% or 514%). The traditional ROI measurements are in the last line (NPV divided by software migration costs and additional data center costs). The results compared with IRR are about the same for AWS Single-AZ (69%) and even higher for Oracle Cloud solutions (not that it matters).

Overall, the IT budget business case from a CFO point of view is non-existent for AWS Multi-AZ and marginal for AWS Single-AZ.

The Financial Case from On-premises to Oracle ADW on Exadata Cloud@Customer X9M is Strong (BE 7 months, NPV \$26.9M, IRR 316%).

The case for Migration to Oracle ADW on Exadata Cloud Infrastructure X9M is also strong (BE 6 months, NPV \$28.6M, IRR 514%).

The Financial Case for Migration to AWS RDS for Oracle Single-AZ is poor (BE 26 months, NPV \$2.7M, IRR 78%).

There is no Financial Case for Migration to AWS RDS for Oracle Multi-AZ.

Only the Oracle Solutions can support Real-time Analytics & Transactional Analytics.

	On-premises "Best-of-Breed" Data Warehousing	AWS Multi-AZ RDS for Oracle Data Warehousing	AWS Single-AZ RDS for Oracle Data Warehousing	Oracle Autonomous Data Warehouse on Exadata Cloud@Customer X9M	Oracle Autonomous Data Warehouse on Exadata Cloud Infrastructure X9M
5-year IT Budget Line Items (\$M)					
Compute & Storage	\$ 14.1	\$ 18.2	\$ 10.1	\$ 3.3	\$ 3.2
ETL Costs	\$ 2.1	\$ 3.6	\$ 3.6	\$ 1.1	\$ 1.1
Database Maintenance	\$ 15.2	\$ 18.5	\$ 18.5	\$ 5.3	\$ 5.3
Operational Support & Application Maintenance	\$ 7.9	\$ 4.7	\$ 4.6	\$ 2.1	\$ 2.1
Environmentals	\$ 0.5	\$ -	\$ -	\$ 2.7	\$ 1.8
Other Datacenter Costs	\$ 4.8	\$ 3.4	\$ 3.4	\$ -	\$ -
Software Migration Costs	\$ -	\$ 0.5	\$ 0.5	\$ 0.3	\$ 0.3
Expected Downtime Costs	\$ 3.4	\$ 3.0	\$ 4.2	\$ 2.2	\$ 1.4
Total	\$ 48.1	\$ 51.9	\$ 44.9	\$ 17.2	\$ 15.4
Financial Analysis					
5-year Net Benefit (\$M)		\$ (3.8)	\$ 3.3	\$ 31.0	\$ 32.8
Net Present Value (NPV, 4%, \$M)		\$ (3.4)	\$ 2.7	\$ 26.9	\$ 28.6
Breakeven (BE, months)		60	26	7	6
Internal Rate of Return (IRR)			78%	316%	514%
ROI (NPV(4%) ÷ Migration Costs)		-66%	69%	893%	1329%

Source: © Wikibon, 2022

Table 1 - IT Budget Financial Analysis of Alternative Solutions to Migrate from an On-premises "Best-of-Breed" Data Warehouse to a Cloud-based System in Preparation for Real-time Analytics & Transactional Analytics. Oracle Autonomous Data Warehousing enables Real-time Analytics without Significant Increase in IT Budget.

Source: Wikibon 2022

Figure 8 and Figure 9 below shows clearly why Wikibon believes every CFO would endorse the Oracle Autonomous Data Warehouse with Exadata Cloud@Customer X9M solution and Oracle Autonomous Data Warehouse with Exadata Cloud Infrastructure X9M. Both cases are low-risk, high-reward decisions. Figure 8

Real-time Analytics Revolutionize Data Warehousing

shows a break-even of 6-7 months with benefits of over \$27 million, and Figure 9 shows an IRR of over 300%.

Wikibon believes that by the most important reason to choose the Oracle Cloud platform is because it is unique in supporting Real-time Analytics at scale now, and enables a clear path to Transactional Analytics in the future.





