Cloud-First Enterprise

The Time is Now

DATOMETRY

EXECUTIVE SERIES

Summer 2018

About This Publication

This series of articles by Mike Waas, Founder & CEO of Datometry, Inc. examines the critical aspects every IT executive needs to be aware of when planning a cloud-first strategy for the enterprise and its data assets.

Foreword

Cloud databases are about to change one of the most important IT markets. Analysts predict within the next decade the vast majority of database installations, currently a USD 40 billion market, will be replatformed to the cloud. In the process, IT leaders will have to re-evaluate their current systems and determine suitable equivalents in the cloud. Data warehouses, widely considered the crown jewels of the enterprise, are at the heart of this tectonic shift.

In this series of articles, I discuss the most important concerns and questions we frequently come across in our work with leading enterprise IT organizations in both Fortune 500 companies as well as mid-level enterprises. These are questions on every IT executive's mind when considering moving data assets to the cloud.

I also investigate common misconceptions around database migration and look at the complexities involved in replatforming data warehouse applications. Finally, I present Adaptive Data Virtualization, a new virtualization paradigm that my team and I have developed in response to working with the stewards of enterprise data over the years.

Michael Waas
Founder & CEO
Datometry, Inc.

“The key challenge [of] a cloud-first strategy is how to shift decades of on-premise data management to the cloud without the risk, expense, and time typically required for such migration projects.”

Table of Contents

Cloud Computing Goes Mainstream—At Last

Well over 10 years ago, Amazon unveiled AWS–its public cloud computing platform–around the same time as the initial release of Microsoft Azure as well as numerous, since-forgotten smaller companies. The long talked-about cloud computing revolution is finally seeing traction with large enterprises as even conservative IT adopters are embracing the public cloud for their core IT infrastructure.

Cloud computing is currently projected to grow by about 20% annually over the coming years and the steep increases in cloud revenue recently reported by leading Cloud Service Providers (CSPs) Amazon, Microsoft, and Google are proof that the cloud revolution is finally underway.

What is remarkable is the fact that despite the striking economics of the cloud offerings, it has taken the better part of a decade for the public cloud to develop significant traction with Fortune 500 companies.

Rigorous Security and Privacy in the Cloud

Cloud security—or rather, the perceived lack thereof—has been the single biggest obstacle that has hampered cloud adoption since the early days of the cloud. Curiously enough, there have been no major security breaches of general cloud infrastructure that would substantiate such fears but the specter of a potential breach at the infrastructure level has been enough to keep large enterprises away from the public cloud.

On the flip-side, security risks have made CSPs invest massively in securing their platforms over the past years. And, it is an investment that is now paying big dividends. Security of any major public cloud is arguably much more advanced and more up-to-date than any on-premises security with CSPs employing some of the world's top security experts working 24x7 on protecting their platforms; in fact, no breach in recent memory can be attributed to security flaws of cloud platforms.

Similarly, privacy concerns are no longer an obstacle either: CSPs are fully accommodating regulations, such as data not leaving a country or a certain geography, with multi-geo offerings that are implemented through an ever-growing network of globally located data centers.

Maturity and Completeness of Cloud Offerings

Cloud offerings have come a long way from just being Infrastructure-as-a-Service (IaaS) offerings primarily composed of storage and compute resources. Modern clouds are rich platforms that offer a myriad of services ready to be consumed on a pay-as-you-go basis. Not only are these offerings an incredible convenience and accelerate application development but they also relieve enterprise IT organizations of technically demanding deployment and operations challenges.

For complex systems such as databases, the CSPs have taken decidedly different approaches bringing various differentiating innovations to market over the past years. However, even in this space, further consolidation and assimilation should be expected, much to the benefit of IT consumers.

> **"Cloud security—or rather, the perceived lack thereof—has been the single biggest obstacle that has hampered cloud adoption since the early days of the cloud."**

In addition, positive news holds for availability and general robustness of cloud platforms. The times of significant outages and downtime—not uncommon in the early days of public cloud—are long gone. The cloud has reached a level of enterprise-readiness and maturity that is well on par with on-premises installations, if not already exceeding it.

Cloud Economics for Enterprise IT

Initially, the main and obvious value proposition of the public cloud was that of reducing an enterprise's Capital Expenditure (CAPEX), and CSPs positioned their offerings as an easy way to shed hardware and move from a model of owning equipment to one of merely renting it. This still required enterprise customers to employ administrators and IT staff and left them with some of their Operating Expense (OPEX).

> **All of enterprise IT is moving from a world in which customers own and operate technology into a world in which they use only APIs, and the underlying technology becomes transparent and its implementation irrelevant.**

This picture is changing rapidly: with richer services, enterprise IT is moving from a world in which customers own and operate technology into a world in which they use only APIs and the underlying technology becomes transparent and its implementation irrelevant. This phenomenon is turning traditional IT on its head.

The typical daily irritations between IT and its business users are now resolved: the CSPs refresh hardware constantly, customers always have the latest software releases at their disposal, and the elastic flexibility of cloud systems eliminates the need for costly over-provisioning. Pay-as-you-go becomes the new normal and classic OPEX is drastically reduced.

The ultimate trigger for enterprises to make the move, however, is software licenses about to expire and upcoming floor-sweep renewals of expensive, large-scale hardware components. These are the ultimate compelling events—completely orthogonal to the above developments and benefits of the cloud as such—that have resulted in enterprise IT leaders taking the plunge to the cloud. Not only have the benefits of the cloud have become too obvious to be ignored, but the penalty for not adopting the cloud has become too severe.

Virtually all enterprises, not just the classic early adopters or technologists, have formulated a cloud-first strategy over the past 12 months and have begun their cloud-native implementations. The industry has truly reached the tipping point.

The time for cloud is now.

Key Takeaways

- Security and privacy in the cloud have surpassed conventional standards
- Maturity and completeness of cloud offerings creates rich IT environments
- Economic penalty of not moving to the cloud increases rapidly

Cloud Databases and The Innovator's Dilemma

Data management in the cloud promises to be the most significant innovation in the database field in decades. IT leaders across all verticals are currently evaluating the various cloud database offerings but find it difficult to assess this new technology with conventional methods at their disposal.

