



Author: Marc Staimer

W H I T E P A P E R

How Oracle Autonomous Database (ADB) Services: ATP, ADW, and AJD **Transform Database IT Outsourcing** (Updated)

How Oracle ADB Services: ATP, ADW, AJD Transform Database IT Outsourcing

Introduction

Database IT outsourcing refers to the hiring of a third-party managed service provider (MSP) to deliver the maintenance, support, and general manually labor-intensive work for the IT organization's operational databases (DB). The apt descriptor here is "general." The Database IT outsourcer is paid to monitor database operations, ensure performance service level agreements (SLA) are being met, and maintain the critical parameters. These MSPs are typically responsible for tuning, indexing, patching, troubleshooting, performing root cause analysis of problems, and applying fixes. Whenever a database alert is generated, it is their responsibility to react quickly and resolve whatever database problem generated the alert.

Database IT outsourcing very rarely includes code development, SQL statements, application performance tuning, application maintenance, or even database performance tuning. It is almost exclusively focused on database management and keeping it running within SLA parameters. In other words, Database IT outsourcing is aimed squarely at taking over the role of the database system DBAs and task-specific DBAs. That's it. They are not concerned with continual database performance optimization, application integration, database servers, database storage, or database networking. All of the other DBA functions still belong to the customer IT organization. The IT organization DBA relief expectations rarely coincide with actual reality.

Most Database IT outsourced DBA functions are performed off-site. In some cases, the IT organization can contract for on-prem personnel, albeit at a premium cost. Database IT outsourcing can monitor database instances running on-prem, in a private cloud, or in a public cloud. As more IT organizations move their databases to public clouds or utilize public cloud database-as-a-service (DBaaS), they find that moving to public clouds or DBaaS does not relieve them of the DBA tasks and responsibilities. Any notion that it would, is quickly dispelled. They still need skilled DBAs. And that is becoming a fundamental problem. There aren't enough of them, especially the knowledgeable, skilled, experienced, and expert ones to go around. That leads many to consider Database IT outsourcing¹.

Database IT outsourcing is not without its issues. The business aims at using economies of scale to leverage their experienced expert DBAs as a service. But because it utilizes more high-priced DBA experts, and the DBA work is labor-intensive, Database IT outsourcing is expensive. Anything that is human manually labor-intensive is going to be costly.

If the IT organization's database administration requirements are nominal, Database IT outsourcing can make sense. For the vast majority, there really needs to be a better, more efficient, and more cost-effective way. This document looks at the DBA expertise shortage problem, why Database IT outsourcing is currently suboptimal, and how the Oracle Autonomous Database (ADB) thoroughly solves the DBA shortage problem both now and in the future while transforming Database IT outsourcing into a more strategic versus tactical service.

¹ There are some DBA outsourcing firms that also host single tenant or a multi-tenant DBaaS as well.

Table of Contents

Introduction2

The Marketplace DBA Expertise Shortage Problem4

 DBA Complexity.....4

 Distinct DBA Roles4

Database Architect 4

Database System DBA..... 4

Task-Specific DBA..... 4

Database Analyst..... 4

Application DBA..... 4

Data Warehouse DBA 5

AI DBA 5

Performance DBA or Analyst..... 5

Cloud DBA..... 5

 Additional Non-DBA On-Prem Administrators – Not Required for DBaaS.....5

Server Admin 5

Storage Admin 5

Networking Admin..... 5

Facilities Admin..... 5

 Trends Exacerbating the DBA Shortage.....6

 DBA Shortage Consequences.....6

Why Database IT Outsourcing Comes Up Short.....6

The Oracle Autonomous Database Solution.....7

 Overview.....7

 Defining the Autonomy Levels in an Autonomous Database8

 How The Oracle Autonomous Database Compares to the Market8

 Oracle Autonomous Database Underlying “Database” Foundation8

 Oracle Autonomous Database Underlying “Hardware” Foundation.....9

 Oracle Autonomous Database Cloud Service10

 Oracle Autonomous Database Tuned Services11

Autonomous Transaction Processing Database (ATP) 11

Autonomous Data Warehouse (ADW) 11

Autonomous JSON Database (AJD) 11

Conclusion 13

For More Information on the Oracle Autonomous Database Services 14

The Marketplace DBA Expertise Shortage Problem

Databases are complicated software infrastructure. They are the heart and soul of most mission-critical applications and workloads. Most modern applications require a database of some kind today. Whether it's e-Commerce, online transaction processing, online analytical processing, genome decoding, protein analysis, seismic processing, molecular analysis, video analytics, AI-machine learning, AI-deep machine learning, or AI-neural networks, a database is generally imperative to a successful application.