Figure 8 - 5-year Net Benefit & Net Present Value for Migration from On-premises to Oracle Data Warehouse Cloud Platforms

Source: Wikibon 2022

Oracle Autonomous Data Warehouse has the Best Business Case by Far for Migration from On-premises to Exadata Cloud@Customer or Cloud Infrastructure X9M



Source: © Wikibon, 2022

Figure 9 - 5-year Net Benefit & Net Present Value for Migration from On-premises to Oracle Data Warehouse Cloud Platforms

Source: Wikibon 2022

Bottom Line:

- The Migration Financial Case from On-premises to Oracle ADW on Exadata Cloud@Customer X9M is Strong (BE 7 months, NPV \$26.9M, IRR 316%).
- Oracle ADW on OCI is also strong (BE 6 months, NPV \$28.6M, IRR 514%).
- The Financial Case for Migration to AWS RDS for Oracle Single-AZ is poor (BE 26 months, NPV \$2.7M, IRR 78%).
- There is no Financial Case for Migration to AWS RDS for Oracle Multi-AZ.
- Only Oracle Cloud Infrastructure Solutions can support Real-time Analytics and a path to Transactional Analytics.
- Executive management should choose between the Exadata Cloud@Customer Exadata Cloud Infrastructure options based on the best place to locate the data, considering performance and compliance factors.

Alternative Cloud Databases

Wikibon assesses that Oracle Cloud Infrastructure Solutions are the only platforms that currently have the functionality and performance to provide a clear path to Real-time Analytics and Transactional Analytics.

However, there are many alternative database vendors. The closest are listed below in order of Wikibon's assessment ability to provide such a path in the future:

1. Oracle MySQL HeatWave

Oracle MySQL HeatWave is by far the best current alternative and complement. Oracle announced HeatWave as a MySQL cloud-based service in 2020, with the extreme performance of MySQL analytics. The performance and price-performance are outstanding. However, MySQL does not currently have Oracle's large-scale and mission-critical capabilities of Exadata, the advanced features of Oracle Enterprise Database, or the Autonomous Data Warehouse level of automation.

Bottom Line: Oracle MySQL Heatwave can complement a strategy using GoldenGate to maintain consistency and recoverability for smaller distributed databases. Oracle has announced the availability of Oracle HeatWave on the AWS IaaS platform, and has a strong development team.

2. Microsoft SQL Server

Microsoft SQL Server is a Tier-1 database that runs on Microsoft Azure. In addition, enterprise IT can run Oracle Database on Azure Infrastructure and connect to Oracle Databases running on OCI. All services can be helpful as a part of a real-time analytic strategy for large customers.

Bottom Line: As a stand-alone platform, Microsoft SQL Server running on Azure does not have the depth of converged databases and the specialized hardware required to run Real-time Analytics and implement Transactional Analytic Software. However, Microsoft has both the development skills and customer base to introduce competing cloud-based products. Wikibon assesses that Microsoft will need 3-5 years to respond.

3. AWS RDS for Oracle - Single-AZ & Multi-AZ

Amazon Web Services provides a Relational Database Service (RDS) for Oracle Enterprise Edition. This cloud offering runs the industry's best database-as-a-service and offers automation of updates. It also runs the full Oracle Enterprise Edition software with support for a single Availability Zone (Single-AZ) and Multi-AZ. The latter synchronizes data between Availability Zones under the covers to provide the service and is twice the price of a Single-AZ.

AWS has improved its large servers with additional cores and memory, which helps process Oracle analytics. In addition, the Flash SSD support has significantly improved. The system runs well for smaller Oracle Database configurations that are not stressed. The AWS architecture is suitable for swapping out idle compute for other users and selling the same compute resources many times over for this type of work. Overall, this lowers the price for users.

However, the fundamental design of AWS RDS does not support large-scale mission-critical Oracle Databases in general or support the low-latency high-bandwidth required for real-time analytics. For example, Wikibon's detailed analysis of performance and cost in the "Cost & Performance Analysis of Oracle Autonomous Data Warehouse" section above shows that the same Oracle Database is three times faster on Exadata 9XM and one third the cost. Many customer interviews echo these results.