This article reviews the market dynamics and presents the argument that cloud databases exhibit the very specific characteristic of a truly disruptive innovation, face the same obstacles as other disruptors, and predict their ability to take substantial market share from the incumbents over the long run.

The Concept of Disruptive Innovation

High-tech experts and investors have used *The Innovator's Dilemma*, authored by Harvard economist Clayton Christensen some 20 years ago, for identifying truly disruptive technology. The framework explains how established vendors get disrupted by upstarts who often innovate on a shoestring budget in ways that seem out of reach for established incumbents. And, once a technology reaches a tipping point, it disrupts the established market and its products—not gradually—but radically and at rapid pace.

The framework identifies two vital ingredients as the key markers of a successful disruptive innovation:

Disruption on price. The dominant, incumbent technology is expensive and requires a highly specialized skill set to operate, and is, therefore, only accessible to a limited upmarket clientele.

Disruption along a new dimension. The innovator does not improve on the incumbent product but innovates along a dimension initially not viewed as critical, or even seemingly detrimental to success.

Why Some Disruptors Succeed—and Others Fail

There are no shortcuts in this process as exemplified by cautionary tales of failed attempted disruptions, such as HP's Kittyhawk, a hard disk drive specifically engineered for disruption. Its creators tried to cut short the period for which a true disruptor typically enters the market as underdog at first. However, by attempting to improve on the incumbents' performance and simultaneously miniaturize the form factor drastically, they created a product with no market at all—and a text book case study for how not to disrupt, now widely taught in business school.

> **Cloud data warehouses have the markers of a truly disruptive technology, both from the technology and market perspective.**

There is an important lesson to be learned that many engineers find counter-intuitive: for disruptors, striving for perfection is one of the highest barriers to entry; rather, innovators are advised to go with a Minimum Viable Product first and let the market dynamics shape the product going forward.

It is important to note that while the framework of the Innovator's Dilemma cannot reliably predict what innovations will disrupt an industry, it neatly points out the weakness of incumbents and how they may get disrupted.

Cloud Databases are the First Real Innovation in the Database Field in Decades

There has been no dearth of innovation in the database field over the past four decades but none of the innovations have been able to truly disrupt the database market. For example, database innovations, such as object-orientation, temporal databases, and unstructured data models have all failed to deliver on their original promise to replace the relational database. These innovations ended up becoming features of the commercially successful incumbents.

Why did they fail to disrupt? It is because the innovations were actually enhancements of an already very powerful product: faster, more scalable and cheaper but lacked to innovate along a new dimension.

In contrast, cloud databases are truly radical and novel:

- They reduce databases to the one singular aspect customers have always appreciated about databases: a query API with a powerful query language.
- They remove the burdensome and exceedingly expensive aspect of having to operate one's own database.

To use a metaphor from the consumer space, cloud databases might just become what car sharing has become to buying luxury sedans: Instead of owning an expensive vehicle with regular maintenance requirements, a new generation of buyers is turning to shared cars with pay-as-you-go billing that gives them access to the latest models, configurations, sizes, and capabilities depending on need and occasion, and always ready, washed and gassed up—convenient and economical.

Cloud databases too offer the latest in hardware advances, software upgrades and bug fixes, security updates, and fully managed services including maintenance operations. Customers need not bother about backups and other much-loathed activities.

How do cloud databases measure up for disruptive potential? In contrast to earlier, incremental but not disruptive innovations, cloud databases appear to have significant potential:

Cloud data warehouses disrupt on price. Their pricing, often paired with go-as-you-pay billing, undercuts traditional offerings by an order of magnitude or more.

Cloud data warehouses disrupt along new dimensions. Instead of improved speeds-and-feeds, they offer a managed services model. They entered the market decidedly as underdogs regarding performance and scalability but are aggressively catching up with the incumbents.

The above points suggest that cloud data warehouses have the markers of a truly disruptive technology, both from the technology and market perspective. Naturally, skeptics and detractors continue to point out the lack of maturity of cloud databases on performance and scalability as a barrier to adoption. Ironically, this too is a marker of the disruptive potential of cloud databases. The enormous investments cloud data warehouse vendors like Microsoft and Amazon.com have been making over the past decade or so are bearing fruit, including improvements to performance and scalability, which suggests reaching the tipping point is imminent.

> **Cloud databases are truly radical as they reduce databases to a query API with a powerful query language and offer a managed services model.**

Key Takeaways

- Disruptive innovation starts with products that cannot compete with incumbent technology at first
- Attempts to innovate and beat incumbents at the same time are fraught with risk
- Cloud databases are the first real innovation in the database field in decades

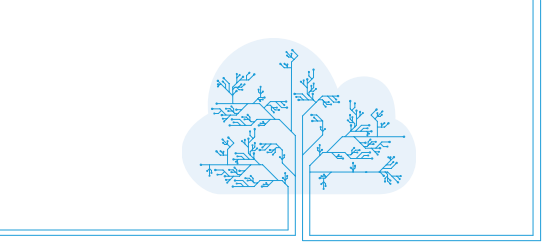

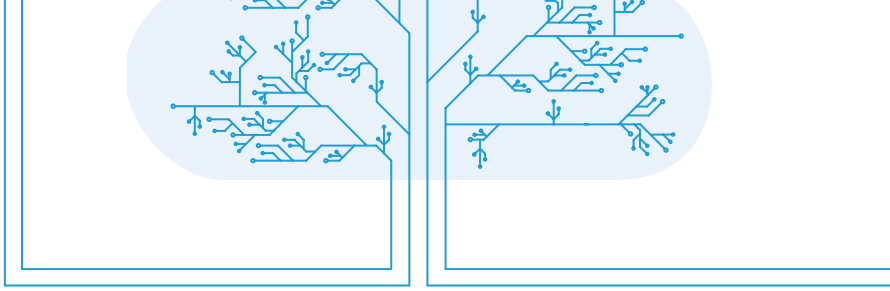

Cloud Databases Shed Operations and Maintenance

Asking any enterprise IT leader about their challenges and satisfaction with their database operations could generate a long litany of complaints big and small. And, while the value of databases is undisputed, the cost of maintaining appropriate service levels for the business, frequent support cases, and even just standard operations weigh heavily on any enterprise.