There are many database types including relational, data warehousing, key value, object, XML, JSON, document, columnar, time series, graphical, spatial, combinations, and more. Regardless of the database type, every database has historically required one or more database administrators (DBA).

DBA Complexity

The DBA is an essential job. They are responsible for performing and/or directing all actions that ensure the databases effectively do what they're supposed to do, without fail. Those responsibilities include: database design; planning; implementation; configuration; ETL (data extraction, transformation, and loading); indexing; capacity provisioning; performance tuning; optimization; SQL coding; data integrity; troubleshooting; root cause analysis; security; backup/recovery; disaster recovery; and business continuity. DBAs are also responsible for working with application design and implementation to make sure the outcome meets the business needs.

Most of these database tasks are manually labor-intensive, requiring extensive cross-functional knowledge, skill, and experience in the database itself plus, servers, memory, I/O, storage, and networking. When the database is in a public cloud, the DBA must have extensive knowledge and expertise in the cloud's virtualized infrastructure, how to tune, troubleshoot, and fix it when the database is not doing what they need it to do. And that's just for a single database. Few organizations are dependent only on a single database today. Large organizations have thousands of databases and legions of DBAs to support them. That's a lot of responsibility for every single database with each one demanding multiple DBAs.

Larger IT organizations separate the vast array of DBA responsibilities into multiple DBA positions as described in the table below.

Distinct DBA Roles
<i>Database Architect</i>
<ul style="list-style-type: none"> • Designs and implements databases. From modeling to physical database design based on access requirements, performance, and database resiliency (backup, recoveries, restores, recovery point objectives, and recovery time objectives.)
<i>Database System DBA</i>
<ul style="list-style-type: none"> • Technical vs. business focus. Emphasis is on implementing, configuring, indexing, tuning, patching, upgrading, troubleshooting, root cause analysis, and fixing.
<i>Task-Specific DBA</i>
<ul style="list-style-type: none"> • Specifically focused on particular admin tasks such as patching, database protection, or testing.
<i>Database Analyst</i>
<ul style="list-style-type: none"> • Generally, a junior DBA similar to an apprentice. Some IT organizations use this term instead of DBA.
<i>Application DBA</i>
<ul style="list-style-type: none"> • Responsible for database application design and support. Skilled at creating, debugging, and optimizing SQL for the application and efficient use of the database.

Data Warehouse DBA
<ul style="list-style-type: none"> As the title suggests, this DBA specializes in extracting optimal data warehouse performance. In addition to typical DBA skills, they also have expertise in analytics, business intelligence, query tools, and especially ETLs.
AI DBA
<ul style="list-style-type: none"> Like the data warehouse DBA, the AI DBA has all of the DBA strengths plus data scientist skills in “R” packages. Must be able to teach AI-machine learning from the correct data. Focus is on making AI machine learning algorithms more efficient and effective.
Performance DBA or Analyst
<ul style="list-style-type: none"> As the title implies, their job is to monitor, improve, and optimize database application performance.
Cloud DBA
<ul style="list-style-type: none"> The DBA functions have not gone away as databases have moved to the cloud. Whether it is a database-as-a-service (DBaaS) or a licensed database running in a cloud instance, DBAs have not been taken over by the cloud service providers, at least not without a separate outsourcing contract (more on this later.) Cloud DBAs must know the capabilities and limitations of running the database in the cloud service provider. For example: AWS limits the EC2 instances ≤ 80,000 IOPS, which can greatly constrain a particular database instance.
Additional Non-DBA On-Prem Administrators – Not Required for DBaaS
Server Admin
<ul style="list-style-type: none"> This admin is responsible for designing, implementing, configuring, managing, operating, patching, troubleshooting, determining root cause to problems, and fixing the database servers. Most DBAs are not infrastructure experts.
Storage Admin
<ul style="list-style-type: none"> Storage tends to be a unique skill set in the data center. This admin is responsible for designing, implementing, configuring, managing, operating, performance optimizing for the database, patching, troubleshooting, determining root cause to problems, and fixing the database storage, providing database data protection and resilience, with as few disruptions as possible.
Networking Admin
<ul style="list-style-type: none"> Just as server and storage administration requires specialized skill sets and expertise, so does networking. These skill sets and expertise is rarely found in a DBA.
Facilities Admin
<ul style="list-style-type: none"> Responsible for racks, cables, transceivers, conduit, power, cooling, UPS, etc. DBAs do not typically have these skills either.

That’s a lot of trained, skilled, knowledgeable, experienced, and expert administrators in high demand. That number is compounded by the fact that each vendor’s database and different database type, requires specific to that database—skills, knowledge, training and therefore additional DBAs. Most IT organizations, even small ones, have several database instances. And they’ll generally have more than one database type as well. Each database instance and type will have a multiplier effect on the number of DBAs. This is especially true for cloud DBaaS providers such as AWS that have a different database type for just about every data type. And every one of them requires a separate learning curve, skills, knowledge, and expertise.

