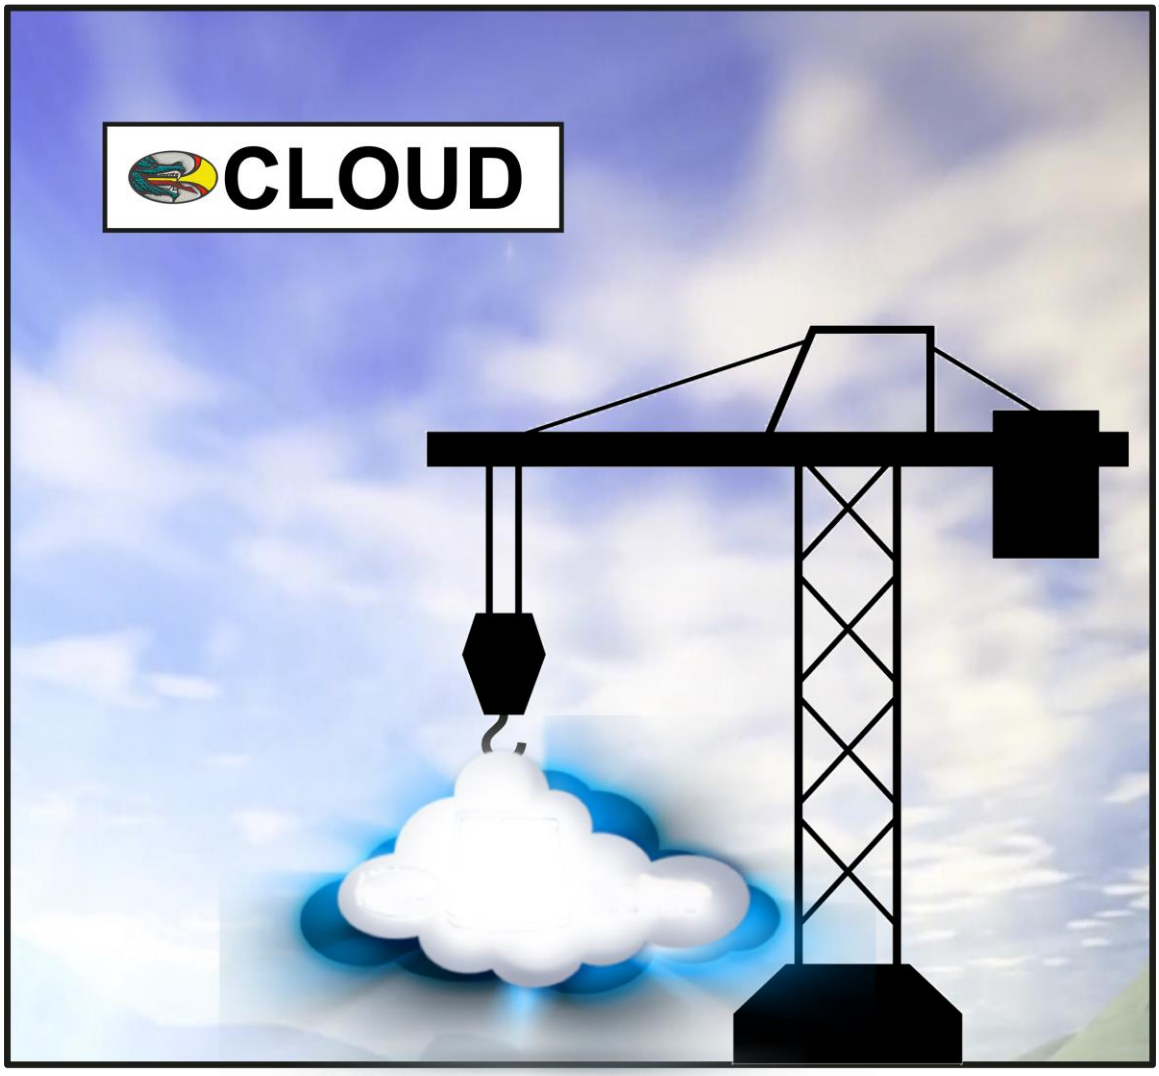

Oracle Builds a Most Compelling Data Platform on a Next-gen Public Cloud

Based on Leading Edge Autonomy, Performance, and Functionality

by, Marc Staimer

2023

Premise

To build a most compelling data platform, Oracle first had to build a compelling next generation public cloud. Compelling technology and services have historically meant being twice as fast, with twice as much capability, and with greater simplicity at half the cost of current market leaders. That definition is just as true for data platforms and public clouds.