Bottom line: AWS messages larger-scale AWS RDS for Oracle customers that the problem is the Oracle Database. AWS often suggests its customers convert to some of the 16 different AWS databases that fit the workload best to solve the problem.

Wikibon does not support this messaging and believes this strategy will worsen data quality, make the ETL processes even more complicated, and increase the administrative costs of managing this environment.

Wikibon assesses that adopting this AWS strategy will delay any enterprise IT ability to move to Real-time Analytics or Transactional Analytics by five or more years. Furthermore, this strategy will also dramatically reduce the quality and value of enterprise data warehouses. As a result, a single source of truth will remain

an unreachable aspiration.

However, many customers want AWS and the full Oracle Database. Wikibon would urge Oracle and AWS to work together to provide an effective platform to run Oracle Database with all its features on AWS with a pathway to Transactional Analytics.

4. IBM DB2

IBM DB2 is a Tier-1 database (together with Oracle and Microsoft SQL Server) that still runs the mission-critical databases for many organizations. DB2 was an earlier pioneer in moving toward Transactional Analytics. However, DB2 has fallen behind in investment and is unlikely to offer a sustainable path to Real-time Analytics or Transactional Analytics.

Bottom Line: Wikibon would urge IBM and Oracle to find ways of using the Oracle hardware and software as a platform to run DB2 together with GoldenGate and avoid the cost and elapsed time of conversion from DB2.

5. SAP Database

SAP is both application software and a database. The SAP HANA database was the first analytic database to support the move to real-time analytics with solid support for many data types. However, SAP's attempt to provide a large-scale mission-critical transactional database has proved elusive.

SAP has previously announced that it will cease support for other databases such as Oracle for SAP application software. SAP has pushed back the data for non-support for Oracle and other databases one year every year, and Wikibon strongly believes it will continue to do so. SAP customers are rightly not wanting the disruption and cost of database conversion. As a result, SAP does not have the scale to compete in the transactional database market, which in turn limits its ability to offer both Real-time Analytics and Transactional Analytics.

Bottom Line: Wikibon believes it is unlikely that SAP will become a database vendor with the ability to support Real-time Analytics and Transactional Analytics. However, the SAP application stack is a prime candidate for enhancement with Real-time Analytics and Transactional Analytics. Therefore, Wikibon expects that SAP will work with Oracle to provide enhanced SAP applications.

6. AWS Aurora

AWS Aurora comes in two flavors designed to be compatible with either PostgreSQL or MySQL. One of these is a forked version of PostgreSQL, and the other is a forked version of MySQL database. Both flavors of Aurora have limited capability to provide an analytic converged database with the transactional database. Previous Wikibon research has shown that the MySQL flavor of AWS Aurora is seven times slower and twice the price of Oracle MySQL Heatwave.

Bottom Line: AWS is a knowledgeable and intelligent Cloud operator and understands the importance of specialized hardware and software for future databases. It also understands that waiting for Open Source solutions is not an option.

However, the hardware and software investment for AWS to become a true database vendor is significant and will take at least three years to come to fruition. Therefore, senior IT executives should pressure AWS to commit the resources to become a true database vendor and not buy the AWS multiple optimized tools strategy.

7. Couchbase

Couchbase is a NoSQL (key/value) database with JSON Document schema and a query language known as N1QL, similar to the widely used SQL language. A core benefit of this approach is programmatic control over

data models. The main platform for Couchbase is AWS at the moment, which is limiting their database performance. Couchbase has worked hard at integrating transactional and analytics in this niche area, both with improved parallel performance and with improved SQL functionality.

However, the transactional side of Couchbase has a way to go before enterprise designers will choose it for mission-critical large-scale systems of record. Wikibon believes the transactional and analytic features, performance at scale, and price-performance do not match Oracle Database Enterprise Edition or Oracle HeatWave.

Bottom Line: Couchbase is performing well in niche markets. However, Wikibon believes that Couchbase does not have the market presence and access to funds to invest in the integrated software and hardware required to turn Couchbase into a viable long-term competitor for the large-scale real-time converged Transactional Analytics marketplace.