Obviously, experienced DBAs can manage more than one database at a time. But the number of databases a DBA can manage varies by database size, tools, database version, database type variances, DBA expertise, database formalization administration, and production versus non-production databases.

No matter how it’s sliced, there simply are not enough DBAs to be had. Especially at a reasonable cost. And the DBA market growth rate is nowhere near that of database instances.

Trends Exacerbating the DBA Shortage

From the late 1970s to the 2010s, the baby boomer generation dominated IT. Not anymore. Large numbers of the baby boomer generation DBAs are retiring from the workforce. That’s a huge amount of skill, knowledge, experience, and expertise walking out the door.

Compounding this loss is a disquieting reality that the technical schools, colleges, and universities are not coming close to filling the gap. It has long been known that Computer Science and Information Systems courses are rarely delivered by the best, brightest, most experienced, and expert personnel. Caused generally by the large disparity of pay, commercial markets provide the best and the brightest versus the low pay teaching in technical schools, colleges, and universities. The result are graduates that typically lack sufficient DBA skills. The IT organizations that hire them must train them and with the exodus of talent, that’s becoming increasingly difficult.

The shortage of DBAs and declining in-house expertise is made that much worse by the rapid growth of applications dependent on databases, such as AI machine learning, micro-services, and more. Databases are essential foundational software infrastructure. Implementing, operating, managing, and maintaining them are more critical than ever.

DBA Shortage Consequences

This disturbing combination of trends is causing IT organizations serious heartburn. The economics of supply and demand have dramatically escalated DBA costs at the same time skills and expertise are decreasing resulting in less effective databases, non-optimal applications, and frustrating performance.

The overworked and highly stressed DBAs, especially those with reduced experience, skills, training, and expertise, results in more human errors, redos, increased outages, poor response time, frustrated users, unhappy customers, and lost revenue.

It is these conditions that lead many Enterprises to consider Database IT outsourcing MSPs.

Why Database IT Outsourcing Comes Up Short

Database IT outsourcing is not a panacea. It’s analogous to putting a Band-Aid on a hemorrhage. It definitely helps, but doesn’t come close to stopping the bleeding. The MSP Database IT outsourcer potentially has

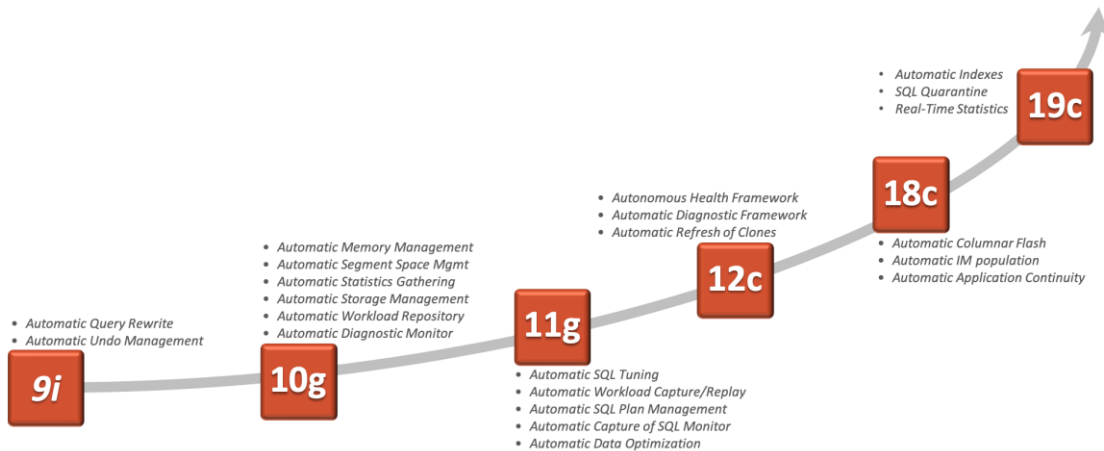


more experienced people; however, there is no guarantee of that. They are facing the same DBA shortage problems as everyone else. Assuming their DBAs are in fact more expert and experienced, then they have the ability to be a bit nimbler and faster at implementation, management, operations, maintenance, and problem solving. That depends on their DBAs not being stretched too thin by the demands of too many database instances and from too many outsourced clients.

Database IT outsourcing is at its root a human-based service. Humans do not scale well. Database IT outsourcing is moving the workload from overloaded IT organization DBAs to the overloaded MSP DBAs providing the outsourcing service. At its root, human manual labor-intensity does not change. There are likely to be nearly as many human errors and potentially more.

