REDISCOVER MAGAZINE

THE CLOUD EDITION

9 Powerful Cloud Truths

From AI to security to multicloud, the cloud secrets no one talks about

Q&A with Kelsey Hightower

Google's cloud guru on Kubernetes as a design principle, open source as a service, and much more

Searching for Greatness in the Cloud

How the cloud helps us rediscover winning products, companies, and technologies

redislabs
HOME OF REDIS

ISSUE II

Microsoft Azure

redis

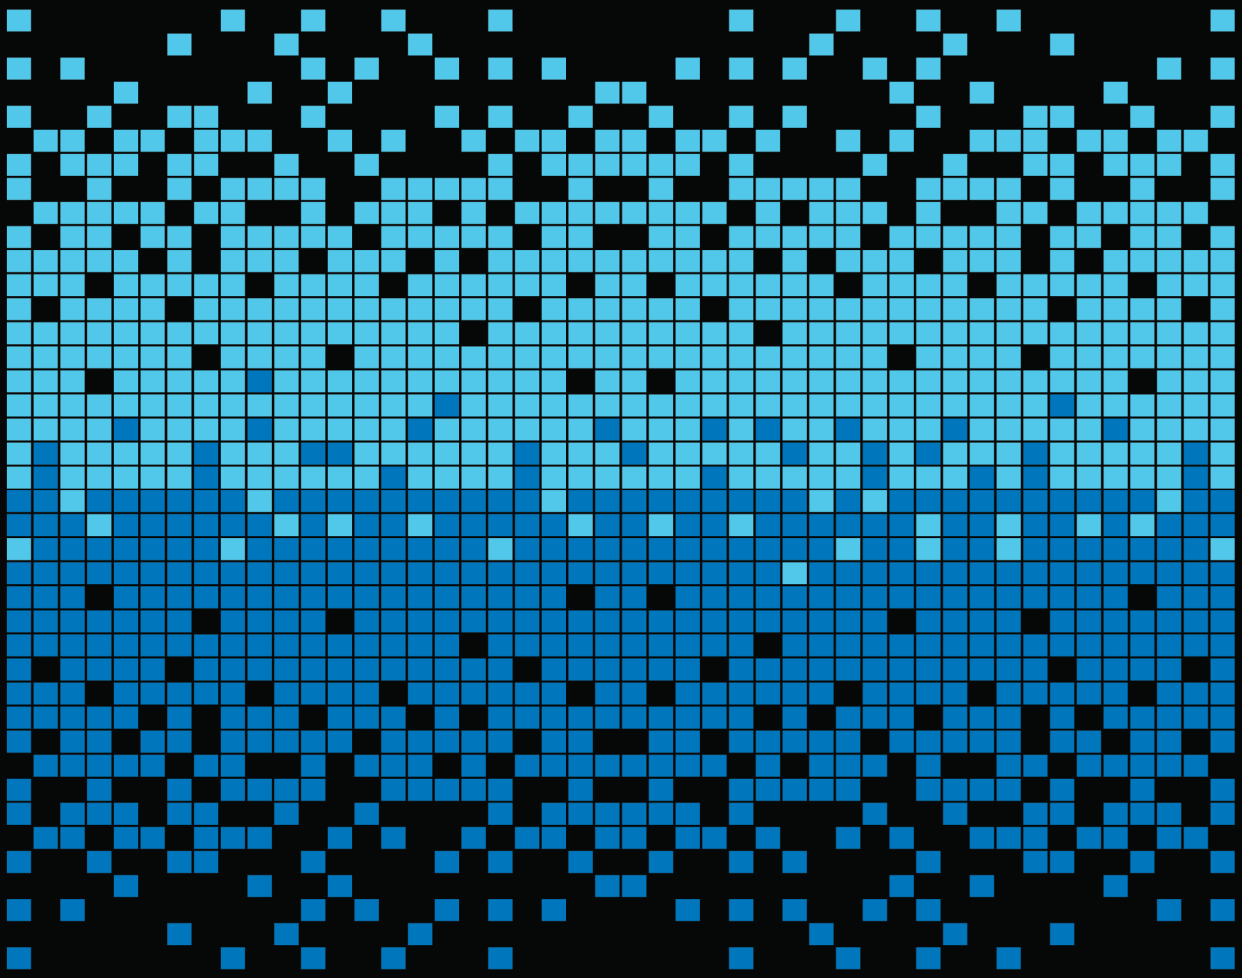

Open-source at global scale

Azure Cache for Redis

Developed in partnership with Redis Labs, Azure Cache for Redis delivers the same unmatched Redis Enterprise performance as a native Azure service with:

- Support for Redis Modules and Redis on Flash
- Industry-leading Azure security, reliability, and global scale
- Unified billing and support

Find out more and how to get started on the Azure Blog: **aka.ms/RedisEnterprise**

TABLE OF CONTENTS

Issue II

Publisher
Mike Anand

Editor-in-Chief
Fredric Paul

Managing Editor
Rusty Weston

Associate Editor
Haley Kim

Art Director
Estella Xian

Design Consultant
Talia Trackim

Contributing Writers
Mike Anand
Charles Babcock
Ofer Bengal
Sanjeev Mohan
Joe Mullich
David Needle
Yiftach Shoolman
Aaron Sun
Allen Terleto
Doug Tidwell

Redis Labs

CEO & Co-Founder
Ofer Bengal

CTO & Co-Founder
Yiftach Shoolman

CRO & GM North America
Jason Forget

Chief Business Development Officer
Taimur Rahid

Chief Financial Officer
Rafael Torres

Chief R&D Officer
Itai Raz

VP Marketing
Mike Anand

VP Human Resources
Adi Stern

EDITOR'S NOTE

Now Is the Time to Rediscover the Cloud

When we decided to focus the second issue of *Rediscover Magazine* on the cloud, it seemed like a no brainer. Way back in early 2020, the cloud was already on a tear, enjoying explosive growth as it upended traditional enterprise technology with the promise of unmatched business agility and flexibility, speeding online products and services to market, freeing development teams to create game-changing innovations and great digital customer experiences instead of worrying about plumbing.

And then COVID-19 hit. Suddenly, yet another "Year of the Cloud" turned into *THE* year of the cloud. As knowledge workers transitioned to working remotely and shoppers turned to delivery services and curbside pickup, the transformational impact of cloud computing shifted from the future to the present. Companies that had embraced the cloud were better able to cope with these dramatic developments, while everyone else raced to catch up.

It's not so much that the global pandemic changed the role of the cloud or the technology itself, but that these external forces accelerated and intensified the cloud's ascendance. Now it's more essential than ever that business and technology leaders possess a clear-eyed understanding of what the cloud is, how it works, and what it can and can't do.