8. Snowflake

Snowflake burst onto the data warehouse marketplace a few years ago. Wikibon has followed them closely and admires the ease-of-use characteristics of the distributed shared access vision for a shared distributed data warehouse of the future.

Snowflake is trying to make ETL from data sources quicker and more automated, and current Snowflake users will welcome these improvements. However, this approach does not address the fundamental requirement for real-time analytics of guaranteed data consistency and provenance and the long-term simplicity of application and business process design that follows.

Snowflake has made a long-term agreement to use AWS resources at a very significant price reduction. However, recent benchmarks show that Oracle HeatWave has ten times the performance at about half the cost of Snowflake. Wikibon believes that the same limitations of the AWS Cloud architecture discussed above are at the heart of this discrepancy.

Bottom Line: Wikibon believes that Snowflake must provide its users with a much better performance platform than it currently does with AWS. Snowflake should either build its own cloud base or jointly invest heavily in a much better SaaS implementation. In addition, Wikibon believes that the lack of a converged transactional database as part of the vision will limit Snowflake in providing Real-time Analytics and Transactional Analytics in the future. These limitations will be especially important for mid-sized and larger organizations.

Wikibon believes that Snowflake is unlikely to invest in changing its current cloud platform strategy significantly and has no plans to introduce or acquire a transactional database. As a result, Wikibon positions Snowflake in last place.

Strategic Requirements for Real-time Analytics

Real-time Analytics requires four fundamental components:

1. First, sharing a single converged database allows minimal interference between transactional processing (mainly row-based) and analytic processing (different data types and mainly column-based). Example: Oracle Enterprise Converged Database.
2. A highly parallelized memory-based analytics architecture allows increased resources to accelerate the elapsed time of Real-time Analytic queries. Example: the Oracle In-memory feature is integrated with Oracle Exadata X9M.
3. Extreme low-latency high-bandwidth hardware environment tightly integrated with the Linux OS and Oracle Database software. Example: Oracle Exadata X9M.
4. Machine Learning to allow optimization of the elapsed time for Real-time Analytic queries. Example: Machine learning in Oracle Autonomous Data Warehouse.

Real-time Analytics Revolutionize Data Warehousing

Figure 5 above shows that the amount of resources required and the elapsed time required to complete data warehouse workloads is about three times faster than “Best-of-breed” on-premises with Oracle Autonomous Data Warehouse with Exadata X9M, four times faster than AWS, and about five times faster than other MySQL alternatives and Snowflake. This capability means queries can process much more data, and more queries can complete in parallel, which dramatically increases the value of the Real-time Analytics running on Oracle.

Like all new technologies, Real-time Analytics will not replace traditional data warehouses overnight, and IT will support ETL processing for many years. The detailed model in Table 3 in the Footnotes shows a steady decline in ETL costs over five years. However, the reduction in ETL and the increase in Real-time Analytics will drive significant improvement in enterprise data quality. This quality improvement can eventually lead to a single source of truth across enterprises.

Bottom line: Only the Oracle Cloud Infrastructure Solutions with Autonomous Data Warehouse can support Real-time Analytics now and can drive improved quality-of-data faster than any other solution available. Enterprises will acquire the IT skills and business familiarity with Real-time Analytics, a prerequisite for Transactional Analytics discussed in the section below.

Strategic Requirements for Transactional Analytics

The difference between Transactional Analytics and Real-time Analytics is that the transactional part of the system issues the real-time queries. Therefore, Transactional Analytics requires all four of the fundamental components needed for Real-time Analytics. In addition, the transactions themselves must run much faster to allow time for the Real-time Analytics to run.

This integration allows the business processes to have real-time data to enable complete automation and synchronize business processes. For example, knowledge of demand and stocks can help dynamic pricing, inventory locations can allow faster or cheaper delivery options, and understanding supply-chain issues can improve on-time delivery rates. These collective capabilities can simplify business, dramatically increase automation and lower costs, improve customer satisfaction, and reduce churn.