The clouds from Amazon Web Services (AWS), Microsoft Azure (Azure), and Google Cloud Platform (GCP) are generally pretty similar architecturally. There are noticeable differences among them in the services they offer and how they price consumption, but these differences are neither large nor compelling. And considering that most mission- and business-critical applications tend to require a high-performance, high-availability database, their cloud database services have generally not been good enough to meet those needs. AWS, Azure, and GCP are also not known for their simplicity. They offer a smorgasbord of specialized services and provide various types of integration services to glue everything together so that data can move

between them—all adding complexity to applications. Outstanding or compelling are not the adjectives commonly associated with these cloud services.

For another vendor to take market share from AWS, Azure, or Google, they have to provide a compellingly better public cloud. That is exactly what Oracle has accomplished with their next-generation Oracle Cloud Infrastructure (OCI) that supports an equally compelling Oracle Data Platform and provides a welcome antidote to the complexity of other clouds.

This research examines whether or not OCI is actually a compellingly better public cloud and how Oracle Data Platform also meets the compelling criteria. And if so, how Oracle made it happen with capabilities such as Autonomous Database, Exadata, and other highly differentiated services.

Research Key Findings

Oracle's drive to create a compelling public cloud began with Larry Ellison, Oracle's CTO, and founder. He wanted convincing differentiation from both an architecture and services perspective. This meant using extremely high levels of automation to deliver the most capable database services, application services, and cloud architecture that delivers industry leading performance, greater availability, security, simplicity, and much lower costs. In other words, it wouldn't be the same as AWS, Azure, and GCP.

Indeed, OCI is not your father's cloud. It's an advanced next-generation cloud purpose-built for cloud applications, architected with RDMA clusters for superior performance and low latency, taking advantage of Exadata as a database platform, with built-in by default zero trust security, autonomous services, and a range of deployment options, including hybrid cloud and multicloud. The result is superior support for many different workloads, including streaming video, AI services, high-performance-computing (HPC), every king of database service, as well as home-grown and packaged Oracle or third-party applications.

While Oracle's video, AI, and HPC capabilities are sexy and eye-catching, it's the day-to-day organizational data generated by transactional applications and in depth data analysis that enables faster-better decisions. This is where Oracle Data Platform has an enormous impact.

The Oracle Data Platform leverages OCI to empower organizations to run applications faster, consolidate enterprise-wide data, and share it with applications and/or people that need it. It additionally simplifies external data source integration, extracting real value from all the data, and ultimately lower costs.

Oracle Data Platform brings together a comprehensive data services suite in an integrated, open, and modular approach that supports virtually any type of data and workload. These capabilities enable organizations to move existing on-premises workloads to the cloud considerably faster and with less effort. Oracle Data Platform also accelerates modern application development and deployment, increases run time efficiencies, improves results, and lowers costs. This creates massively synergistic results beyond what can be achieved with individual OCI services or other public cloud counterparts.

A crucial element in many Oracle Data Platform deployments is Oracle Autonomous Database. It's based on the Oracle Database, the most popular enterprise databases on the planet, having been ranked number 1 for decades on DB-Engines Database Rankings[1]. To put that in perspective, 19 of the top 20 Fortune Global 500 banks run significant portions of their business on the Oracle Database and Oracle Exadata systems. Autonomous Database adds a market unique level of autonomy to the Oracle Database which is discussed in detail within this research.

Autonomous Database and Oracle Database are the indisputable leaders of databases because of their:

- Production proven performance that is markedly faster than any other converged database with far lower latencies, much greater IOPS, and higher throughput—particularly when running on Oracle Exadata systems.
- Substantially higher uptime and availability with unmatched essential functions such as Oracle Real Applications Clusters (RAC) and Autonomous Data Guard for unparalleled rapid disaster recovery.
- Converged databases that support OLTP, OLAP (data warehouse), time series, JSON (document) including a MongoDB-compatible API, XML (object), key value, spatial, graphic, block chain, data lake,

[1] They also have MySQL open source database – also owned by Oracle – ranked number 2 for years.

and AI/machine learning (ML), instead of requiring a dozen different single-purpose databases as in other clouds.