Cloud databases appear to hold the key to better future: by reducing databases to an API, they let enterprises shed the operations, support, and maintenance burden altogether. Indeed, cloud databases are requiring IT leaders to rethink their database ecosystem and, if approached in the right way, offer fundamental benefits that traditional database environments have never been able to deliver.

Cloud Databases: All the Benefits, None of the Hassles

Cloud databases are redefining database technology with only a powerful query interface, data import and export capabilities. Tasks typically associated with operating databases, such as backup and restore, storage administration, and a variety of activities needed to keep on the lights 24x7 are now wrapped in Software-as-a-Service (SaaS), or more specifically Database-as-a-Service (DBaaS) offerings. The database can now be taken care of as part of a fully managed service and no longer the enterprise's problem. Not only does this relieve enterprises from having to handle some of the most onerous database chores but reduces database operations to a minimum, driving down CAPEX and OPEX to the great benefit of the enterprise.

Effectively, enterprises are trading ownership of a highly complex hardware and software stack for the simplicity and convenience of a powerful API. This is a daunting proposition for enterprises and vendors alike and a true quantum leap for the industry. The shift from owning technology to just using APIs-on-demand is almost like giving up car ownership for a car sharing service where maintenance, insurance, and even the fuel for the vehicle is included with the plan, all while enjoying extreme flexibility and convenience at low hourly rates. Using APIs on demand may appear to be an unusual idea at first, but only until the overwhelming benefits become clear.

The advantages of the latest cloud technology also include removing the need for database performance tuning and query optimization.

Cloud Technology Upends the Database Market

The cloud database market is evolving rapidly as all the major Cloud Service Providers (CSPs)—Microsoft, Amazon, and Google—have been making substantial investments in database research and have built out sizable development teams for their cloud data products in recent years. Smaller, independent 3rd party database vendors have entered the race too.

> **Cloud databases are requiring IT leaders to rethink their database ecosystem.**

It is important to note that the traditional database market is dominated by a handful of vendors and every one of the products have been on the market for decades. What this means is that the current disruption of the traditional database industry is happening in a highly mature industry and, as such, the bar for the disruptors is high. But, the prize is clearly worth the effort: within the next decade, the entire database market is expected to be completely replatformed to the cloud. That is, a USD 40 billion IT market segment is up for grabs offering a once-in-a-lifetime opportunity for vendors to acquire market share that has been previously been locked down by competitors for decades.

For enterprises, this is an important and extremely favorable trend that ensures vendor commitment with highly competitive product development and pricing strategies. This also means that product deficiencies—in particular, performance

and scalability—are top priorities for vendors.

Own the Database, Own the Business

For CSPs, the stakes are even higher: the decision which cloud database to adopt automatically anchors an enterprise with a specific CSP. Unlike compute or storage that can be abstracted well through various Platform-as-a-Service (PaaS) mechanisms and are, therefore, easy to move between clouds, databases impose high switching cost and quickly become the center of gravity in any cloud strategy. This means acquiring customers early on in their journey to the cloud is paramount for CSPs.

An increasingly richer cloud services environment calls for a strong virtualization and abstraction of database APIs to enable applications—existing or new—to interoperate with different databases.

In this new paradigm of cloud shift, enterprises and IT leaders are now more than just stewards of the database infrastructure as they set direction and agenda for a critically important agenda of the business: control of data, their most valuable asset. And, cloud databases are developing into the most important investment any enterprise will make within the next few years.

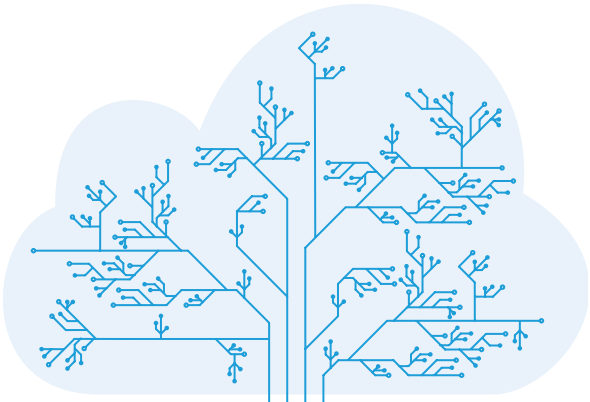

An increasingly richer cloud services environment calls for a strong virtualization and abstraction of database APIs to enable applications—existing or new—to interoperate with different databases.

Key Takeaways

- Cloud databases reduce complex technology to simple-to-use APIs
- The database market is undergoing a complete replatforming
- Databases are critical element in defining an enterprise's cloud strategy

Cloud Databases: The Advantage of No More Performance Tuning

For DBAs and enterprise IT, one of the key challenges is the tuning of databases and query workloads. It has become a core IT discipline in its own right over the past decades and represents significant OPEX in every IT budget.

With the advent of modern cloud-native databases, this situation is about to change. The need for tuning and optimizing queries is fast becoming unnecessary as elasticity and flexibility of these systems enable DBAs to compensate and remedy performance issues effectively at low cost.

Discipline of Database Tuning and Workload Optimization

One of the many artifacts of the database industry developed over the decades is the curious discipline of tuning databases and optimizing workloads or individual queries with the aim of improving their performance. DBAs spend countless hours on analyzing of workloads, understanding workload patterns, tinkering with configurations, and troubleshooting the system. This kind of tuning encompasses a wide range of activities, most notably:

Monitor and observe. Determine patterns in usage of the system and anticipate resource bottlenecks, then put in place strategies to mitigate problem situations.

Optimize workloads. Analyze and modify queries or workloads to ensure an ever-increasing query load combined with a gradual or even rapid increase in data volume can be processed within the available time.