Future Transactional Analytics

Real-time Transactional Analytics combines faster transactions and Real-time Analytics and offers exciting longer-term opportunities. Because the transactions are running faster, time is created for the transaction systems to issue low-latency queries to augment the transactions. The data available to these queries is always up-to-date. Developers can update the transactions in real-time to improve the business process, such as checking for fraud, establishing optimum dynamic pricing, offering alternative payment plans, running ML samples, or advising sales of opportunities or customer dissatisfaction in real-time. Transactions can grow in scope to automate complete business processes, and the real-time capabilities will excite RPA vendors such as UiPath. SaaS vendors, including Oracle, are also likely to lead initial development efforts.

Wikibon believes that applications driven by Transactional Analytics will create enormous business value. Wikibon expects data processing to adopt highly parallel, low-cost, low-power, ARM-based, heterogeneous architectures connecting CPUs, neural engines, GPUs, and other specialized engines with shared DRAM. In general, we believe that processing will become more distributed and move closer to the point of data creation. We expect new techniques we call remote micro-requests will emerge as a standard model for distributed analytics. Tesla, for example, is already using this technique in its vehicles

In summary, faster, parallel computers will speed up transactions, lower transactional and analytic costs, and enable Real-time Analytics and Transactional Analytics. Wikibon believes this breakthrough will lead to dramatic enterprise automation and efficiency improvements.

Wikibon believes that migration to Real-time Analytics and Transactional Analytics will not be a single step but a step-by-step iteration of innovations and learnings that will eventually lead to dramatic changes in enterprise efficiency and effectiveness.

Conclusions & Recommendations

The results of the Wikibon concludes that IT creates today's traditional data warehouses by using ridiculously circuitous and expensive ETL processes. The data warehouses as currently designed are outdated, asynchronous, lead to multiple copies of data, and fail to provide a single version of the truth.

Wikibon has defined Real-time Analytics and believes it will significantly reduce and eventually eliminate the requirement for ETL. Modern hardware and software technologies can provide enterprises with converged real-time data warehouses that are an extension of the transactional database. In addition, the availability in the cloud of large amounts of resources can accelerate complex queries in real-time. As a result, IT should plan to eliminate ETL progressively and instead make enterprise data available to users and applications in real-time. Over time, enterprise data can become more accurate, simpler to manage, and provide the enterprise with a single version of the truth with proven provenance and compliance.

Wikibon has also defined Transactional Analytics, which integrates Real-time Analytics as part of the transaction systems. Wikibon believes Transactional Analytics will be of great value to the lines of business and the supporting application developers to help fully automate business processes.

Oracle is the leading transactional database vendor for large-scale mission-critical Systems of Record. Oracle has pursued a strategy of a single converged database that supports all required data types (including transactional and analytics) and uses modern in-memory techniques to create a data warehouse database able to provide a single version of the truth. Oracle has realized that it must integrate the infrastructure, operating system, virtualization, and database software to provide the performance and latency necessary to make this vision a reality.

Wikibon believes this approach will dramatically improve IT's ability to support the business with increased automation and customer service. Oracle Autonomous Data Warehouse can enable Real-time Analytics without any significant increase in IT budget.

This research is the first of two research papers. This paper aims to show that database systems can be fast enough and cheap enough to make Real-time Analytics and the eventual elimination of ETL. As a result, the vision of a data warehouse with a single version of the truth becomes viable.

The second paper will explore Transactional Analytics, the infrastructure, and software requirements and quantify the potential business benefits of using these techniques.

Action Item

Wikibon strongly advises Executive Management to embrace Real-time Analytics, set a 5-year goal of eliminating at least 90% of all ETL operations, and embrace the vision of a single version of the truth with provenance and compliance. In addition, Wikibon recommends IT development start integrating Real-time Analytics with existing mission-critical Systems of Record.

Wikibon recommends implementing Real-time Analytics on Oracle Autonomous Data Warehouse on Exadata Cloud@Customer or Exadata Cloud Infrastructure X9M because only these platforms have the converged database architecture, performance, price performance, and level of investment to be effective converged real-time database platforms.