- Efficiency supporting an enterprise's entire database fleet, regardless of data and database type, with extensive multi-tenancy, container databases (CDB), and pluggable databases (PDB) that eliminate many data silos.
- Tightly integrated, best-in-class security capabilities. This security has decades of continuous improvement based on production hardened best practices that apply to all database types instead of fragmenting security operations across multiple specialized databases and integration services.
- Relentless innovation. An example of which is the recently announced JSON Relational Duality capabilities. That duality solves a major problem for developers by enabling them to store data and build applications accessible by both relational and/or JSON paradigms.
- Industry leading predictable cloud total cost of ownership (TCO) reductions including platforms, administration, and consumption—consistently shown to be lower than other public cloud competitive services.

How Oracle Built OCI into the Most Compelling Next Gen Public Cloud

OCI's Unique Autonomy

Mr. Ellison added an extremely important principle that is essential to making OCI a compelling public cloud—autonomy. Oracle Autonomous Database is a prime example of how Oracle is using AI/ML technologies to create a differentiated cloud. Autonomous Database was designed and built from the ground up for the cloud with AI/ML-based automation, running on high-performance Oracle Exadata infrastructure in OCI. Autonomous Database incorporates more than 40 years of Oracle best practices for database tuning, scaling, maintenance, patching, security, and many other areas. Its autonomy operates continuously, reacting in real-time to workload changes or the underlying Exadata systems. Those reactions are much faster than any DBA or hardware administrator can provide. This autonomous high performance, availability, security, and intuitive self-service, makes the Autonomous Database the simplest cloud database service anywhere in the world.

There are currently three Autonomous Database cloud services, each tuned slightly differently for specific workloads.

1. Autonomous Transaction Processing
   - Tuned for transactions and database consolidation.
2. Autonomous Data Warehouse
   - Tuned for data warehouse and data lake analytics.
3. Autonomous JSON Database
   - Tuned for JSON documents.

Regardless of how an Autonomous Database instance is tuned, each comes with the entire complement of Oracle Database capabilities, with autonomy solving a major IT problem—the shortage of knowledgeable, experienced, and highly skilled database administrators (DBAs) which keeps getting worse as experienced baby boomer DBAs retire and leave the work force. Autonomous Database's automation means that IT teams and DBAs no longer have to be adept at manually performing many traditional and labor-intensive activities such as database provisioning, tuning, patching, securing, and scaling. This has the substantial positive side effect of radically reducing human errors that impact business operations. And when DBAs no longer have to waste their precious time on labor-intensive database administration, they can focus instead on the strategic priorities of the business—such as new projects, new services, new products, more differentiation, increased productivity, revenues, and profits. As a result, organizations win twice.

OCI's Unparalleled Hardware Infrastructure Delivers Unparalled Performance

One outcome of Autonomous Database's automated services is they frequently run faster than non-automated equivalents. The ability the Oracle Data Platform's various services to automatically scale up performance when needed, and back down when not needed, allows much faster time-to-actionable insights, greater confidence in results, and accelerated product or service innovation delivery to the market.

Delivering on the Autonomous Database's incredible performance requires the right infrastructure to support it. That infrastructure has to effortlessly run all workloads very fast–often concurrently, efficiently utilize resources, elastically scale-up and scale down—preferably without human intervention and do all of this without hiccups. That's why it exclusively runs on Exadata Cloud Infrastructure in OCI and Exadata Cloud@Customer systems in enterprise data centers.

Oracle Exadata systems are co-engineered with the Oracle Database and Autonomous Database to deliver the lowest latencies, highest IOPS, and most throughput for any type of database workload. It further requires deep engineered integration of the entire software stack with the complete hardware infrastructure down to the components, including the microcode. Doing this necessitates comprehensive cooperation and close collaboration between Oracle and the hardware infrastructure vendors. It's source code-level co-engineering.

AMD is one of Oracle's most important hardware partners in this endeavor with their AMD EPYC™ processor. After extensive evaluations, Oracle determined it provided the ideal combination of performance, flexibility, and cost for Oracle Autonomous Database and Exadata requirements in OCI. To that end, every current generation Exadata database server on OCI is built using AMD EPYC processors.