Database IT outsourcing is primarily partial outsourcing of the tactical manually labor-intensive DBA functions. Its focus is on outsourcing general DBA routine administration work. Database IT outsourcing most frequently targets monitoring, troubleshooting, root cause analysis, fixing, patching, data protection, and testing. The rest is generally left up to the client IT organization. The result is increased costs for limited benefits. It's also why when it comes to Database IT outsourcing, according to the Computer Economics report², few Enterprises have whole heartedly embraced it. However, many feel it may become necessary as the DBA shortage grows worse.

Oracle recognized the DBA shortage problem several years ago and set out to solve it. For the past 20 years since the release of Oracle Database 9i, Oracle has been continuously automating database technology within the Oracle Database. Its flagship Autonomous Database is the culmination of those efforts.



The Oracle Autonomous Database Solution

Overview

The Oracle Autonomous Database (ADB) is an industry unique public cloud-based database-as-a-service (DBaaS) that puts more than five decades of human DBA expertise into the DBaaS. Doing this solves two problems. The first major problem it solves is the shortage of knowledgeable DBAs. When the DBaaS service has the built-in expertise and autonomy, it no longer compels the DBA to be adept at manually labor-intensive database administration. In fact, it doesn't require them to do that work at all. The second problem it solves is the dearth of DBA cycles available for strategic work. That's work focused on the business side with new projects and application development and optimization that focuses on decreasing costs while increasing productivity, revenues, and profits. When DBAs no longer have to waste significant time on foundational databases administration they can focus on the strategic priorities of the business.



² Computer Economics: DBA Outsourcing not Widely Embraced, but May Become Strategic

Defining the Autonomy Levels in an Autonomous Database

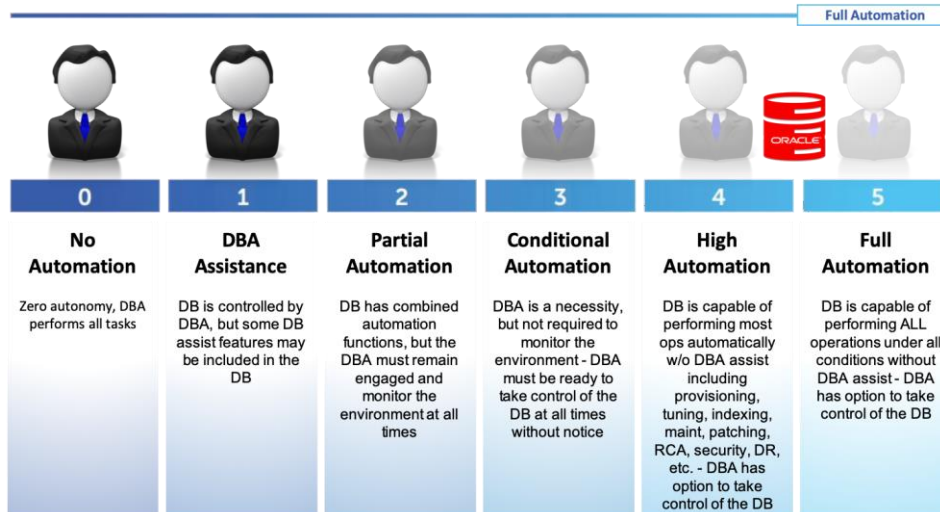


Chart 1: Database Autonomy Levels

There are six levels of database autonomy from level 0 through level 5 as illustrated in the chart above.

How The Oracle Autonomous Database Compares to the Market

Oracle initially release the Oracle Autonomous Database DBaaS in the fall of 2017. The current version has an autonomy level greater than 4 and close to the highest level 5. No other database or DBaaS comes close. The vast majority have an autonomy level of 0 or 1 at most. That includes every single AWS database: Aurora, Redshift, DynamoDB, ElastiCache for Memcached or Redis, DocumentDB, Keyspaces for Apache Cassandra, Neptune, Timestream, and QLDB. Microsoft too with SQL Server and Azure SQL Server; IBM DB2; Teradata and more. Even the much-hyped Snowflake Database is at best a level 2 autonomy. None of these databases or DBaaS, or any other database or DBaaS, come close to the Oracle Autonomous Database automation capabilities. Oracle Autonomous Database transcends the barriers of everything else on the market today.

There are two primary reasons why the aforementioned databases and DBaaS managed service providers (MSP) are so far behind the Oracle Autonomous Database:

1. Underlying database foundation
2. Hardware infrastructure foundation

Simple logic makes it clear, an autonomous database is only as good as the underlying database and hardware infrastructure upon which it's based on. A deeper look into the Oracle Autonomous Database makes it crystal clear why it's clearly superior.