Footnotes

vCPU Calculations for Oracle Database Platforms

Calculations of Total Workload vCPUs & % Compute Utilization Required by Oracle Data Warehouse Cloud Platform											
1	2	3	4	5	6	7	8	9	10	11	12
On-Premises "Best-of-Breed" vCPUs Required	On-Premises "Best-of-Breed" % vCPU Utilization	AWS RDS for Oracle vCPUs Required	AWS RDS for Oracle Instances	AWS RDS for Oracle VPUs/ Instance	AWS RDS for Oracle vCPUs Required	AWS RDS for Oracle % vCPU Utilization	Oracle ADW w Exadata X9M vCPUs Required	Oracle ADW w Exadata X9M Base vCPUs	Oracle ADW w Exadata X9M Burst vCPUs	Oracle ADW w Exadata X9M Total vCPUs (Base + Burst)	Oracle ADW Exadata Cloud@Customer & Cloud Infrastructure X9M % Compute Utilization
25%		0%	db.r5b.xlarge	4	28	20%	65%				65%
25% Savings from Virtual Sharing of CPUs & Memory		0% Savings from Virtual Sharing	db.r5b.2xlarge	8	48	18%	65% Savings from Virtual Sharing of CPUs & Memory				
			db.r5b.4xlarge	16	80	11%		45%			
			db.r5b.8xlarge	32	256	10%			28%		
			db.r5b.16xlarge	64	448	13%				73%	65% Savings from Virtual Sharing of CPUs & Memory
			db.r5b.24xlarge	96	384	10%					
933		1,244			1,244	12%	435	196	120	316	

Source: © Wikibon 2022

Table 2- vCPU Calculations and AWS RDS for Oracle Instance Choices for On-premises Data Warehouse (base, Blue), AWS RDS for Oracle (Orange), and Oracle Autonomous Data Warehouse with Exadata X9M (Red).