Size and plan hardware resources. Understand when it is time to address performance issues with new hardware by determining the trade-off between peak performance needed and how much under-utilization can be tolerated during off-hours.

Query and database tuning have become an integral part of any database operations and gained notoriety for being an extremely hard-to-solve problem.

Challenges in Database Tuning and Query Optimization

The typical starting point for database tuning exercises is a slow query or workload, in short, a business problem that needs solving in a fixed time, specifically a shorter time than it currently takes.

The crux of the problem is that DBAs have a fixed hardware footprint at their disposal, that is, servers with a given number of cores, a limited amount of memory, and a predefined I/O subsystem, such as DAS or SAN. To solve the original business problem in a specific amount of time, DBAs need to be creative and be able to squeeze every last iota of slack from the database workloads, tune up the physical database design, and, if all else fails, persuade their business users to modify their original queries. Each of these tasks is laborious and requires DBAs to have deep expertise in a wide set of IT disciplines, not to mention the negotiation skills needed to convince users to change long-standing business processes.

This necessary and frequent tuning over time makes for highly unfavorable economics for the enterprise due to the following factors:

- DBA hours are costly and constantly modifying the system causes frequent disruptions to the business.
- The resulting system becomes increasingly complex making future tuning increasingly harder.

> **The elasticity of cloud databases is resulting in a new economic model for enterprises in which resources are only being provisioned and used as needed and excess capacities can be shed quickly.**

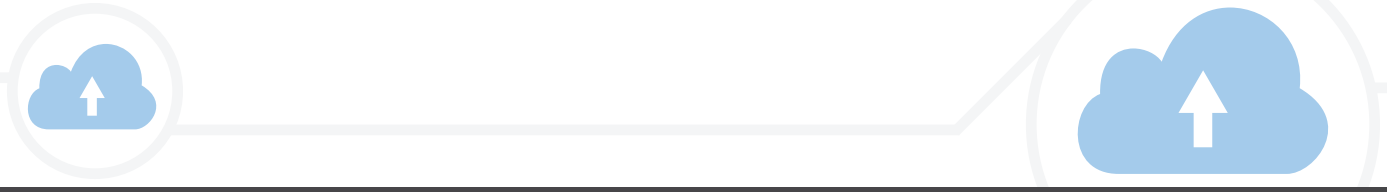

- Applications optimized this way are no longer portable between database systems and cement vendor lock-in to the current platform.

What is more, tuning is subject to diminishing returns on investment: optimizing for the last percentage point is exponentially more expensive than using broad strokes to solve the very basic problems.

> **Moving to a cloud database has just become orders of magnitude easier: there is no need for millisecond-parity between the old and new databases.**

Cloud-Native Databases Make Database Tuning Unnecessary

The elasticity of cloud-native databases is bringing a completely different and new dynamic to IT: DBAs can rapidly add and remove resources on demand.

The elasticity of cloud-native databases is resulting in a new economic model for enterprises in which resources are only being provisioned and used as needed and excess capacities can be shed quickly. This lets DBAs take advantage of fluctuations in database workload patterns, for example, scale the system out for end-of-quarter reporting, and scale back down to a fraction of system size the day after, and so forth.

However, elasticity not only addresses the questions of seasonal system utilization. A much-overlooked aspect of elasticity is its impact on tuning. Instead of spending thousands of dollars in labor to make a query, say, 10% faster, DBAs can leverage the elasticity of cloud-native databases and simply scale out the system by a specific increment, such as an additional server, as needed.

The economics of this simple trick are striking. Instead of increasingly customizing the database, constantly tuning the physical design of the database, and tuning queries—operations that drive up complexity and TCO—the DBA can simply move the slider on the control panel, provision a few more units, and put performance and throughput problems to rest at once. To be clear, this does not eliminate the need for DBAs but simplifies their jobs as it takes care of the most challenging—and frustrating—problems in database maintenance.

Advantages to Adopting Cloud-Native Databases

The elasticity of cloud-native databases has immediate and drastic ramifications for enterprise IT moving or considering moving to the cloud:

- Selecting a cloud-native database is not a matter of speeds and feeds. Not fast enough? Scale out the database with a few clicks. This lets solution architects focus on the features they need and on how aligned a specific Cloud Service Provider is with the enterprise cloud strategy.
- Moving to a cloud-native database has just become orders of magnitude easier: there is no need for millisecond-parity between the old and new databases. The challenge of whether the new version of the database or the new database will be as fast as the old is no longer a consideration.

What this means for the DBA and the enterprise is the significantly less aggressive customization of applications to the underlying database, faster technology adoption cycles, and short time to value.

Key Takeaways

- Tuning of databases long seen as a necessary evil
- Elasticity of cloud databases greatly reduces need to tune individual queries
- Cloud databases increasingly engineered without any means for tuning at all

“Cloud databases are the first real innovation in the database field in decades.”

Database Migration Paradox: The Surprising Truth About Replatforming

With the on-premises segment of the database market hurtling toward its demise, enterprises of all sizes are moving their data and database assets to the cloud and have begun to face one of the biggest logistical IT challenge in decades: how to migrate databases to the cloud? Truly, database migration is one of the most daunting initiatives any CIO faces.

Paradoxically, the biggest challenge in database migrations is not the actual migrating of the database content but the replatforming of the applications that interoperate with the database. This is frequently underestimated, even by experienced and seasoned practitioners, so much so that Gartner estimates that 50% percent of all migrations fail, go over budget, or run late because of this paradox.

Given the fact that in the next 5 to 10 years, virtually every database instance—currently on premises—will have to be replatformed, it is imperative for CIOs and IT leaders to understand the challenges the Database Migration Paradox poses and how Datometry Adaptive Data Virtualization™ can be used to greatly reduce time and cost involved while significantly mitigating the enormous risk associated with migrations.

Anatomy of Database Migration Projects

Database migration projects are among the most elaborate and complex transformation projects that every enterprise IT department faces as companies are headed for the cloud.