This synergistic combination of hardware and software delivers unequaled performance based on the published stats.

- Average read latency of ≤ 19 µs.
- With just 8 database servers, which is ¼ of what Exadata Cloud Infrastructure supports.
  - ≥ 22.4 million Oracle Database read IOPS.
  - ≥ 5.9 million 8K SQL write IOPS.
- When fully scaled out, Exadata systems in OCI supports.
  - Up to 2.8 TB/s in scan throughput.
  - Up to 3.2 PB of database capacity.
  - Up to 32 PB capacity with hybrid columnar compression (HCC).
  - Up to 32 EB capacity with HCC and sharding.

For both Autonomous Database and Exadata Database Service on OCI, Exadata systems deliver 25x better IO latency than all-flash Amazon RDS and 50x better IO latency than all-flash Azure SQL, while scan rates can be more than 100x faster than Amazon RDS or Azure SQL. By any measure, Exadata is the fastest Oracle converged database platform in the cloud or anywhere else. And since every cloud database service, including Autonomous Database, is priced based on how long a scalable number of processors are used, faster performance and automatic scaling of database consumption equates directly into lower cost. Time really is money in the cloud. For end users, faster performance means superior application response times that radically increase productivity, save money, and accelerate revenues as described in detail in IBM's illuminating research published in 1982 and 1997 "The Economic Value of Rapid Response Time" (see additional information in Appendix A).

The Value of the AMD EPYC Processors

AMD EPYC™ processors provide the highest core count *and* the fastest x86 cores. They have augmented multi-processor performance and scalability through the latest advancements in die-to-die connectivity with the AMD Infinity Fabric™. EPYC processors in the current generation of Exadata systems in OCI also provide exceptional memory bandwidth with 8 memory channels per socket delivering a peak 410 GB/s DRAM bandwidth. I/O bandwidth is also industry-leading with 8x16 PCIe Gen4 links, equating into 128 PCIe 4.0 lanes per socket and 160 per dual socket. Each bandwidth link can support up to 64GB/s bi-directionally for a total of 512GB/s per socket.

AMD EPYC processors don't just increase Autonomous Database's performance envelope, they also raise the bar on the CPU scalability powering Exadata's scale-out architecture. Each Exadata Cloud Infrastructure database server comes with two AMD EPYC processors with 64 cores per socket, giving each server a total of 128 cores and 256 threads. Oracle reserves two cores per database server to manage the hypervisor, so each server has 126 cores / 252 threads available for database processing. Exadata Cloud infrastructure then combines 2 to 32 database servers to create a high-availability database cluster with 252 to 4,032 cores and hundreds of terabytes to petabytes of database storage capacity. From an Autonomous Database point of view, AMD EPYC processor-powered Exadata systems deliver the parallelism and performance needed to deliver best-in-class performance for both transaction processing and analytics.

OCI uses AMD EPYC processors with a greater number of speedier cores to process transactions and run analytical queries faster than any other database or cloud database service. The large number of processor cores in Exadata Database Servers also allows organizations to run complex analytics and AI/ML analyses in the database without impacting the performance of other databases also running on the system. Oracle additionally leverages AMD EPYC processor's large L3 on-chip cache to store and access data much quicker—resulting in significantly faster query response times.

Oracle's move to AMD EPYC processors gives it 2.5 times more database processing cores than the previous Exadata generation. This has led to smoother customer operations and efficient operations with no downtime – especially critical for business and mission-critical applications. This leads to better overall customer experiences, faster time-to-actionable-insights, and faster time-to-market. The larger number of AMD EPYC processor cores per database server and extensive amounts of shared storage enable efficient consolidation of previously siloed databases, which means that customers need fewer servers, have a reduced carbon footprint, and best of all, much lower customer costs. Cloud time equates to cost, and faster processing equals less processing time—translating into lower costs.

The requirements that Larry Ellison set for OCI are incredibly difficult to achieve. And yet, Oracle achieved every one—autonomy, much faster performance, more advanced functionality, simplicity, always on security, and considerably lower cost relative to other clouds. OCI is, as of 2023, the most compelling next generation public cloud currently available.