This issue is packed with insights and intelligence designed to help you rediscover how the cloud can transform your organization, give you context on the concepts and technologies behind today's most important cloud technologies, and prepare for a cloud-driven future. In "9 Cloud Truths (That Nobody Talks About)," for example, you'll learn the real story about today's cloud challenges and opportunities (page 12). Google Cloud guru Kelsey Hightower, meanwhile, shares his insights on what really matters in the cloud (page 24), and Gartner Analyst Sanjeev Mohan offers some well-informed cloud predictions (page 8). And that's just scratching the surface.

Want to learn more about Redis Labs in the cloud? Check out **RedisLabs.com/cloud**. If you're looking for more about the power of rediscovery—including free access to the premiere issue of *Rediscover Magazine*—visit **https://redislabs.com/rediscover-magazine/**.

Thank you!

Fredric "The Freditor" Paul
Editor-in-Chief

A PUBLICATION FROM redislabs HOME OF REDIS

Rediscover Magazine

Questions, comments, or concerns for the Rediscover Magazine editorial team? Send an email to RediscoverMagazine@RedisLabs.com

Interested in Redis, Redis Enterprise, or Redis Labs? Visit RedisLabs.com or call (415) 930-9666

RETAIL CLOSES THE SALE ON DIGITAL TRANSFORMATION

What industry is leading the charge to the IT systems of the future, and who's falling behind? According to an exclusive Redis Labs survey of more than 500 AWS re:Invent attendees, retail is closing the sale when it comes to digital transformation, while education appears to be a slow learner.

Among many other things, we asked respondents about their strategies for the cloud, Kubernetes, microservices, and NoSQL databases. Then we assigned numeric values to their descriptions of their strategies and calculated a Digital Transformation Index (DTI). The index helps assess how far along an organization is on its digital transformation journey.

The full results will be released soon, but here are some spoilers:

- Across all industries, cloud strategies are the most advanced with a DTI of 3.64, while NoSQL came in last at 2.45. (The DTI ranges from 5 down to 0—or most to least transformational.)
- As noted, retail topped the list of industries with a DTI of 3.49, while education sat in the back of the class at 2.20.
- You might think the tech industry would be a leader, but it came in fifth with a DTI of 3.29.
- Finance scored even lower than tech (3.07), perhaps because most banks use highly structured data and need ACID transactions, so they're less likely to quickly shed their reliance on relational database.

That's only a small sample of the wisdom revealed in the survey. The report also shows how different industries make their database choices (another spoiler: cost is important), the popularity of putting a cache in front of a relational database, and which programming languages they use to build database applications. Keep an eye out for the complete survey results and analysis.*—Doug Tidwell*

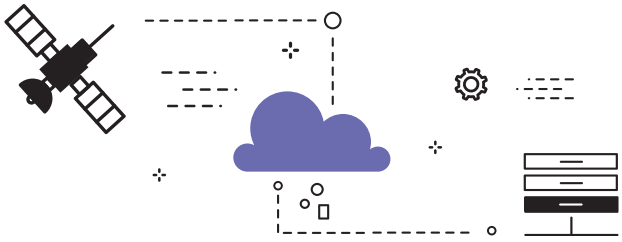

THE CLOUD ABOVE THE CLOUDS

It turns out there's a need for computational and communications power in outer space—particularly on the International Space Station. Sure, processing scientific workloads on a supercomputing platform orbiting 254 miles above the ground makes operational sense, but nothing comes easily or inexpensively in space. NASA wants to use cloud computing for data collection, but among many other things must contend with the presence of radiation, which is like Kryptonite for computer processing outside the atmosphere unless devices are specially hardened.

Hardships aside, the IEEE recently noted that cloud technology is crucial for space exploration, because astronauts and robots rely on readily available information to tip the scales between success and failure of a mission.

But exploration isn't the only reason that cloud providers are scrambling to launch space clouds. Microsoft's Azure Space cloud, for example, intends to "supply a multi-orbit, multi-band, multi-vendor, cloud-enabled capability" to help control a network of communications satellites. For its part, AWS Aerospace and Satellite Solutions foresee helping IoT communications companies blanket low-space orbit with thousands of sensor-equipped satellites to provide low-latency internet and high-resolution Earth observation. Meanwhile, IBM projects that edge computing may become prevalent in orbit because it saves time and energy to process data in space rather than transmit it back and forth to earth. Big Blue even sees potential for rolling out blockchain in space!*—Rusty Weston*

THE CLOUD VS. THE EDGE?

In the near corner, wearing white, we present our returning champion, the Cloud. According to Research and Markets, the cloud is expected to grow 17.5% a year to reach $832 billion by 2025.

In the far corner, clad in neon-red, meet our challenger, the Edge, much smaller and younger but growing twice as fast: 34% a year headed to a modest $15.7 billion in 2025, per Markets and Markets. So is this a fair fight for enterprise architecture supremacy?

Like the cloud, edge computing can sound a bit vague, but it's where you crunch data as close as possible to where it's created—when there's no time to backhaul it to the cloud or a data center. Edge applications typically perform discrete functions, often working with IoT data, such as monitoring the safety of an oil refinery in the Gulf of Mexico or the temperature of a Chicago highrise. Fueled by an enticing array of 5G, AI, and Internet of Things technologies, intriguing edge applications include various combinations of robotics, computer vision, AI, and deep learning in places that can't function without real-time data, such as hospital logistics, public and environmental safety, factories and manufacturing plants, city streets, and even autonomous vehicles.

While cloud vs. the edge sounds like a blockbuster Marvel franchise, in real life they're more like shoes and jackets vying to be your favorite apparel. Nearly everything processed at the edge is destined for the cloud, and almost everything else now in data centers is headed there too. So what seems like a sharp distinction today is likely to blur over time as new tools emerge to help enterprises seamlessly manage workloads and data on either platform.*—Rusty Weston*

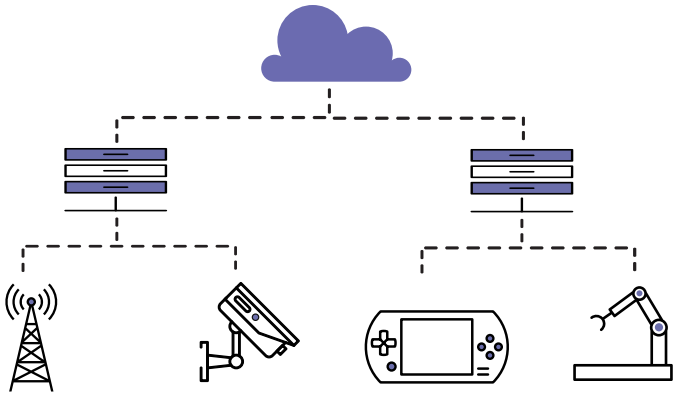

CAN JEDI PROTECT US IN THE REAL WORLD?