Oracle Autonomous Database Underlying "Database" Foundation

The Oracle Autonomous Database DBaaS applies its extensive administration services automation on top of the very latest Oracle Database software. The Oracle Database has been and continues to be the Enterprise market leader by any measure per Gartner³ and DB-Engines⁴. It utilizes the very best Oracle Database practices with policy engines, decision trees, and advanced AI-machine learning algorithms to automate provisioning, clustering, disaster protection, tuning, indexing, elastic resource scaling up and down, patching, upgrading, and security.

The Oracle Autonomous Database DBaaS takes advantage of and automates many of the most advanced, popular, industry-leading, production-proven, built-in Oracle Database capabilities including:

- Exadata Database Machine
- Real Application Clusters (RAC)

³ 2019 Gartner Magic Quadrant for Operational Database Management Systems

⁴ db-engines ranking

- Active Data Guard (ADG)
- Flashback
- Recovery Manager (RMAN)
- Database Vault
- Multi-tenant Container Databases (CDB) with pluggable databases (PDB)
- Multi-concurrent Database types utilizing the same or different data – i.e. no data movement or ETLs required⁵
 - Relational
 - OLTP
 - OLAP
 - Data Warehouse
 - Object
 - XML
 - Key Value
 - JSON
 - Document
 - Time Series
 - Graphic
 - Spatial
 - Blockchain
 - IoT
- Advanced AI machine learning algorithms

<p style="text-align: center;">Classification</p> <p>Naïve Bayes Logistic Regression (GLM) Decision Tree Random Forest Neural Network Support Vector Machine (SVM) Explicit Semantic Analysis</p>	<p style="text-align: center;">Regression</p> <p>Linear Model Generalized Linear Model (GLM) Support Vector Machine (SVM) Stepwise Linear Regression Neural Network LASSO</p>	<p style="text-align: center;">Feature Extraction</p> <p>Principal Comp Analysis (PCA) Non-negative Matrix Factorization Singular Value Decomposition (SVD) Explicit Semantic Analysis (ESA)</p>
<p style="text-align: center;">Clustering</p> <p>Hierarchical K-Means Hierarchical O-Cluster Expectation Maximization (EM)</p>	<p style="text-align: center;">Attribute Importance</p> <p>Min Description Length Princippal Comp Analysis (PCA) Unsupervised Pair-Wise KL DIV CUR Decomposition for Row & AI</p>	<p style="text-align: center;">Text Analysis</p> <p>Sentiment Analysis Algorithms Support Text Tokenization & Theme Extraction Explicit Semantic Analysis (ESA) for Document Similarity</p>
<p style="text-align: center;">Time Series</p> <p>Forecasting - Exponential Smoothing Includes Popular Models e.g. Holt-Winters w/trends, seasonality, irregularity, missing data</p>	<p style="text-align: center;">SQL Analytics</p> <p>SQL Windows, Patterns, Aggregates</p>	<p style="text-align: center;">Statistical Functions</p> <p>Basic Statistics: min, max, median, stdev, t-test, F-test, Pearson's, Chi-Sq, Anova, etc.</p>
<p style="text-align: center;">Anomaly Detection</p> <p>1-Class SVM</p>	<p style="text-align: center;">Association Rules</p> <p>A Priori / Market Basket</p>	<p style="text-align: center;">R Packages</p> <p>3rd party R Packages through Embedded Execution Spark Mlib Algorithm Integration</p>
	<p style="text-align: center;">Predictive Queries</p> <p>Predict, Cluster, Detect, Features</p>	

Oracle Autonomous Database Underlying “Hardware” Foundation

The Oracle Autonomous Database is an Oracle cloud service that runs on the Oracle Exadata hardware platform. Exadata is co-engineered with the Oracle Database to deliver the lowest latencies, highest IOPS, and most throughput of any database platform anywhere in the world. Oracle has engineered unmatched performance, functionality, security, reliability, availability, and serviceability in this synergistic relationship. There are more than 70 unique Oracle Database net incremental capabilities only available on the Exadata platform. And only Oracle enables its database to offload computationally intensive processes from the Exadata database servers to the Exadata storage servers. And since every DBaaS service, including the Oracle Autonomous Database,



⁵ Most databases are a singular specialized type of database. That means each database requires its own specialized skills, knowledge, training, and expertise. The Oracle Database is the only database that incorporates every major database type into a single converged multi-tenant database with no data silos or separate database architectures to implement, operate, manage, and maintain.

prices based on vCPUs utilized over time (OCPUs in the Oracle Cloud), faster performance directly equates into lower cost. Time really is money.