Changing the database infrastructure in a Fortune 500 enterprise comes with a set of challenges, some typical and some specific to the enterprise. Conventional wisdom has it that database migrations take at least 12 months, but it is important to note that these projects can drag on for several years depending on the complexity of the system in question. For example, in the case of replatforming a data warehouse, considered the crown jewels of the modern, data-driven enterprise, a replatforming project takes typically three to five years.

In addition, such projects frequently involve teams of 100 or more personnel, often highly specialized consultants, for an extended period of time. Without taking opportunity cost into account, migration projects can quickly exceed USD 10–15 million, with larger projects running up USD 50 million or more.

Another critical aspect that determines the chances of success for such projects is the average tenure of a CIO which currently runs about 4.3 years. If not carefully planned and aligned, replatforming projects have the potential to become stumbling blocks in an executive's career.

In conclusion, the time, cost and risk of conventional migrations pose significant challenges to any enterprise: exceedingly long durations disrupt business processes and run the risk to lose their alignment with overall corporate strategy as projects extend over multiple years; pose a heightened risk to competitiveness of the enterprise; and can negatively

> **Modernization of database applications takes up 80% of the time, constitutes majority of the cost, and adds significant risk in replatforming projects.**

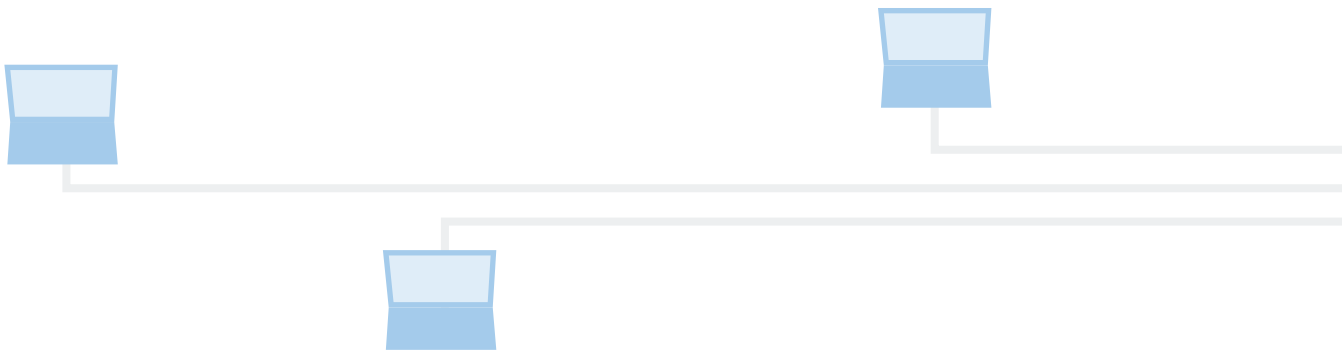

impact the careers of the decision makers. It should come as no surprise that IT leaders have historically avoided engaging in these projects.

Migration Paradox: Transferring Database Content is Easy, Modernizing the Applications is the Hard Part

A database replatforming initiative consists, broadly speaking, of three major components:

- Transfer and conversion of the schema
- Movement of the actual data
- Modernization of every application to work with the new database

Schema and data transfer together represent the migrating of the actual content of the database and are increasingly well supported by tools and utilities often provided free of charge by the database vendors in anticipation of the windfall they will experience if the enterprise replatforms to their product. For example, both Amazon.com and Microsoft advertise tools in this space: Database Migration Services and Azure Data Factory respectively.

In most cases, the transfer of the database content does not pose significant challenges although intricate corner cases can hamper the successful application of said tools, and, in such situations, manual intervention may be needed to make up for the shortcomings of these utilities.

However, what stumps IT practitioners routinely is the question how to modernize their applications. Many learn a painful lesson at this point: it is not the migrating of the database content that is the hard part of the migration after all; rather, the

> **A large Global 2000 Food Retailer is in its fifth year of rewriting business applications with no end in sight.**

modernization of applications is what takes up 80% of the time, constitutes the majority of the cost, and adds significant risk to the project. For example, a large Global 2000 Food Retailer is in its fifth year of rewriting business applications with no end in sight.

And, even more surprising, even third-party vendor applications that are often touted by their vendors as highly portable frequently become stumbling blocks as they are highly customized to make better use of the current database or use abstractions, such as stored procedures, that are intrinsically linked to the current database.

If one adds to this mix a decade or so of custom application development and its queries using system-specific syntax embedded in a wide variety of applications from maintenance jobs to Excel sheets authored by analysts far from the actual database, enterprises are often overwhelmed with the prospect of a database modernization.

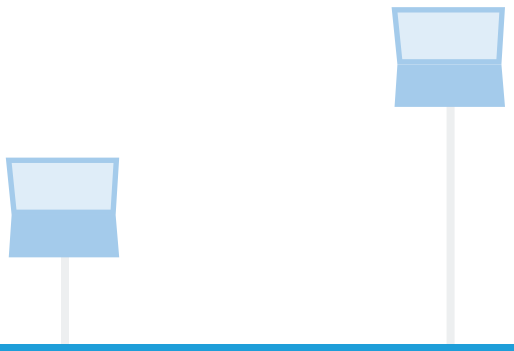

Key Takeaways

- Transferring content of databases accounts only for small fraction of cost in database migration
- Rewriting applications is the lion's share of expense and risk in a migration
- Need a new way of virtualizing the database to keep applications intact

Application Modernization: The Good, the Bad, and the Ugly

Enterprise IT leaders are facing enormous challenges in how best to modernize their database or data warehouse applications as enterprises are moving to the cloud and on-premises databases need to be replaced with cloud-native systems. Database application modernization is a logistical nightmare because it means having to modify potentially 1,000s of existing applications with each application posing specific challenges.

Understanding the Difficulty of Application Modernization

Application modernization, the process of making an application work with a database different from the one it was originally written or configured for, is really a misnomer: instead of modernizing and improving the functionality, application modernization is primarily about changing applications to work around functional discrepancies between the old and new database.