When stories broke in late 2020 about a series of devastating and long-undetected hacks of US federal agencies, all thoughts turned to the need for rock-solid cloud security. In the old days, a Congressman might have thumped a committee table with a wingtip shoe and bellowed that "If we can't stop foreign actors from plundering our secrets, we won't have anything left to protect."

In 2021, we're counting on JEDI to save us. And that's only one reason why enterprises should care about the heavily litigated, decade-long, $10 billion JEDI (Joint Enterprise Defense Infrastructure) contract with the Department of Defense with a range of cloud-computing services—including security and much more.

The winner, Microsoft's aspirationally named Azure Government Top Secret cloud, is intended to help federal agencies keep a lid on their intellectual property and secure data with stronger protections than provided by FedRAMP, for example. The big question: Will JEDI bring new hope to the rest of the galaxy?

After all, if JEDI can protect the technology behind Star Wars anti-ballistic missile systems, imagine what it could do for more mundane secrets like accounts receivable workloads and our personally identifiable information? To put it another way, how will JEDI's world-class cloud cybersecurity trickle down to the enterprise? Will it raise the bar for other cloud providers?

We'll see, but we looked it up, and Microsoft will not "trust the Force" to keep us safe. Instead, Microsoft Azure Government applies the same underlying technologies as regular, everyday Microsoft Azure but ladders up protections for top secret data. That may be better than what your own security teams can muster, but it's a long way from lightsabers, blasters, and Jedi mind tricks.*—Rusty Weston*

A NEW WAY TO SELL SOFTWARE IN THE CLOUD

BY OFER BENGAL
REDIS LABS | CO-FOUNDER AND CEO

Before the cloud era, when IT was installed and managed on-premises, it was enterprises' responsibility to deploy and maintain sophisticated software systems such as databases in their data centers.

That was a big, complex job. Servers can go down. The amount of data can grow, and you may need to cluster multiple servers. All that was handled by database administrators (DBAs) and eventually operations and DevOps teams, who worked for the enterprise.

When the cloud started, it initially provided infrastructure only (primarily compute and storage servers), but soon added data services, which are basically automated versions of the activities that DBAs, DevOps, and operations teams once performed.

In the early days, the cloud providers offered a few basic data services of their own and created marketplaces where independent software vendors (ISVs) could offer additional cloud services. Today, those marketplaces feature hundreds and hundreds of data services.

As the cloud industry matured, cloud providers developed advanced versions of the most popular data services in the market, including emerging data models, and offered them in the form of a Database-as-a-Service (DBaaS). On top of that, cloud providers used their power to offer discounts for committing to spend a certain amount with them—typically excluding ISV offerings in the marketplaces.

This creates a conundrum for enterprise buyers. Cloud provider service offerings can lower costs and help buyers streamline and unify procurement processes. But these built-in services may not be best of breed, and finding and using alternatives in the marketplace typically involves more friction than using the cloud providers' own options.

Buyers' decisions to go with the cloud provider's own offerings has a dramatic negative effect on third-party ISVs offering their services in the cloud provider's marketplace. This can be a live-or-die question.

So, what can these third-party ISVs do to compete? The obvious solution is to offer a better service than what's available from cloud providers. Now, not every ISV is able to do that, but some can. There *are* better solutions out there.

But even companies with a significantly better offering may not always be able to completely overcome the structural barriers erected by the cloud providers. That's one reason many ISVs choose to partner with the cloud providers—demonstrating to them that offering strong third-party software can increase their own market share. That's the approach that Redis Labs has taken with Microsoft Azure and Google Cloud.

The challenge for ISVs is to create a strong enough differentiation between their product and what the cloud providers offer to create an incentive for them to partner.

For database vendors, that means adding unique value-add capabilities like Active-Active Geo Distribution, on top of core capabilities like unlimited scaling and high-availability. Active-Active—the ability to deploy a database in multiple geographic locations, read and write in each one independently, and keep them all synced—is in great demand by enterprises, and is available with Redis Enterprise.

With their virtually unlimited resources, cloud providers can theoretically develop almost any capability, but it may require major product architectural changes, and it may take a long time. In the meantime, ISVs like Redis Labs can create the *next* competitive advantage. This reality presents cloud providers with an incentive to partner.

Finally, it's important to highlight the importance of hybrid and mulitcloud deployments. Many enterprises want to deploy their IT systems beyond a single cloud. So being able to support on-premises, hybrid, and multicloud environments can be another way for third-party software providers to bring value to enterprise customers beyond what the cloud providers offer. By building better software and software powered by unique, in-demand technology, ISVs can attract enterprise buyers and show the cloud providers that partnerships can be mutually beneficial.

THE RISE OF THE REAL-TIME DATA PLATFORM

BY YIFTACH SHOOLMAN
REDIS LABS | CO-FOUNDER AND CTO

More and more database and big data companies have started to position themselves as data-platform providers. Although every company uses a slightly different term—MongoDB's data platform, Cloudera's enterprise data platform, Snowflake's data cloud, and Databricks' lakehouse platform—they all boil down to the same thing.

A data platform serves as a one-stop-shop for all your data needs, and lets you ingest and store data, and run operational transactions as well as analytics queries on your data from a single platform—often in multicloud/hybrid-cloud deployments. Many data platforms started as a single-purpose database supporting operational or analytics use cases with a single data model, such as relational, document, key-value, graph, time series, and so on. Eventually, they evolved into multi-purpose databases supporting operational and analytics workloads using multiple data models. Of the top-ranking databases on db-engines.com, the top eight have all classified themselves as multi-model databases.

A multi-model database simplifies application development by providing a single set of APIs to access multiple data models. It also simplifies operational complexity, as DevOps teams can use the same deployment template with minor configuration tweaks to set the scaling factors, replication sets, data persistence, and consistency schemes. Finally, a multi-model database can lower costs, because only one database is needed for most workloads.

So how is a data platform different from a multi-model database? Essentially, a data platform is the "programmability option." Every data platform provides users with tools to program data flows and triggers across cluster nodes and data models, using an embedded or external database engine.

Redis has always been considered more than a simple key-value datastore. Some people even used to call it the data-structure store, thanks to Redis core's 10 native data structures, including Strings, Hashes, Sets, Sorted Sets, Lists, and others.

Over time, the need to support modern data models has significantly grown. With the introduction of Redis 4.0 in 2017, we opened Redis core to accept new functionality in the form of Redis modules. Since then, we have introduced five new data models that together map nicely to the new way of modeling modern applications.

Today, applications that need real-time performance can store objects in JSON with RedisJSON, index and query data with RediSearch, map connections in your data with RedisGraph, provide advanced monitoring with RedisTimeSeries, and use RedisBloom to solve probabilistic problems.