All other databases and DBaaS run on commodity-off-the-shelf generic server and storage hardware. There is no co-engineering, automatic database performance optimization, synergistic database and hardware security, or database processor offload to storage. This is especially true when utilizing public cloud-based DBaaS. Commodity hardware and sharing that hardware with multiple applications and tenants is a key element in keeping the cloud service provider's overhead costs down. May be an adequate strategy for small databases. But it can become a serious problem as those databases scale and performance requirements increase. There are typically hard performance scalability limits per instance, requiring extensive database sharding rework to scale. Even those DBaaS that can scale usually have constrained performance limits. These limits are fixed performance per platform category. If the DBA guesses incorrectly, they have to move their database to another platform category. That move is not automatic.

The Oracle Autonomous Database has no hard limits. It can start very small and move to very large. The customer only pays for what they use when they use it. It elastically scales up and down. And because it's based on a converged database that supports every major database type, there are no multiple database subscriptions or duplicate costly data silos. There is no data fragmentation like an organization can encounter when running dozens of different isolated, specialized databases—each with its own security, patching and management requirements.

Oracle Autonomous Database Cloud Service

Oracle applies highly advanced automation machine-learning driven capabilities on top of these impressive foundations to create a unique cloud service called Oracle Autonomous Database. The extensive automation and autonomy makes it "self-driving, self-securing and self-repairing." Self-driving means all of the database and infrastructure management, monitoring,



indexing, and tuning processes are automated.

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

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

Self-repairing means database healing, again without DBA expertise or intervention thus providing a significant reduction in outage downtime from unplanned maintenance. The Oracle Autonomous Database typically has a customer experience of zero downtime for a normal database operational lifecycle including patching, and averages fewer than 2.5 minutes of downtime in any month experiencing unplanned outages. It is critical to note that Oracle does not put caveats around those numbers in its SLAs. There are none of the exceptions that are so prevalent in the Database Cloud industry. This translates into the Oracle Autonomous Database delivering the highest uptime in the industry.

The Oracle Autonomous Database delivers all of this unprecedented automation elastically and cost justifiably. Elastically means that while completely online, the Oracle Autonomous Database auto-scales up as well as down and customers only pay for actual CPU usage, or true "pay-for-usage." This is a huge boon to organizations that have substantial variability in their database usage. Cost justifiably means the Oracle Autonomous Database saves far more than it costs. It does this by enabling customers to focus their DBAs on generating more profitability with revenue producing and/or cost reducing applications.

For the Database IT outsourcer MSP, the Oracle Autonomous Database does two things. First, it solves their DBA shortage. Database IT outsourcers are equally affected by the increasing shortage of DBAs. Second, it

transforms their business by moving it up the stack for higher value and more profitable services. The MSP no longer has to waste their time on low value database admin tasks including implementation, operations, management, tuning, patching, troubleshooting, and maintenance. They can focus their expertise for their clients on high value strategic projects that leverages the Oracle Autonomous Database.

Oracle Autonomous Database Tuned Services

The Oracle Autonomous Database has been tuned for three different cloud database services: Autonomous Transaction Processing, Autonomous Data Warehouse, and Autonomous JSON Database. All three are all based on the Oracle Database and run on the extremely high performance Exadata platform. All three are available in the Oracle Cloud Infrastructure, Oracle Exadata Gen2 Cloud@Customer, and Oracle Dedicated Region Cloud@Customer.

Autonomous Transaction Processing Database (ATP)

The Oracle Autonomous Transaction Processing is a cloud database tuned for best possible transactional performance. It eliminates the complexity of operating and securing high-performance transactional databases.

The Oracle Autonomous Transaction Processing cloud database service automates provisioning, configuring, tuning, scaling, patching, encrypting, and repairing of databases. The Oracle Autonomous Transaction Processing cloud database includes all of Oracle's advanced database options, such as real application clusters (RAC), multitenant, partitioning, in-memory, advanced security, and advanced compression. This Oracle cloud database service is built to support everything from simple web apps to large and sophisticated applications that are critical for business operation. And it has low pay-per-use pricing. The Oracle Autonomous Transaction Processing cloud database service is available in the Oracle public cloud infrastructure, Oracle Exadata Cloud@Customer, and Oracle Dedicate Region Cloud@Customer.


Autonomous Data Warehouse (ADW)

The Oracle Autonomous Data Warehouse is a cloud data warehouse cloud database service tuned for best possible analytics and throughput performance. It fundamentally eliminates the complex operations of a data warehouse while securing the data and enabling much faster development of data-driven applications.

The Oracle Autonomous Data Warehouse automates provisioning, configuring, securing, tuning, scaling, patching, backing up, and repairing of the data warehouse. It is unlike any other fully managed DBaaS warehouse solutions, which only patch and update the service. The Oracle Autonomous Data Warehouse features elastic, automated scaling, performance tuning, security, a broad set of built-in converged database capabilities that enable simpler queries across multiple data types, machine learning analysis, simple data loading, data visualizations, and low pay-per-use pricing. The Oracle Autonomous Data Warehouse cloud database service is available in the Oracle public cloud infrastructure, Oracle Exadata Cloud@Customer, and Oracle Dedicate Region Cloud@Customer.