Source: © Wikibon 2022

Detailed Data by year and IT budget Line Items

Real-time Analytics Revolutionize Data Warehousing

5-year IT Budget for Continuation of On-Premises "Best-of-Breed" IT Infrastructure, with Support and Services for Data Warehousing						
vCPUs	933	933	933	933	933	Total
	Year 1	Year 2	Year 3	Year 4	Year 5	
Compute	\$ 1,721,479	\$ 1,721,479	\$ 1,721,479	\$ 1,721,479	\$ 1,721,479	\$ 8,607,394
DB Storage (Tier1 & IOPS)	\$ 1,108,078	\$ 1,108,078	\$ 1,108,078	\$ 1,108,078	\$ 1,108,078	\$ 5,540,392
Network	\$ 422,125	\$ 422,125	\$ 422,125	\$ 422,125	\$ 422,125	\$ 2,110,625
Operational Support & Application Maintenance	\$ 1,582,809	\$ 1,582,809	\$ 1,582,809	\$ 1,582,809	\$ 1,582,809	\$ 7,914,043
Oracle Database Maintenance	\$ 3,035,193	\$ 3,035,193	\$ 3,035,193	\$ 3,035,193	\$ 3,035,193	\$ 15,175,966
Environmentals	\$ 108,000	\$ 108,000	\$ 108,000	\$ 108,000	\$ 108,000	\$ 540,000
ETL (Software & Staff) costs	\$ 966,898	\$ 966,898	\$ 966,898	\$ 966,898	\$ 966,898	\$ 4,834,491
Expected Downtime Costs	\$ 684,932	\$ 684,932	\$ 684,932	\$ 684,932	\$ 684,932	\$ 3,424,658
Total Cost On-premises	\$ 9,629,514	\$ 9,629,514	\$ 9,629,514	\$ 9,629,514	\$ 9,629,514	\$ 48,147,569
Source: © Wikibon, 2022						
5-year IT Budget for Migration to AWS RDS for Oracle IT Infrastructure Single-AZ, with Support and Services for Data Warehousing						
vCPUs	1,244	1,244	1,244	1,244	1,244	5-year Total
	Year 1	Year 2	Year 3	Year 4	Year 5	
Compute	\$ 519,148	\$ 1,386,006	\$ 1,501,911	\$ 1,501,911	\$ 1,501,911	\$ 6,410,886
Network, Storage & IOPS	\$ 297,726	\$ 794,860	\$ 861,330	\$ 861,330	\$ 861,330	\$ 3,676,576
ETL Costs (Software & Staff)	\$ 966,898	\$ 808,375	\$ 600,856	\$ 600,856	\$ 600,856	\$ 3,577,840
Oracle Database Maintenance	\$ 2,860,000	\$ 3,459,500	\$ 4,059,000	\$ 4,059,000	\$ 4,059,000	\$ 18,496,500
Operational Support & Application Maintenance	\$ 1,582,809	\$ 1,068,396	\$ 641,037	\$ 641,037	\$ 641,037	\$ 4,574,317
Migration Cost	\$ 422,082	\$ 105,521	\$ -	\$ -	\$ -	\$ 527,603
On-premises Costs during Migration	\$ 2,447,962	\$ 917,986	\$ -	\$ -	\$ -	\$ 3,365,947
Expected Downtime Costs	\$ 1,095,890	\$ 822,260	\$ 776,712	\$ 776,712	\$ 776,712	\$ 4,248,288
Total Cost -AWS RDS	\$ 10,192,516	\$ 9,362,904	\$ 8,440,846	\$ 8,440,846	\$ 8,440,846	\$ 44,877,957
Savings compared with "Best-of-Breed"	\$ (563,002)	\$ 266,610	\$ 1,188,668	\$ 1,188,668	\$ 1,188,668	\$ 3,269,612
Cumulative Savings	\$ (563,002)	\$ (296,392)	\$ 892,276	\$ 2,080,944	\$ 3,269,612	\$ 3,269,612
Source: © Wikibon, 2022						
5-year IT Budget for Migration to AWS RDS for Oracle IT Infrastructure Multi-AZ, with Support and Services for Data Warehousing						
vCPUs	1,690	1,690	1,690	1,690	1,690	5-year Total
	Year 1	Year 2	Year 3	Year 4	Year 5	
Compute	\$ 1,038,297	\$ 2,772,012	\$ 3,003,822	\$ 3,003,822	\$ 3,003,822	\$ 12,821,773
Network, Storage & IOPS	\$ 436,603	\$ 1,165,628	\$ 1,263,102	\$ 1,263,102	\$ 1,263,102	\$ 5,391,535
ETL Costs (Software & Staff)	\$ 966,898	\$ 808,375	\$ 600,856	\$ 600,856	\$ 600,856	\$ 3,577,840
Oracle Database Maintenance	\$ 2,860,000	\$ 3,459,500	\$ 4,059,000	\$ 4,059,000	\$ 4,059,000	\$ 18,496,500
Operational Support & Application Maintenance	\$ 1,582,809	\$ 1,068,396	\$ 769,245	\$ 641,037	\$ 641,037	\$ 4,702,524
Migration Cost	\$ 422,082	\$ 105,521	\$ -	\$ -	\$ -	\$ 527,603