In 2020, with the introduction of RedisGears (a low-latency distributed programmability engine) and RedisAI (a fast inferencing engine), Redis has become a real-time data platform. Today's Redis lets you control any data flow, set up relevant data triggers, run aggregations and complex mathematical functions, as well as perform analytics queries across multiple data models in a fully distributed manner, powered by AI, while keeping the sub-millisecond processing time of Redis.

Our customers use Redis as a real-time data platform in a variety of ways:

1. A large financial service company runs real-time, AI-based predictions on time-series data, using RedisGears and RedisTimeSeries.
2. A market intelligence company uses Redis for real-time analytics, powering RediSearch and RedisGears for dynamic indexing and querying of billions of fast-changing data points.
3. A government agency runs its real-time facial-recognition application based on AI search algorithms using RediSearch, RedisGears, and RedisAI.

Of course, these use cases merely scratch the surface—Redis users around the world are busy finding new ways to leverage Redis' power as a data platform.

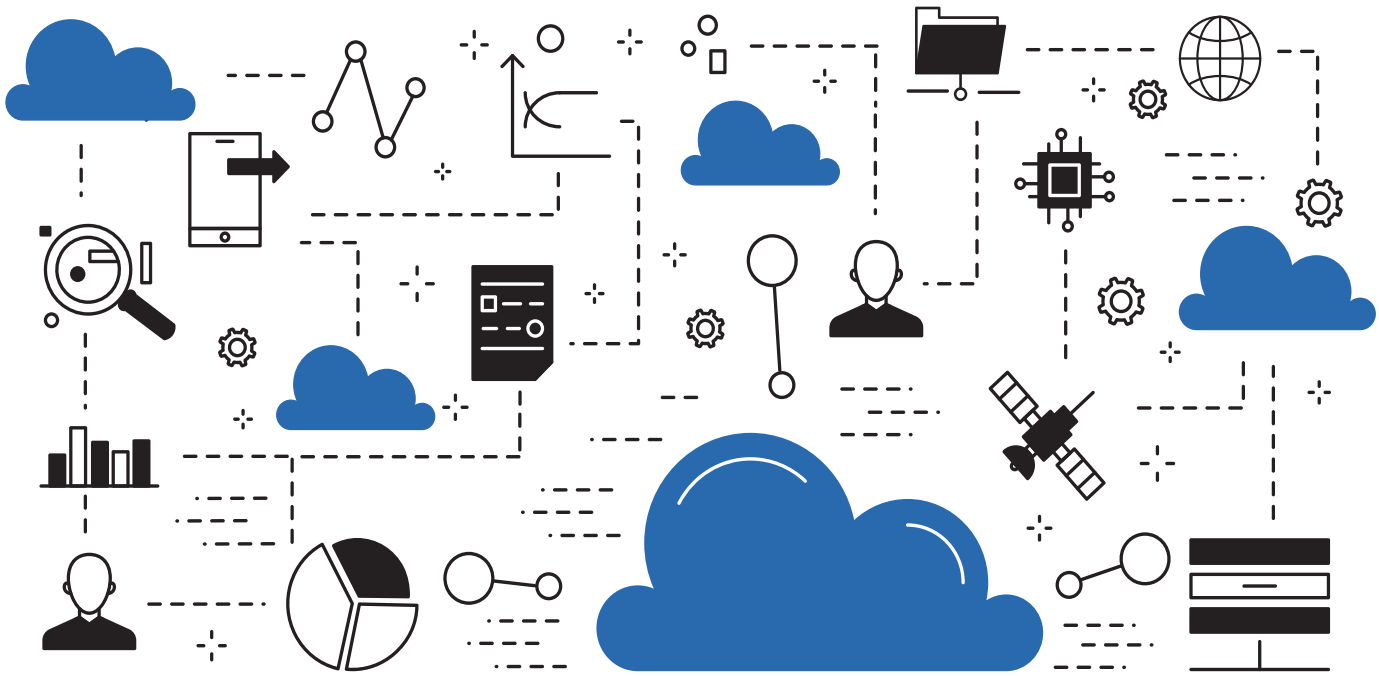

5 CLOUD TRENDS TO WATCH

2020 was a banner year for the cloud, and the momentum is accelerating

BY SANJEEV MOHAN
GARTNER | VICE PRESIDENT, ANALYST

2020 was an incredibly difficult year, but there's no denying that it inadvertently proved the cloud model. Companies already in the cloud benefited tremendously, while those that were not pushed ahead as fast as they could. Movement to the cloud is only going to accelerate in 2021, driven by five key trends:

1. Customer experience comes into focus

Snowflake's blockbuster IPO showed the world that it's not all about technology—it's about the customer experience you provide to your users. Techies like me love to be able to configure a product to our heart's delight, but that's not always the right thing for organizations.

As a result, we are starting to see more emphasis on platform. Customers don't necessarily want to be in the business of integrating a lot of best-fit pieces—at the end of the day, they want ease of use, ease of maintenance, and a user experience that makes them eager to do more with the technology, rather than spend time gluing it all together.

2. On-premises isn't going anywhere

It's clear that growth will come from the cloud—not on-premises—but not everything is going to move to the cloud just yet. There are a number of cases where it makes sense to stay on-premises, like IoT, because it's cost-prohibitive to move such massive amounts of data. But hybrid models have become a really important component strategy for everybody, and with that comes a brand new space, which

is data orchestration and data ops.

This doesn't just mean hybrid cloud, but also multicloud. How do I build my future-state architecture if the data moves between different places? You need end-to-end orchestration of the data pipeline.

3. More emphasis on data governance

The California Privacy Rights Act (CPRA) was passed in November 2020. This more-stringent data privacy regulation aligns more closely to the European Union's GDPR. It's becoming the blueprint for the other 49 states to be more sensitive about protecting people's private data. And it's not just the United States—other countries are also enacting laws that prohibit data from leaving their borders, so I'm very interested in how quickly the cloud scales globally to handle data privacy regulations.

This trend is potentially the most important, because cloud providers have not put as much effort into data governance as they probably should. Fortunately, we have some really good, credible, and independent companies that have started to corral this space of data governance. Right now, though, there are too many products and too many players, so there's a lot of confusion.