It All Comes Together in the Oracle Data Platform

Oracle Data Platform is the epitome of a public cloud data platform. It's a comprehensive suite of tools, services, and applications specifically architected for businesses to manage their data easily, efficiently, and cost effectively throughout its lifecycle. Oracle Data Platform combines multiple products and services for transaction processing, data warehouse, data lakes, AI/ML, and more. It is an integral part of how OCI provides unmatched performance and secure data management. It is not only the definition of comprehensive, but also the most open and collaborative in its design—allowing customers to replace or extend capabilities with those of their own choosing.

Oracle Data Platform empowers users to connect multiple data sources from within and beyond their organization to create a complete "single-source-of-truth" for their enterprise. Data can come from cloud and on-premises databases, files, unstructured, and semi-structured data stores and be analyzed by myriad different analytical and AI/ML techniques. It specifically helps organizations discover deeper insights, better predict outcomes, increase employee productivity by automating processes, increase customer engagement with targeted offers, and meet both revenue and TCO goals.

Oracle Builds a Most Compelling Data Platform and Next-gen Public Cloud

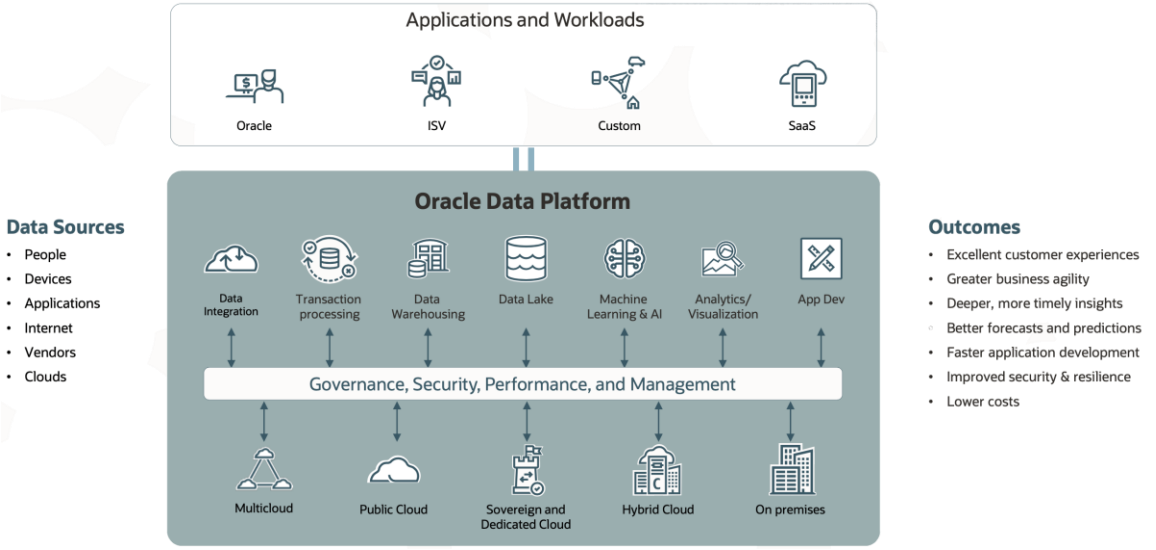

Getting this best-in-class market-leading performance requires deeply engineered integration of the entire software stack and the complete hardware infrastructure down to the components, including the microcode. Doing this necessitates comprehensive cooperation and close collaboration between Oracle and the hardware infrastructure vendors. It's source code-level co-engineering.

AMD is one of Oracle's most important hardware partners in this endeavor with their AMD EPYC™ processors, with every current Exadata database server in OCI is built around it. After extensive evaluations, Oracle determined it provided the ideal combination of performance, flexibility, and energy efficiency for Oracle's Autonomous Database and Exadata requirements in OCI. All of this adds up to the most compelling data platform on an equally compelling public cloud that leverages Oracle Autonomous Database, Exadata, and myriad OCI services running on AMD EPYC processors.

However, OCI is not just limited to Autonomous Database on Exadata when it comes to meeting Larry Ellison's blueprint for the most compelling next generation public cloud. Not by a long shot. The performance of AMD EPYC processors is available for a number of other crucial services within and beyond the Oracle Data Platform, including MySQL HeatWave, Hadoop and Spark data lake capabilities in the Oracle Big Data Service, support of the Oracle VMware Cloud Solution, and flexible deployment of VM and bare-metal compute shapes. A key example is Oracle MySQL HeatWave.