Autonomous JSON Database (AJD)

The Oracle Autonomous JSON Database is a new cloud document database service built for JSON-centric developers and tuned for the best possible document analytics processing. It is the latest Oracle Autonomous Database cloud service. The Oracle Autonomous JSON Database automates provisioning, configuring, tuning, scaling, patching, encrypting, repairing of databases, eliminating database management and delivering 99.995% availability. The Oracle JSON Database is available in Oracle Cloud Infrastructure, Oracle Exadata Cloud@Customer, and Oracle Dedicated Region Cloud@Customer.



The Oracle Autonomous JSON Database makes it incredibly simple for developers to create JSON-centric applications. It features simple document APIs, serverless scaling, high performance ACID transactions, comprehensive security, and like all Oracle Autonomous Database cloud services, low pay-per-use pricing. Each of Oracle’s Autonomous Database cloud services are vastly superior to other DBaaS offerings. The Oracle Autonomous Transaction Processing Database has 50x lower IO latency than AWS RDS, irrespective of the relational database utilized, and up to 100x lower IO latency than Azure SQL. And it’s orders of magnitude more automated. The Oracle Autonomous Data Warehouse has at least a 2X performance advantage compared to Snowflake or AWS Redshift. And again, there is no comparison on the autonomy. Oracle’s latest Autonomous JSON Database blows away the MongoDB Atlas service. Since this is Oracle’s most recent service, a deeper look illustrates the clear advantages.

Unlike the Oracle Autonomous JSON Database, the MongoDB Atlas service does not have its own public cloud infrastructure. MongoDB Atlas runs on generic AWS infrastructure, meaning there are no special optimizations for MongoDB, no source-code level engineering like what exists between Oracle Autonomous Database, Exadata and OCI. MongoDB is just another database on AWS, whereas Autonomous JSON is part of the core Oracle Database on OCI. That should matter to developers because it will make their lives easier. While the biggest difference is that Oracle Autonomous JSON Database cloud database service is *autonomous* and MongoDB is not, there are several other major Oracle Autonomous JSON Database advantages:


Capabilities	Oracle Autonomous JSON Database	MongoDB Atlas
Max Document Size	32 MB	16 MB
Max nested depth of documents	1024 levels	100 levels
Indexes per collection	unlimited	64
Compound index fields	unlimited	32
Full document index	JSON Search Index	∅
Server-side functions	Functions, procedures, triggers	Not recommended*
Multi-document transactions	Always ACID	ACID only upon request via explicit API calls
Transaction duration	unlimited	60 seconds default
Transaction size	unlimited	maximum of 1000 documents*
Aggregation data size	unlimited	100 MB RAM + explicit allowDiskUse param
Serverless auto-scaling	✓	∅
SQL access over JSON documents	✓	∅
Comprehensive security (e.g. Virtual Private DB, Data Redaction, Custom DB Roles)	✓	∅
Low Price (8 OCPUs for Oracle - M60 for MongoDB Atlas)	\$2.74 / hour	\$3.95 / hour

*Recommendations as per MongoDB documentation: [link1](#), [link2](#)

Contrary to common wisdom, MongoDB Atlas costs considerably more than the Oracle Autonomous JSON Database. In fact, MongoDB Atlas costs > 44% more per hour (\$3.95 vs \$2.74 USD) based on the MongoDB Atlas dedicated cluster at M60 tier which is comparable to 8 OCPUs for Oracle Autonomous JSON Database.

The lower cost continues beyond the cost/hour. The Oracle Autonomous JSON Database is completely elastic and serverless like all of the Oracle Autonomous Database services. That’s not the case with MongoDB Atlas which relies on fixed hardware shapes. If the MongoDB Atlas needs a bigger shape for unexpected spikes, it must be disrupted and migrated to the larger shape. The Oracle Autonomous JSON Database automatically scales up for more processing when needed and down when no longer needed. And it does so while online. No downtime or disruption required. It additionally includes backup and simple connectivity to business intelligence (BI) tools that costs more for MongoDB Atlas.

Then there’s the performance. Remember, Oracle Autonomous Database services run on Exadata. Exadata is engineered to deliver the best possible performance for Oracle Databases. This is also true for the Oracle Autonomous JSON Database. Based on the same cost configurations above and using the industry-standard [YCSB](#) benchmark for JSON databases, the comparison below demonstrates that the Oracle Autonomous JSON Database delivers approximately 2 – 3x better performance compared to MongoDB Atlas results published by MongoDB ([here](#)), as seen in the charts below.