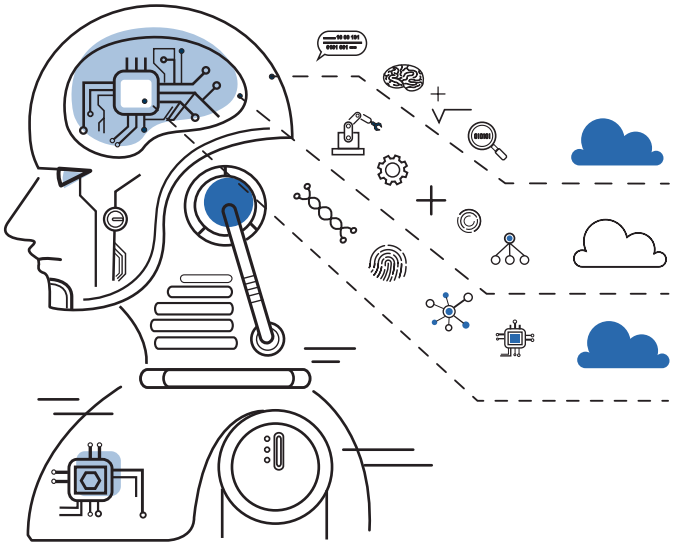

4. AI and machine learning continue to accelerate, especially in the cloud

Machine learning and AI have become far more accepted, and I see new developments with automation. But with increased focus on data governance, there is now a lot of emphasis on things like explainable AI. How can I build my machine-learning models but still ensure that I am doing responsible, ethical AI?

We also can't ignore the fact that AI and machine learning benefit the most from the cloud. If I have a mission critical, real-time, low-latency application requirement, I may not want to move to the cloud right away because I have dedicated hardware. It may be over-provisioned and I may be paying too much for it, but it's mission critical. But when I'm doing machine learning, I'm still exploring for new insights. I don't know how long it will take. I don't know if I will need 100 CPUs, 20 GPUs, or maybe a field programmable gate array. I will likely need a lot of resources, but for a very short amount of time. I definitely need an environment that can elastically scale up or down. So AI and machine learning products are actually best suited to leverage the elasticity and scalability that we get from the cloud.

5. Open source software keeps rising

Open source software is now considered to be secure and enterprise-ready, and we are starting to see a lot of growth there. Many options like PostgreSQL and its extensions like geospatial and time series are now state of the art. They give you familiar open standards that you can change if you want to move to something else. Many organizations I've spoken to now embrace policies to employ open source first, commercial next—and this shift of mindset is relatively new.

Finally I've also noticed that all the vendors are increasingly focused on how data is delivered to end users. For example, being able to view data through a knowledge graph when you are using machine learning and other technologies is relatively new. We are also beginning to see the goal of reducing copies of data being realized. This also supports the data minimization requirements of the EU GDPR. Instead of replicating data and creating more data silos, we are finding ways to control who sees what data through dynamic access control. Sharing data in a secure and governed manner without making copies is, to me, where Data-as-a-Service comes to life. Newer data sharing platforms are helping create secure and governed data exchanges and marketplaces.

Sanjeev Mohan is a VP Analyst researching data management and analytics strategies within the Gartner for Technical Professionals group.

BOTTOM UP AND TOP DOWN SALES: GROWTH, SCALE, AND PREDICTABILITY

BY MIKE ANAND
REDIS LABS | VICE PRESIDENT, MARKETING

Selling enterprise software comes with a built-in structural conflict. Early on, software companies often leverage a product-oriented approach to attract and satisfy developers who help create the adoption and validate the product-market fit. But achieving sustainable growth, scale, and predictability in the enterprise software market also requires an enterprise-wide approach that focuses on business value.

The conundrum for marketing leaders is how to build and execute the low-touch, bottom-up strategies that attract technically oriented developers, while also incorporating the top-down, sales-driven motions that speak to the needs of strategic business decision makers.

Many start-ups view this top-down motion as an expensive, nice-to-have addition that can wait until everything else is in place. But that bottom-up funnel simply can't take you all the way. The issue, of course, is that combining those two approaches is as much art as science.

Product-led bottom-up growth is now seen as the most efficient way to seed and acquire new users. The rise of cloud computing accelerated this groundswell by removing the friction involved in building, buying, and deploying software, putting developers firmly in control of the product adoption journey.

But developers' distaste for traditional marketing is well documented, so the path to victory starts with a few key steps in a different direction:

First, you need to hire the right leader and the right team. Credibly reaching technical buyers needs a deep understanding of how developers think, how the technology underpinning your product works, and how it meets their needs.

Next, you have to make things easy for your customers by streamlining the sign-up and onboarding experiences to require as few clicks as possible.

Then you need to radically simplify your messaging. Instead of lofty promises and clever slogans, developers crave things like great documentation, hands-on training, technical case studies, and practical sample apps.

Finally, test, test, test. What's driving signups? Who is signing up? What is their customer journey? Which nurtures are working, and why? What are the conversion rates, and who is converting—and who isn't? Once you can answer these questions, you're ready to continuously iterate on your marketing motions to maximize results.

But to fully crack enterprise-wide adoption beyond the initial product-driven flywheel, companies need to focus on use cases and value vs. features. That means product teams must learn to build for enterprise buyer needs. It means increasing investments in product marketing to create outside-in messaging and solutions for key vertical markets. And everyone needs to obsess over usage data to understand customer journeys and conversion drivers.

Next, account-based marketing (ABM) should be the main focus in order to drive pipeline, but it must be paired with one-to-many lead-generation activities to bring in new people. Smart partnerships with the cloud vendors and other channel leaders are essential to improve your win percentage.

> ***"But to fully crack enterprise-wide adoption beyond the initial product-driven flywheel, companies need to focus on use cases and value vs. features."***

Finally, if bottom-up marketing is about creating a product-driven flywheel, top-down is about leveraging the tailwinds for pipeline generation to spear the biggest fish. "Solve not sell," driven by use cases, agile go-to-market tactics leveraging ABM, and hybrid pricing to create usage-based growth—as well as flexibility to discount at the top—are all essential for success.

LEARNING FROM THE PANDEMIC TO SUCCEED BEYOND IT

BY ALLEN TERLETO
REDIS LABS | FIELD CTO

The response to COVID-19 was a systemic shock to the global economy that will influence enterprise operating models and consumer behavior beyond the pandemic. Based on conversations with business and technology leaders, I've found organizations that outperformed market expectations often emphasized the importance of *business agility*. Early in the pandemic, these enterprises pivoted to *digital* channels by quickly adjusting their operating models so their consumers could safely and conveniently access their products and services. Their business agility and digital readiness proved to be the difference between perseverance and catastrophe as, in many cases, these channels became the only direct option for consumer interaction.

For more than a decade, operating models have been slowly shifting towards digital processes, technology, and culture. This *digital transformation* movement promised increased business agility through application modernization, leaner processes, and agile methodologies. However, while there were plenty of success stories, things didn't always go as planned. It turned out that transforming existing operations was expensive, complex, and often political, so many enterprises delegated digital investment towards greenfield projects and flashy new mobile applications instead of replatforming their legacy technology stack. These workarounds were common until the growing movement for cloud migrations exposed their brittle architectures.