Oracle MySQL HeatWave

Oracle MySQL HeatWave is the most advanced cloud database service for MySQL, the world's most popular open source database. It deeply integrates the MySQL OLTP database with HeatWave OLAP, AutoML, Autopilot, and data lakes into a single code base. AutoML automates and accelerates the machine learning lifecycle and Autopilot automates many manual functions.

MySQL HeatWave also takes advantage of the AMD EPYC processors in OCI, which is one of the reasons its performance and TCO/performance blows away the competition. The proof is once again in the performance and cost numbers.

Based on the industry standard 4 TB TPC-H[2] benchmark (OLAP), MySQL HeatWave on OCI delivers 1100X better performance at ½ the cost of the Amazon Aurora MySQL service. The advantages continue with dedicated OLAP cloud database services. In larger 30TB TPC-H[2] benchmark comparisons versus Amazon Redshift, Snowflake, Google BigQuery, and Azure Synapse, MySQL HeatWave delivers 3.7x to 15x better performance and 20x to more than 70x better price/performance as reported in Wikibon's recent TCO and TCO/performance research. In machine learning, MySQL HeatWave AutoML trains models 25X faster than Amazon Redshift ML at 1% of the cost.

Customers that have switched from competitor clouds to MySQL HeatWave on OCI have claimed up to 10 times the performance at less than 1/3rd the cost. That's again, very compelling!

[2] Benchmark queries are derived from the TPC benchmarks, but results are not comparable to published TPC benchmarks results since these do not comply with the TPC specifications.

Problems Solved by Oracle Data Platform

- Eliminates database, data warehouse, data lake, and ML performance issues, accelerating response times, actionable insights, time-to-market, and time-to-revenues.
- Eradicates the conspicuous decline in database performance as data proliferates, especially performance decreases exacerbated by sprawl.
- Completely reverses mediocre-to-poor user productivity caused by slow and variable response times by accelerating them consistently.
- Obsoletes all other public cloud database services performance by accelerating results, time-to-actionable-insight, time-to-market, and time-to-revenue.
- Obliterates sub-optimal database availability from planned and unplanned disruptions/outages—including patching and upgrades that degrade user productivity by making it all non-disruptive with best-in-the-industry uptime and availability. With public cloud unprecedented availability at 99.995% uptime.
- Radically reduces public cloud costs making it much more affordable.

Don't Other Clouds Offer AMD EPYC Processors?

AMD EPYC processors are available in other clouds; however, NONE have done the cooperative co-engineering of the entire database software and hardware stack that delivers the unparalleled performance and scalability of the Oracle Data Platform, Oracle Autonomous Database, and Oracle MySQL HeatWave in OCI.

Conclusion

Oracle set out to build a significantly faster, more autonomous, more scalable, more reliable, more secure, and lower cost public cloud. In other words, Oracle set a goal of making their Oracle Cloud Infrastructure compellingly better than any of their competitors.

Equally, Oracle aimed to create a comprehensive, open, and compellingly better data platform that enables the storing and analyzing of data from all sources with ease.

Leveraging Autonomous Database, Oracle Database, Exadata, and MySQL HeatWave running on AMD EPYC processors in OCI, they have wildly succeeded at both efforts.

For IT organizations seeking higher productivity at lower costs, this partnership is definitely strong and growing to include additional services. Whether it be demanding high performance applications, workload consolidation, or developing data-driven applications there is no more effective and efficient line of cloud database services than those powered by the AMD and Oracle partnership in OCI.

For More Information

Go to:

AMD EPYC Processors

Oracle Data Platform

Oracle Autonomous Database Video Explainer

Oracle Autonomous Databases

Oracle MySQL HeatWave Video Explainer

Oracle MySQL HeatWave

Oracle Cloud Infrastructure

Appendix A: Rapid Response Time Economic Value

Calculating performance productivity impact:

- Performance has a substantial and measurable impact on productivity.
  - Response time has a direct correlation on user productivity, quality-of-work, and time-to-market.
  - It was determined that the maximum application response time before user productivity declines precipitously is 3 seconds. Anything over 2 second response times caused user attention to wander.
  - Application response times that are less than 3 seconds promptly increase user productivity, quality-of-work, and time-to-market.
  - Reducing response time to ~ .3 seconds more than doubles productivity versus 2 seconds. Productivity gains are substantially greater depending on the user's level of expertise.