On-premises Costs during Migration	\$ 2,447,962	\$ 917,986	\$ -	\$ -	\$ -	\$ 3,365,947
Expected Downtime Costs	\$ 1,095,890	\$ 575,942	\$ 448,288	\$ 448,288	\$ 448,288	\$ 3,016,695
Total Cost -AWS RDS for Oracle	\$ 10,850,540	\$ 10,873,358	\$ 10,144,311	\$ 10,016,104	\$ 10,016,104	\$ 51,900,417
Savings compared with "Best-of-Breed"	\$ (10,850,540)	\$ (10,873,358)	\$ (10,144,311)	\$ (10,016,104)	\$ (10,016,104)	\$ (51,900,417)
Cumulative Savings	\$ (10,850,540)	\$ (21,723,898)	\$ (31,868,210)	\$ (41,884,314)	\$ (51,900,417)	\$ (51,900,417)
Source: © Wikibon, 2022						
5-year IT Budget for Migration to Oracle Autonomous Data Warehousing on Exadata Cloud@Customer X9M, with Support & Services for Data Warehousing						
vCPUs	-	-	-	-	-	5-year Total
	Year 1	Year 2	Year 3	Year 4	Year 5	
Base & Burst vCPUs	\$ 320,925	\$ 446,504	\$ 446,504	\$ 446,504	\$ 446,504	\$ 2,106,942
Oracle Exadata Network & Storage	\$ 182,794	\$ 254,323	\$ 254,323	\$ 254,323	\$ 254,323	\$ 1,200,084
ETL Costs (Software & Staff)	\$ 566,244	\$ 230,385	\$ 172,789	\$ 115,193	\$ 57,596	\$ 1,142,207
Oracle Database Maintenance	\$ 1,317,888	\$ 1,003,118	\$ 1,003,118	\$ 1,003,118	\$ 1,003,118	\$ 5,330,358
Operational Support & Application Maintenance	\$ 324,228	\$ 451,100	\$ 451,100	\$ 451,100	\$ 451,100	\$ 2,128,630
Migration Cost	\$ 316,562	\$ -	\$ -	\$ -	\$ -	\$ 316,562
On-premise Costs	\$ 1,835,971	\$ 216,000	\$ 216,000	\$ 216,000	\$ 216,000	\$ 2,699,971
Expected Downtime Costs	\$ 1,659,247	\$ 179,775	\$ 157,303	\$ 134,832	\$ 112,360	\$ 2,243,516
Total Cost Exadata OCI X9M	\$ 6,523,859	\$ 2,781,205	\$ 2,701,137	\$ 2,621,069	\$ 2,541,001	\$ 17,168,271
Savings compared with "Best-of-Breed"	\$ 3,105,655	\$ 6,848,308	\$ 6,928,377	\$ 7,008,445	\$ 7,088,513	\$ 30,979,298
Cumulative Savings	\$ 3,105,655	\$ 9,953,963	\$ 16,882,340	\$ 23,890,785	\$ 30,979,298	\$ 30,979,298
Source: © Wikibon, 2022						
5-year IT Budget for Migration to Oracle Autonomous Data Warehousing on Exadata Cloud Infrastructure X9M, with Support & Services for Data Warehousing						
vCPUs	316	316	316	316	316	5-year Total
	Year 1	Year 2	Year 3	Year 4	Year 5	
Base & Burst vCPUs	\$ 279,065	\$ 446,504	\$ 446,504	\$ 446,504	\$ 446,504	\$ 2,065,082
Oracle Exadata Network & Storage	\$ 158,952	\$ 254,323	\$ 254,323	\$ 254,323	\$ 254,323	\$ 1,176,242
ETL Costs (Software & Staff)	\$ 555,445	\$ 230,385	\$ 172,789	\$ 115,193	\$ 57,596	\$ 1,131,408
Oracle Database Maintenance	\$ 1,317,888	\$ 1,003,118	\$ 1,003,118	\$ 1,003,118	\$ 1,003,118	\$ 5,330,358
Operational Support & Application Maintenance	\$ 281,938	\$ 451,100	\$ 451,100	\$ 451,100	\$ 451,100	\$ 2,086,339
Migration Cost	\$ 316,562	\$ -	\$ -	\$ -	\$ -	\$ 316,562
On-premises Costs during Migration	\$ 1,835,971	\$ -	\$ -	\$ -	\$ -	\$ 1,835,971
Expected Downtime Costs	\$ 873,288	\$ 164,932	\$ 144,315	\$ 123,699	\$ 103,082	\$ 1,409,315
Total Cost Exadata OCI X9M	\$ 5,619,108	\$ 2,550,362	\$ 2,472,149	\$ 2,393,936	\$ 2,315,723	\$ 15,351,277
Savings compared with "Best-of-Breed"	\$ 4,010,406	\$ 7,079,152	\$ 7,157,365	\$ 7,235,578	\$ 7,313,791	\$ 32,796,292
Cumulative Savings	\$ 4,010,406	\$ 11,089,558	\$ 18,246,923	\$ 25,482,501	\$ 32,796,292	\$ 32,796,292
Source: © Wikibon, 2022						

Table 3 - Wikibon Detailed Financial Analysis of Four Alternatives for Migration to

Oracle Database Cloud Services

Source: © Wikibon 2022

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