Autonomous JSON Database with 8 OCPUs compared to MongoDB Atlas on M60
 Industry-standard Yahoo Cloud Serving Benchmark (YCSB)
 Source of MongoDB results: [Mongo YCSB test results](#)

Autonomous JSON processing 2-3 times faster than MongoDB Atlas means much less time to complete a job. Less time processing equals less money spent. In the cloud, time truly is money.

JSON developers spend the majority of their time not even coding. Instead they are grappling with issues such as bug fixes, security concerns and storage requirements. Oracle Autonomous Database can solve all of this and help simplify the lives of developers so they can execute more code and spend no time on the unproductive and boring tasks that make up a developers’ life.

Because the Oracle Autonomous Database services are based on a converged, multi-modal database, it addresses the time sink of spending countless cycles haggling with half a dozen APIs from several different databases. A single API experience empowers developers to write more useful code instead of writing complex separate API calls and dealing with the difficult problem of data fragmentation. Simple and convenient translates to more coding and less stress. Autonomous JSON Database takes care of all the grunt work and makes “more work, less talk,” a tangible reality for developers.

For every Oracle Autonomous Database cloud service, the unparalleled high performance achieved on OCI reduces costs significantly over all other DBaaS offerings. Competitive database suppliers have yet to come out with anything remotely comparable to any of these Oracle Autonomous Database services.

One more thing. Oracle offers a free cloud tier in the Oracle public cloud for each of the Autonomous Database services. This enables customers to try out the service without any financial obligation.

Conclusion

Databases are labor-intensive, require much expertise, for the error prone tasks of indexing, re-indexing, tuning, monitoring, troubleshooting (root cause analysis), patching, fixing, securing, backing up, and upgrading the database. Experienced, skilled, and expert DBAs are expensive. The severe and growing DBA shortage is getting worse over time, as reported by many Enterprise accounts. This is compounded by the increase in data, database instances, and database types with each requiring at least one DBA and frequently more. The retirement of the baby boomer generation plus insufficiently trained DBAs coming out of school is making the shortage worse.

Database IT outsourcing has proven to be less than a satisfactory answer. It offloads some of the DBA tasks, but not enough, while costing far more than it saves.

The Oracle Autonomous Database with tuned services for Autonomous Transaction Processing, Autonomous Data Warehouse, and Autonomous JSON Database is the first and currently only database cloud in the market that builds DBA expertise, knowledge, skills, automation, machine learning, and machine action into the database instead of requiring it in the DBA. It automates the vast majority of the database manual labor-intensive tasks for implementation, operations, management, and maintenance by making the database self-driving, self-securing, and self-healing. And it does all of this cost justifiably. Thus, it solves the immediate and long-term DBA shortage issue. In so doing, it empowers DBAs and database IT outsourcers to focus on the strategic aspects of using the Oracle Autonomous Database to reduce costs, increase revenues, and profits.

In summary, all three of the Oracle Autonomous Database cloud services solve the DBA shortage for both IT organizations and DBA outsourcing MSPs. It transforms the DBA role into one that is strategic. And it does so with more automation and autonomy than any other database cloud service. Most importantly, the Oracle Autonomous Database services do so at a surprisingly affordable cost that's lower than its non-autonomous competition.

For More Information on the Oracle Autonomous Database Services

Go to: [Oracle Autonomous Database](#)

Go to: [Oracle Autonomous Transaction Processing Database](#)

Go to: [Oracle Autonomous Data Warehouse Database](#)

Go to: [Oracle Autonomous JSON Database](#)

Paper sponsored by Oracle. **About DSC:** Marc Staimer, as President and CDS of the 22-year-old DSC Analyst Firm in Beaverton, OR, is well known for his in-depth and keen understanding of user problems, especially with storage, networking, applications, cloud services, data protection, and virtualization. Marc has published thousands of technology articles and tips from the user perspective for internationally renowned online trades including many of TechTarget's Searchxxx.com websites and Network Computing and GigaOM. Marc has additionally delivered hundreds of white papers, webinars, and seminars to many well-known industry giants such as: Brocade, Cisco, DELL, EMC, Emulex (Avago), HDS, HPE, LSI (Avago), Mellanox, NEC, NetApp, Oracle, QLogic, SanDisk, and Western Digital. He has additionally provided similar services to smaller, less well-known vendors/startups including: Asigra, Cloudfenna, Clustrix, Conduvis, DH2i, Diablo, FalconStor, Gridstore, ioFABRIC, Nexenta, Neuxpower, NetEx, NoviFlow, Pavilion Data, Permabit, Qumulo, SBDS, StorONE, Tegile, and many more. His speaking engagements are always well attended, often standing room only, because of the pragmatic, immediately useful information provided. Marc can be reached at marcstaimer@me.com, (503)-312-2167, in Beaverton OR, 97007.