Consequently, database modernization projects actually burden applications with additional workarounds that inadvertently increase complexity and provide substantial room for error. This article examines scenarios of varying degrees of difficulty when modernizing SQL queries and categorizes them into three groups: the good, the bad, and the ugly.

The Good

A large group of syntactic modifications fall into this category and as they are easy to spot, IT practitioners are quick to offer rewrites. Most prominently, this category includes discrepancies in keywords, such as abbreviations, or workarounds for elegant shortcuts the new database does not offer.

Workarounds are typically straightforward and pose limited risk. Even though at first sight they might appear as pure textual differences, few can be dealt with using textual manipulation tools only, for example, Regular Expressions, a widely used tool often considered a programmer's Swiss army knife for text manipulation.

For queries in this category, the adjustments needed are of some local nature, that is, they can be made without a holistic understanding of the query.

The Bad

Queries in this category use powerful, often non-standard constructions that require elaborate rewrites and subject matter expertise when adjusting them for cloud databases. Examples include proprietary clauses like Teradata's QUALIFY or extensions that pre-date standards and, therefore, deviate from the language definition of virtually any other SQL dialect. A subtlety that has proven particularly error-prone is system-specific interpretations of standard clauses such as the ordering of NULL values relative to other data; the same query may return different results when executed on the on-premises database and its cloud counterpart, respectively.

These queries require a full and detailed semantic understanding of all components and changes are no longer locally contained but usually need a complete restructuring of the original query.

The Ugly

Finally, an even more complex class of database application queries are those that contain advanced—and, often, proprietary—control flow features, such as recursive queries or stored procedures. If the cloud database does not provide primitives of

> **Database modernization projects actually burden applications with additional workarounds that inadvertently increase complexity and provide substantial room for error.**

similar expressivity, a manual rewrite requires more than just an adjustment of the SQL text. Rather, the application code in which the SQL text is embedded needs to be modified to compensate for the lack of expressiveness.

In this case, even the surrounding application code is exposed to the risk of introducing software defects. The problem is usually compounded by the fact that the original application developers may no longer be with the company and the code may not have been touched in years.

> **Static modernization adapts existing business processes to the current version of the target database and ignores the fact that cloud databases are undergoing drastic and rapid development.**

Application Modernization is Inherently Counter-Productive

The challenges described above reflect the current state-of-the-art in an industry that is bracing for a tsunami of enterprise data management replatforming initiatives over the next decade. At the core of the problem is the notion of modifying database applications to make them work with a new database and executing the modifications as quickly as possible and in a way that causes the least disturbance to the business.

The modernization approach can be referred to as static as it attempts to make changes to applications in a stand-alone manner and is independent of the application behavior at run time. Static application modernization besides being inherently insufficient and impacting the business negatively, also has the following significant shortcomings:

- Static modernization adapts existing business processes to the current version of the target database and ignores the fact that cloud databases are undergoing drastic and very rapid development. Changes introduced to deal with any current limitations of the database will need to be revisited at a later point, effectively making another modernization necessary.
- Defects, inevitably introduced in any such large-scale project, proliferate insidiously in static modernizations as templates and clauses are widely copied across all affected applications.
- Enterprises effectively trade the vendor lock-in they have been seeking to escape, in some cases for decades already, for yet another vendor lock-in. Obviously, this is done in the hopes the new target platform will prove to be the right choice and no further migration is needed. Given the rapid market dynamics in the cloud, this may not be a sound risk mitigation strategy.

The effective alternative to the above challenges is Datometry Adaptive Data Virtualization technology which eliminates the need for conventional migration procedures.

Key Takeaways

- Application migration requires sophisticated modeling of queries to ensure correctness
- Static rewrites frequently compromise application code to compensate for missing features in target database
- Virtualization platform with real-time statement and data translation eliminates risk and cost

Virtualizing the Database: A New Paradigm

Cloud databases are about to redefine the database landscape in a fundamental way. They are cost-effective and easy to operate on one hand and offer unprecedented flexibility and scale on the other. One of the biggest challenges IT leaders face in the years ahead is the question how to replatform to this new kind of technology and move data assets to the cloud without any significant disruption to the business.

Datometry Adaptive Data Virtualization (ADV) is a next-generation virtualization technology which makes databases effectively interchangeable without code changes to the applications. It presents a new paradigm in data management, one that offers a radically different way to adopt cloud databases at a fraction of cost, time, and risk of conventional database migrations.

Indeed, it marks the beginning of a new era in data management enabling IT leaders to tap the competitive advantages afforded by cloud database technology.

The Global Database Market is Set for a Tectonic Shift Due to Cloud

In a recent research publication, Matt Aslett of 451 Research forecast the demise of the on-premises database market as we know it. According to his analysis, the USD 40+ billion on-premises database market is expected to contract to a fraction of its original size as cloud databases expand their market share dramatically and displace on-premises databases almost entirely.

Given the enormous interest and uptake cloud databases have seen in just the past 12 months, it is conceivable that this impending transformation of the IT landscape will actually happen even faster—once it reaches a tipping point—as databases are core to any enterprise infrastructure and, hence, the adoption of cloud databases is critical to a successful cloud strategy. An increasing commoditization of cloud database technology and with it the decline in price will likely increase the pace of this tectonic shift.

> **As per 451 Research, USD 40+ billion on-premises database market is expected to contract to a fraction of its original size as cloud databases expand their market share dramatically and displace their on-premises predecessors almost entirely.**

Adaptive Data Virtualization Accelerates Adoption of Cloud Databases

Datometry ADV technology allows applications originally written for the current database or data warehouse to run natively and in real-time on a different database. By acting as an abstraction of the underlying database, ADV eliminates switching cost and vendor lock-in while enabling enterprises to preserve long-standing investments in application development. The speed of adopting cloud databases with ADV means that enterprises can immediately unlock a wide array of benefits available in cloud databases.