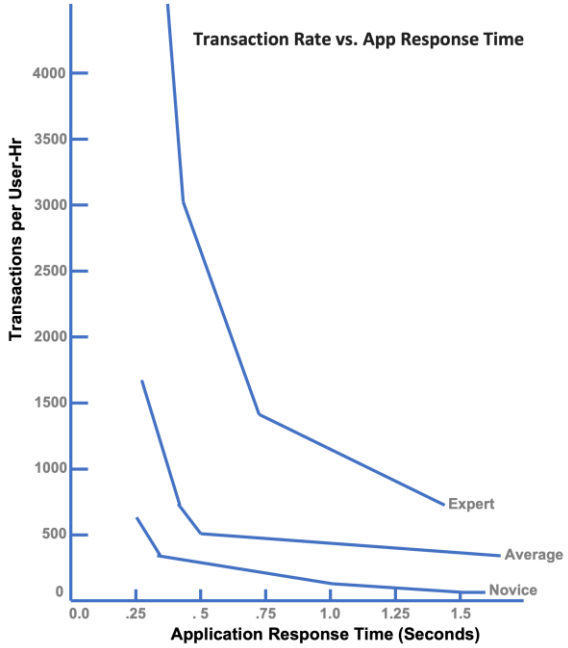

  - Faster response times mean shortened project schedules and higher work quality.
  - ≤ .4 seconds equates into what is called the Doherty Threshold. The Doherty Threshold is when response time becomes addictive whereas > .4 seconds is painful and users' attentions begin to stray.

Application Response Time (Sec)	Transactions per Hr	Task Time (Min)	Time Saved per Task (Min)	Time Saved per Day
3	180	60	-	-
2	208	51.9	8.1	1h/4m/48s
1	252	42.9	17.1	2h/16m/48s
0.6	279	37.7	22.3	2h/58m/24s
0.3	371	29.1	30.9	4h/7m/12s

  - Determine application response times for each service under consideration.
  - Compare productivity rates.
  - Divide FTE costs by productivity to calculate FTE cost per transaction.
  - One alternative is to compare the time required to complete a defined set number of transactions.
  - Multiply the time saved by FTE average hourly cost.

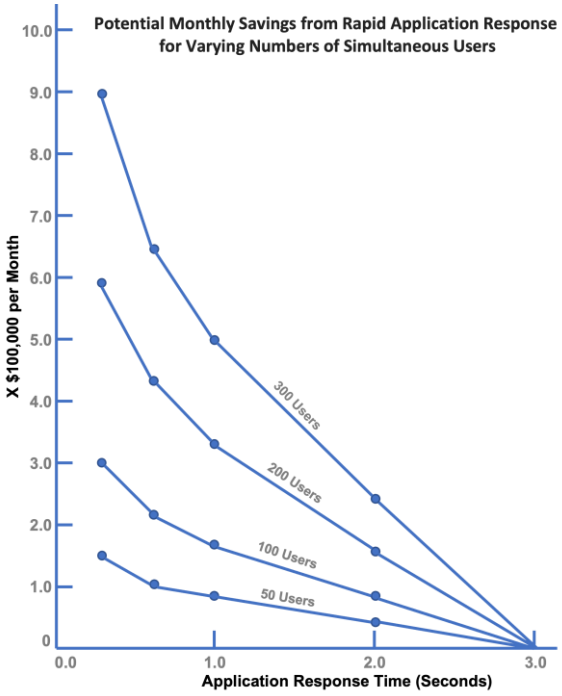

Time-to-market revenue acceleration increases top line revenues and botom line profits

- Based on current schedules estimate the following:
  - Amount of revenue for each week or month schedule is moved up.
  - Project how much time the reduced application response time performance will accelerate the time-to-market. This can be derived from the increase in productivity based on application response time. If the developers can more than double their productivity, they can more than cut in half the amount of time to complete their project.
  - Apply the projected market growth rate to that revenue for a set period, anywhere from 1 – 10 years. Compare the total revenues to what is would have been had the schedule not been accelerated. The differences are the unique gains. If the database cloud service delays time to market, then the differences are the unrecoverable losses.
  - Example from a large microchip manufacturer:
    - By accelerating delivery of their chip to market by one quarter they were able to realize unique revenues > than $100 million upfront and five times that amount over 3 years.