Cloud migrations promised increased business agility through on-demand scalability of infrastructure, platforms, and cloud-native services. They renewed the call for application modernization, since the "lift-and-shift" model rarely worked where it mattered most—the data tier. Though legacy databases are expensive to support, and least capable of meeting digital-performance requirements, they are also the most challenging to replatform due to decades of built-in complexity reliant on in-house expertise. To overcome these on-premises anchors, enterprises began to widely-adopt a new architectural style based on lightweight and loosely coupled business-oriented *microservices*.

> *"I've found organizations that outperformed market expectations often emphasized the importance of business agility."*

Microservices enable incremental migration of subdomains from monolithic technology stacks. The movement towards this architecture promised increased business agility by allowing microservices to operate on their own release cycles, embrace end-to-end product-ownership, and adopt a *DevOps* culture. Enterprises reduced time-to-market for new service development from months to days. Microservices also reduced barriers to data tier cloud migrations since they primarily rely on cloud-native NoSQL databases, according to a 2021 IDC InfoBrief.

Throughout this year's executive briefings, prior investments into these movements were repeatedly referenced as a key differentiator for perseverance through the pandemic.

Their systemic impact on technology, process, and culture provided leaders with the requisite business agility to pivot their operating models in time to make a difference. It was through digital readiness, cloud-native technology, application modernization, and building a DevOps culture that these enterprises were able to reinvent the customer journey composed primarily of digital channels.

Final thoughts

The pandemic provided a preview of the digital-first economy and shortened expectations of when we will reach it. As we turn the corner on COVID-19, it's clear that digital's impact on culture and business will only continue its acceleration. This pandemic showed us that organizations can no longer afford to deprioritize modernizing the data tier as part of their digital transformation and cloud migration initiatives.

In the digital-first economy, data will be the lifeblood of the organization, which makes databases the heart of enterprise architectures. By reinventing the customer journey across digital channels using cloud-native technology, enterprises can increase their business agility and be better prepared for whatever lies ahead.

9 CLOUD TRUTHS

That No One Talks About

Cloud nine or cloud 911? After a pandemic-driven year of sometimes frantic cloud adoption, it's time to rediscover the essential cloud truths lost in all the bustle.

BY JOE MULLICH

2021 marks the cloud's 15th year of existence, making it a promising teenager on the cusp of maturity. But ready or not, the pandemic has supercharged the demand for cloud-based services: IDC reports that 80% of enterprises now intend to double the speed of their shift to a cloud-centric infrastructure. And the scramble to the cloud isn't going away when COVID-19 does: Gartner predicts cloud spending will jump another 18% in 2021.

In the midst of this gold rush, however, some fundamental truths about the cloud can easily get overlooked or papered over. That doesn't mean the cloud isn't delivering huge value, or that companies should hit pause on their cloud initiatives. But it can be difficult to focus on critical nuances when everyone is scrambling to transition to the cloud faster than ever before. So, to make it easier to understand just what you're getting into when you move critical applications and workloads to the cloud, here are nine important cloud truths that don't always get the attention they deserve:

1 The cloud may not be as cheap as you think—and that's OK

The idea that running your business in the cloud is always cheaper than using on-premises data centers has always been more myth than reality. "There is a lot of BS, to use a technical term, about saving money by using the cloud," says consultant Joshua Greenbaum, Principal for Enterprise Application Consulting (EAC).

But many companies haven't gotten the message. In fact, according to a 2020 Flexera study, companies bust their cloud budgets by an average of 23%—a figure expected to rocket up to 47% in 2021. One problem is that companies waste some 30% of their cloud spend due to overprovisioning resources, and a lack of cloud infrastructure performance tuning can also lead to cost overruns.

These issues can be mitigated with the help of experienced cloud DevOps talent or better management tools, but not everyone is willing to make those investments. "People are biased toward saving money rather than spending money," says Corey Quinn, Chief Cloud Economist with the Duckbill Group. "People misunderstand the nature of their business, and so they focus on shaving off pennies, rather than moving their business forward."

Focusing on cost at the expense of strategic advantages falls under the category of penny wise, dollar foolish. Enterprises spend just a tiny fraction of their revenues on application hosting, so viewing cloud investments only through the lens of cost savings is limiting and not always strategic.

"A lot of the public cloud software forces companies to standardize on best practices, and that exercise can save a lot of time, complexity, and money," says Greenbaum. "But it's ridiculous to move to the cloud just to try to get cheap software. That won't move your business forward. You want more from the cloud than cheap software. A lot more."

2 The real cloud payoff is measured in speed, flexibility, and agility

"Treating the cloud just as an auxiliary to your existing data center can be a useful transitional step," Quinn says. "But too many companies get stuck there, declare victory, and miss out on the real benefits of the cloud. Theoretically, the upper bounds of cloud savings is 100%, but you can make multiples of that by hitting the right market at the right time with the right product."

The real payoff from cloud initiatives comes from bolder innovation, faster time-to-market, and better business agility. For instance, when you need to scale up workload processing power in the cloud, additional compute can be added in just a few clicks—or even automatically—making it vastly easier to experiment with new technologies and processes.

The cloud allows new applications to bypass the time-consuming provisioning and procurement processes that make on-premises technology so irksome. For example, DBS Bank in Singapore was able to leverage its cloud foundation to quickly spin up new, self-service applications, including one that allowed house-bound customers to submit their documents digitally. As a result, the bank landed tens of thousands of new online equity trading accounts while rivals floundered.

Growth, not savings, is the truly important part of the ROI equation. Disruptive technologies such as advanced analytics and artificial intelligence, for example, are underpinned by the cloud. In addition to speed, scale, and flexibility, the cloud's ability to manage and process huge amounts of data in real time and host apps that can unlock insights about customers and opportunities will be worth far more than reducing the number of servers on premises.

3 You won't optimize your cloud journey unless you optimize your data layer

The temptation is understandable: Just move your legacy databases with their existing data layers to the cloud and be done with it. What companies often don't realize is that seeming short-term gain can lead to long-term pain.

For example, legacy monolithic applications—and their data—running in your data center may need to be broken up into microservices in order to take full advantage of moving to the cloud. Not taking the data layer into account can create data silos and unnecessary complexity. Legacy applications often rely on one relatively simple data model. But the cloud makes it easy to leverage a multiplicity of data models, and you won't get the full benefits if you're stuck with an obsolete, suboptimal data layer that doesn't address a variety of data models, such as graph, time series, JSON, and more. The average company requires three different databases to manage the data across their various hybrid and public clouds, relying on APIs to share data among them.

Most importantly, traditional databases are often too slow and inflexible to keep up with web-scale demands in the cloud. At a time when making real-time prediction decisions has become a delineator of success, storing data at rest for later analysis is a non-starter. Your cloud applications need a fast, highly scalable, highly available database platform to deliver real-time application performance no matter the amount of data or the number of users.