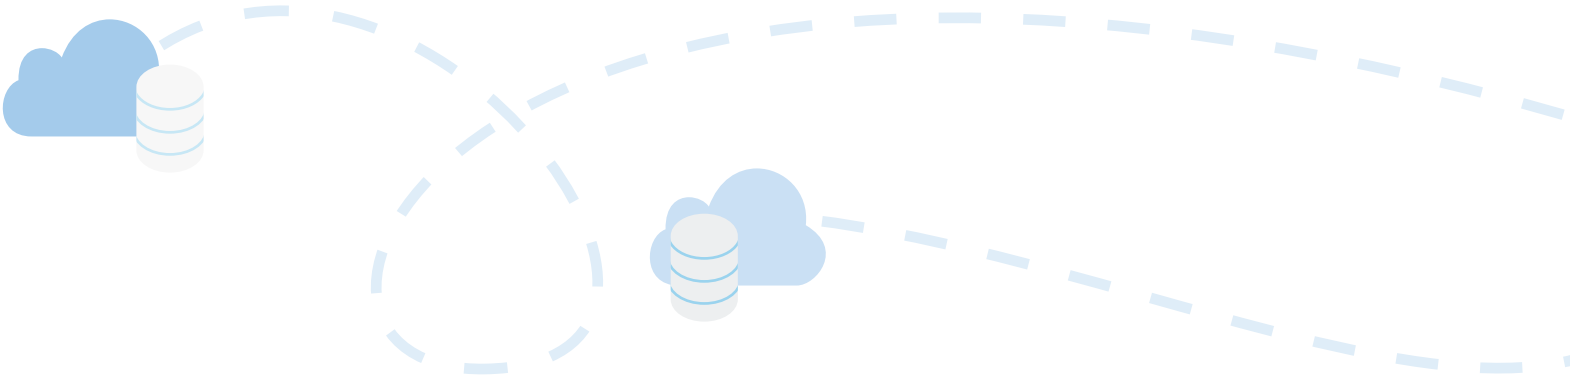

> **Replatforming databases is the most important commitment to a specific cloud service provider all CIOs will have to make in the next decade.**

With ADV, the immediate impact to an enterprise's bottom line can be as much as USD 10-20 million which means enterprises cannot afford to ignore this digital transformation accelerator. In customer surveys among Global 2000 enterprises as well as global and national system integrators, ADV is frequently noted for its potential to become the single most impactful component in the implementation of a cloud-first strategy for enterprise data management.

Future-Proofing the Enterprise Through Adaptive Data Virtualization

IT leaders have understood the true long-term potential and staying power of ADV: by deploying ADV in their enterprise data management stack, enterprises can replatform to new technology rapidly and—if needed—correct course by either reverting back to their previous technology stack or, in the more likely scenario, move on to new technology.

With its compelling economics, ADV technology is set to become an important cornerstone of enterprise cloud strategy by supporting rapid time to value, future-proofing the enterprise, and offering core virtualization tenets of consolidation, management, and simplification of IT.

The future of Datometry ADV is expected to be very similar to other virtualization technologies, such as storage and network virtualization. It is the company's vision for ADV to become the ubiquitous data fabric managing all communication between databases and applications.