A modern, cloud-native data layer must be able to scale with your application, and your organization, without causing your costs to soar as well. High availability of the data layer is also non-negotiable. So is responsiveness. If your cloud app responds to customers in the blink of an eye, you could be too slow! To satisfy users today, an app has to respond in about 100 milliseconds, including sending the data back and forth, which doesn't leave much time for the database to do its thing. Finally, in a global economy, the data layer must support applications and customers no matter where they're located, with no data loss, degradation of availability, or performance issues.

4 Most apps still aren't in the cloud, and some will never be

A few years ago, people were quick to tick off lots of reasons certain applications couldn't run in the cloud—security and compliance, serious data durability, and strict latency requirements were often considered deal breakers. Now, there are fewer "stories of what doesn't belong in the cloud," Quinn says.

Yet, only about 20% of applications now run in public and private clouds or edge environments, according to Accenture. Over the next five years, that number will rise to 80%. Given the cloud's ability to scale dynamically, workloads that are highly variable or unpredictable—or where self-provisioning is important—are obvious cloud use cases. But the migration should never be by rote. Deciding if an app belongs in the cloud requires a careful analysis of the application and factors such as latency requirements, its value to the business, and more.

There are still plenty of examples of apps that shouldn't be in the cloud. High performance computing may be the next "it-can't-run-in-the-cloud" story to fade away, but for now HPC applications can be a stumbling block, at least economically. "Any application with massive data egress charges doesn't belong in the cloud," Quinn says. "Even moving between two availability zones in the same region comes at a cost of $.02 per gigabyte, which adds up to a lot of money when transferring multi-petabyte workloads."

Latency can be a critical determinant of app placement. Data processing at the edge, often based on IoT inputs, is not yet commonplace, but interest is rising in industries as varied as manufacturing, healthcare, financial services, and communications. In those specialized circumstances, once the data is analyzed it is shuttled to the cloud for further analysis in data lakes or cloud-based data warehouses.

5 AI in the cloud is hot—but requires the right architecture

Companies are salivating about artificial intelligence in the cloud. From fraud detection to medical diagnosis to facial recognition, the speed, flexibility, and agility of the cloud combined with AI promises to transform all sorts of applications and industries, providing compelling insights, automating full lines of business, generating new revenue streams, and making people's lives easier.

Cloud-based AI databases must simultaneously ingest, explore, analyze, and visualize fast-moving, complex data within milliseconds. The goal is to lower costs, generate new revenue, and integrate AI models so that businesses can make more efficient, data-driven decisions, act with greater agility, and decrease risk while exploring new opportunities.

But as companies begin to shift from AI and machine-learning (ML) pilot programs to actual production systems, many don't see a potential bottleneck looming on the horizon: the need for an enterprise feature store and inferencing engine. In this context, "features" are highly curated sets of data that inform predictive models. For example, a fraud model might include a feature based on unusual credit card purchases. A feature store, meanwhile, catalogs and serves up all available features—ready for use by any machine-learning model. Centralizing these features in the cloud speeds development, improves model accuracy, and eases version tracking.

So, when you think about AI in the cloud, don't forget that creating "intelligence" and value takes more than just training models. You also need to consider inference and features. And because many AI use cases are real-time, you need to keep your features and inferencing close to your database in order to minimize application latency.

Moving to the cloud is changing IT jobs—and unleashing creativity

6

As workloads shift to the cloud, changes to the IT organizational chart are following suit. Although virtualization already ushered out many hardware-focused jobs, the cloud is making even more dramatic changes to rack-and-stack roles.

According to the 2020 IDG Cloud Computing Survey, 67% of organizations have added new cloud roles and functions. Data center specialists are being replaced with cloud administrators or cloud architects. (If companies can find them. It's not easy to hire "unicorns" with skills in traditional IT, public and private clouds, security, and governance.)

Even if people's job titles remain static, the daily work they do is changing significantly. Offloading management of infrastructure to a cloud provider reduces the amount of control on-site personnel have over applications. Employees must develop new skills in contracts, vendor management, and helping end-users deal with new processes.

These workforce disruptions won't come without pain. Some jobs will be lost as companies close or de-emphasize their data centers. But there's a big upside, too. As noted earlier, the real opportunity for IT departments comes with offloading mundane chores so key personnel can focus on more strategic tasks. "There is a crying need for all enterprises to cut the umbilical cord of excessive operational maintenance cost," Greenbaum declares. For example, a Harvard Business Review white paper recounts how after one CIO moved his enterprise ERP system from an on-premises data center to the cloud, he was able to exchange five application administrators for five IT relationship managers who worked with the business units to identify their business requirements.

Another CIO says cloud capabilities changed the mindset of his team. The cloud allowed them to take snapshots of the environment and spin them up for tests or development projects, and then instantly revert back to a version in the cloud if needed. This gave them the "psychological safety" they needed to try new ideas.

7 Everyone says they have a "multicloud strategy," but it's mostly just talk

Clouds are like potato chips: Hardly anyone can stop at just one. Some 93% of enterprises say they have a multicloud strategy, according to the 2020 State of the Cloud Report by Flexera. On average, the study reports, enterprises use 2.2 public clouds and 2.2 private clouds.

Using multiple cloud providers can help companies avoid vendor lock-in, leverage best-of-breed features and capabilities, and boost leverage on pricing. Using a second cloud for data backup can be a sensible disaster-recovery strategy for worst-case scenarios. "However, few companies really approach multicloud from a strategic standpoint," Greenbaum says. "They just end up there."

This can happen for many reasons. A software vendor may nudge them to a specific cloud. Or a company might make a big acquisition and inherit three different ERP systems scattered in different clouds. Perhaps a rogue subsidiary cuts its own cloud deal. Or maybe it's a shadow-IT situation where an ambitious development team spins up its own instances without clearing it with corporate.

So rather than implement an actual multicloud strategy, most companies tend to cluster their infrastructure on one cloud, with bits and pieces running on other clouds for a variety of reasons. But while it's most common to have different applications running on different clouds, companies are increasingly choosing to distribute a single app across multiple clouds to create globally distributed service that delivers low latency to end users no matter where they are located.

The nirvana of multicloud—when the same app running on multiple clouds can have its data seamlessly traverse clouds, or failing over one from cloud to another—is extremely complex and still fairly rare. To make it work, companies need to coordinate the connections among the various cloud services and platforms and any on-premises resources. Notably, moving stored data from one cloud to another for processing can wreak havoc on application latency rates and network bandwidth.