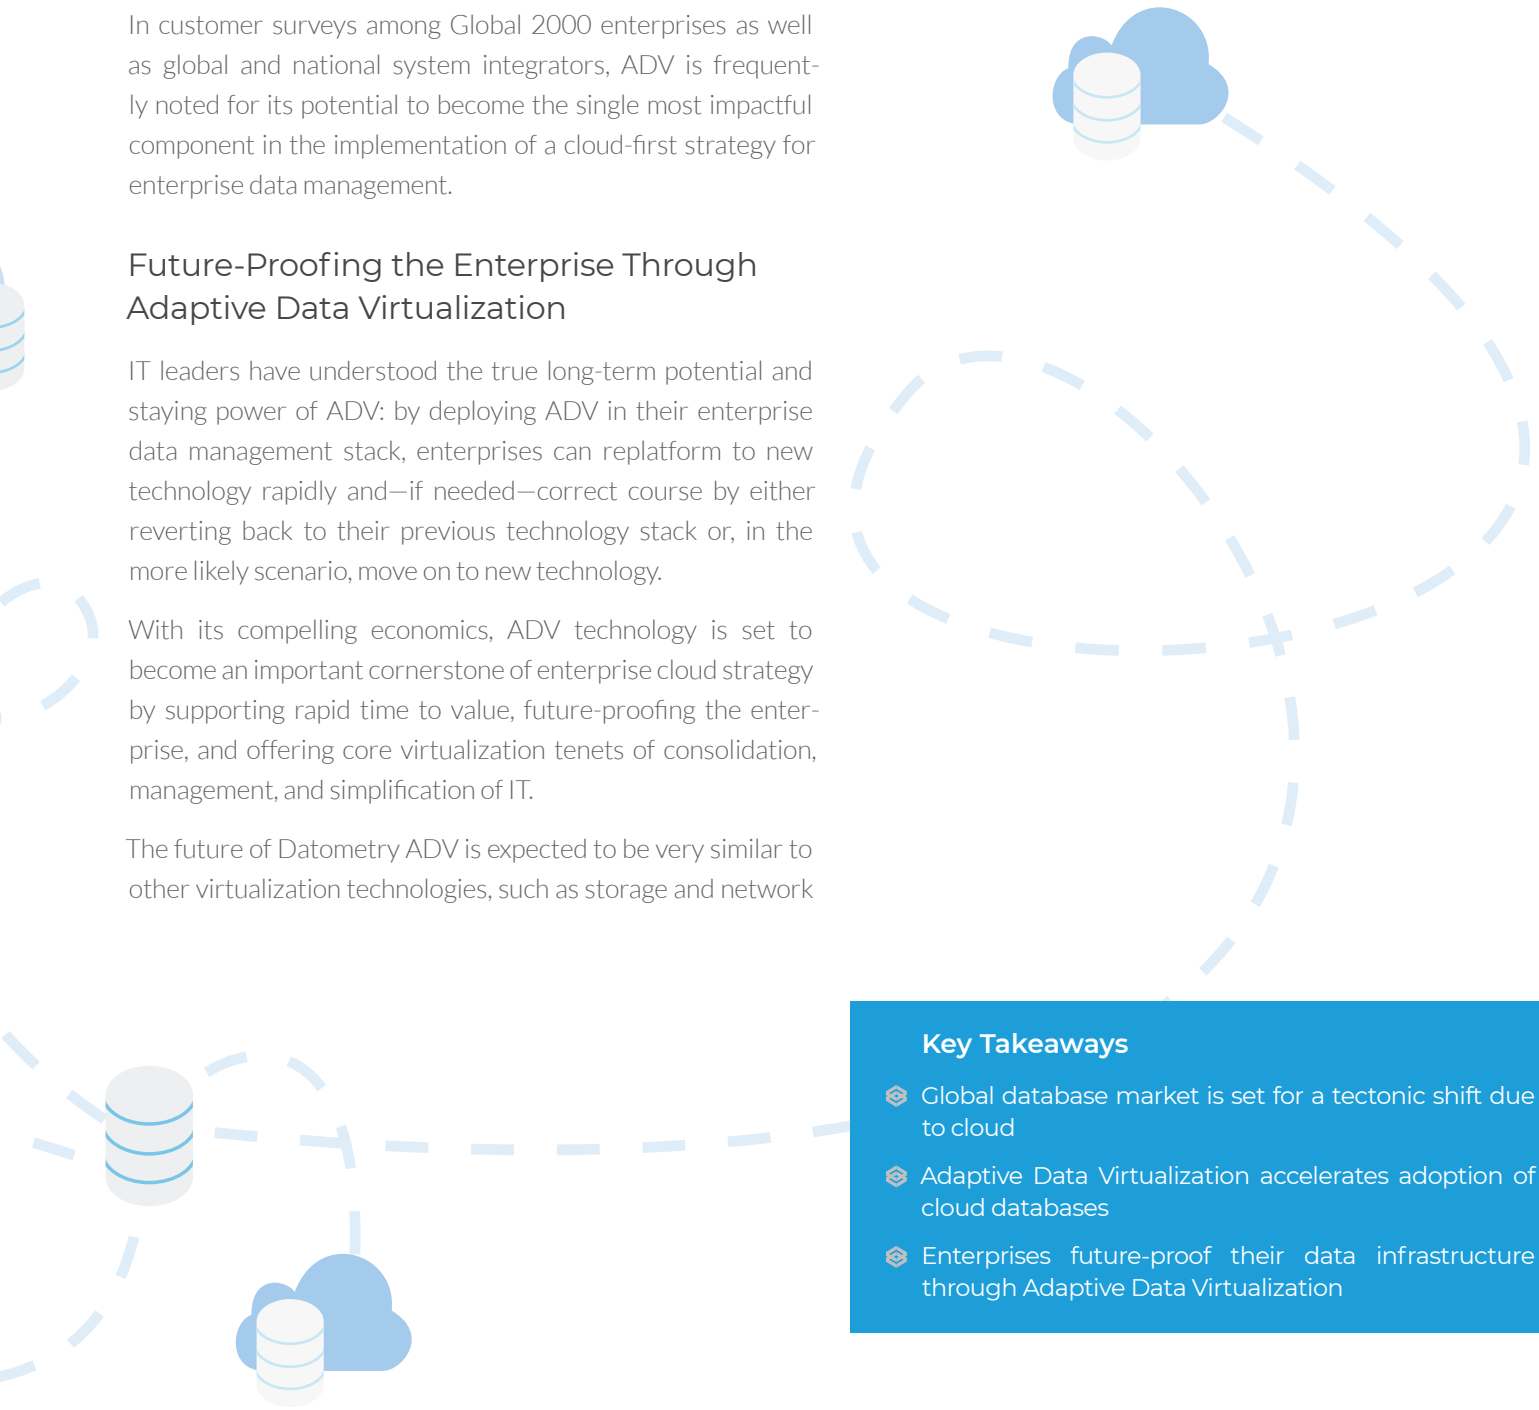

Key Takeaways

- Global database market is set for a tectonic shift due to cloud
- Adaptive Data Virtualization accelerates adoption of cloud databases
- Enterprises future-proof their data infrastructure through Adaptive Data Virtualization

ABOUT DATOMETRY

Datometry is revolutionizing how global enterprises can gain competitive advantage by up-leveling their data management to become cloud-native quickly and effectively. Its flagship product, Datometry® Hyper-Q™ accelerates the digital transformation of enterprise IT by eliminating the costs and risks of adopting new data management technologies and lets the enterprise realize significant savings, innovate faster, and increase the speed of business. Hyper-Q is powered by Datometry's proprietary technology developed by an expert team of database researchers and engineers. Datometry is headquartered in San Francisco, California.

Awards

20 Top Data Integration Companies 2017

30 Most Innovative Companies 2017

CEO Today USA Awards 2017

Top 20 Most Promising Virtualization Solution Providers 2017

Top 20 Most Promising Virtualization Solution Providers 2016

Microsoft Accelerator Cohort Member 2016

GIA Shengjing 2015 Global Innovation Award

References

Russom, P. 2015. **Four Drivers of Data Warehouse Modernization**, TDWI. Retrieved from https://tdwi.org/articles/2015/06/16/four-drivers-of-data-warehouse-modernization.aspx

Anderson, E. Warrilow, M. 2016. **Market Insight: Cloud Shift—The Transition of IT Spending from Traditional Systems to Cloud**, Gartner. Retrieved from https://www.gartner.com/doc/3321217/market-insight-cloud-shift-

Aslett, M. 2017. **Database market set for seismic shift driven by accelerated cloud and SaaS Adoption**, 451 Research. Retrieved from https://451research.com/report-short?entityId=93718

Columbus, L. 2017. **Roundup of Cloud Computing Forecasts**, Forbes.com. Retrieved from https://www.forbes.com/sites/louiscolumbus/2017/04/29/roundup-of-cloud-computing-forecasts-2017/#586a177b31e8