There are other multicloud issues as well. "You are at the whim of both the software vendors' upgrade, update, and API-management strategy, as well as the maintenance and support strategy of the cloud vendors themselves," Greenbaum warns. "There is a crying need for a single pane of glass to manage multiple cloud environments, and that doesn't exist right now. You have to put in a lot of sweat equity to manage all those moving parts."

If you have a cloud security issue, it's probably your own fault

8

Admit it: Putting data in the cloud is scary. A recent Bitglass survey of IT professionals found that almost three quarters (73%) were extremely or very concerned about security in the public cloud. Another 20% were moderately concerned. That sounds ominous—until you realize that three quarters of those same IT professionals also acknowledged that cloud-based apps and software were just as secure as or even more secure than on-premises systems.

The real point? Putting data *anywhere* is scary for IT professionals.

The truth is the cloud is more secure than on-premises data centers. Unlike many enterprises, cloud providers have no choice but to spend heavily on the latest versions of security apps and services, and are usually two-to-three generations ahead of most enterprises. For example, few on-premises systems encrypt their data in transit and at rest as all the major cloud providers now do.

Enterprises are slowly but surely catching on to this complex reality. (In the Bitglass survey, 28% said clouds were more secure, while 27% said they were less secure.) As the CIO of a large consumer goods company put it in a *Harvard Business Review* white paper, "This is not our grandfather's cloud where it was completely vulnerable from a security standpoint."

When cloud security issues do occur they are almost always due to customer error—misconfigured implementations, poor user management, and other user mistakes that the cloud provider does not control. The real question, then, is not whether the cloud is secure, but whether your company is using the cloud in a secure manner by implementing proper procedures and controls. Because the cloud is only as safe as you make it.

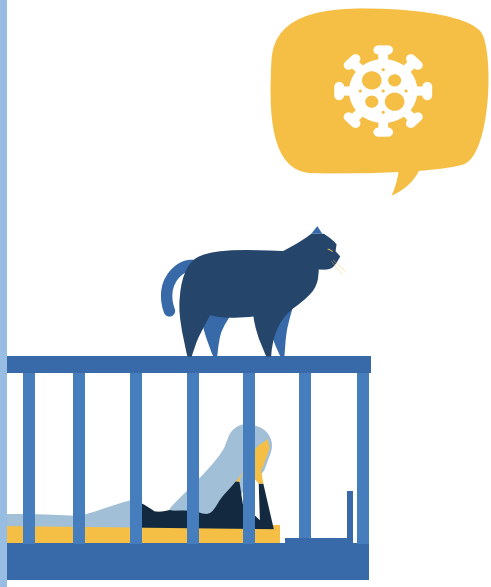

9 The COVID cloud rush will have big but unexpected repercussions

As noted at the beginning of this story, the COVID-19 pandemic caused lots of companies to expand their use of the cloud dramatically as they sought to support and secure a newly remote workforce, ensure resiliency, and optimize IT costs. In a telling example, Deloitte Consulting's Chief Cloud Analyst David Linthicum recalls clients whose servers went down during quarantine, and no one was allowed into their buildings to restore them. "People realized that having data in the cloud was much safer than having it on premises," he says.

A hallmark of the pandemic spending spree is the cloud's surging share of the enterprise pie. Driven by greater needs for cost efficiency and business continuity, Sid Nag, Gartner Vice President, predicts cloud spending will hit 14.2% of the total global enterprise budget in 2024, up from 9.1% in 2020. With remote work likely to remain widespread even after the pandemic eases, many companies are looking for modern cloud-native applications to help them do their jobs.

However, there might be an unexpected hitch—in the rush to cope with COVID-19, companies didn't always think things through strategically. The emergency expansion of cloud usage also resulted in lots of unanticipated complexity.

With little time for coordination between cloud migration and cloud development teams during the pandemic, some companies now face a lot of work fixing mistakes they may not even realize they made. So even when the pandemic eventually eases, enterprises will likely continue to boost cloud spending to address issues stemming from rushed cloud adoption as well as to drive ongoing business and operational transformation.

MICROSERVICES AND THE CLOUD:

It's Complicated

BY CHARLES BABCOCK

Inside the complex, symbiotic relationship between microservices and the cloud—and the role of cloud-native technologies in creating this powerful combination.

As companies work to modernize their technology approaches, microservices and the cloud are two of the most powerful trends shaping their choices. In fact, the two technologies get mentioned in the same breath so frequently that it would be easy to conclude they're different versions of the same thing. Compounding that misconception: many cloud applications are built as microservices.

Nonetheless, it's useful to remember the cloud and microservices are separate, distinct, yet complementary technologies. Understanding the differences—and also how they work together—can help you better appreciate and take advantage of their respective strengths. Like the Apollo spacecraft and the Saturn V rocket that lifted it into orbit, microservices and the cloud make a powerful team. But each must be optimized for its own role before they can be combined for an ambitious moon shot.

Microservices = application building blocks

A microservice-based architecture is all about creating software as a set of modular, loosely coupled, independently deployable services—which may be written in different programming languages and use different data-storage techniques—instead of a single monolithic block of code. These microservices typically communicate with each other via APIs.

This makes it possible for related microservices to exist in multiple locations instead of all having to run on the same server. Critically, unlike a monolith, where modifying even a small piece of code might require rewriting and redeploying the entire application, a microservice architecture lets developers change or scale specific parts of the application without affecting the rest.

Ideally, each microservice should represent a discrete business process, typically just one facet of a larger and more complex process. Spotify, for example, was an early adopter of microservices to help keep its more than 144 million subscribers happy since at least 2014. The music streamer's microservices architecture breaks down its service into a collection of self-contained single-purpose microservices, such as a search engine, content tagging, and behavioral analytics used to create playlists and power its highly praised recommendation engine. Notably, each microservice is supported by an autonomous full-stack development team, making it easier to scale, as well as easier to test, deploy, and monitor.

APIs let microservices talk with each other

Critically, microservices can increase flexibility and developer productivity, boost innovation, speed time to market, and make it easier to scale. They let developers build applications based on well-defined APIs that render the microservice easily accessible to many different applications, devices, and outside systems—essential if you want your service to work with other services without having to specifically design those interactions.

For example, if Spotify wanted to add a genre, offer a holiday special, or experiment with its pricing structure, it could simply modify existing microservices or add new ones without having to rebuild everything from scratch or disrupting other parts of the existing application.

Setting developers free

That's a really big deal. Microservice architectures promise to significantly improve developer productivity, enabling them to reduce application backlogs and focus on creating innovative new products and features.

In the not-so-distant past, IT teams struggled with what was known as the application backlog: the dozens of change requests to existing applications that piled up as business conditions changed. But one change couldn't be made without affecting many other parts of the monolithic business app, each of which would then require its own adjustments. In the real world, that meant only the most pressing changes were ever implemented.

With microservices in the cloud, however, that

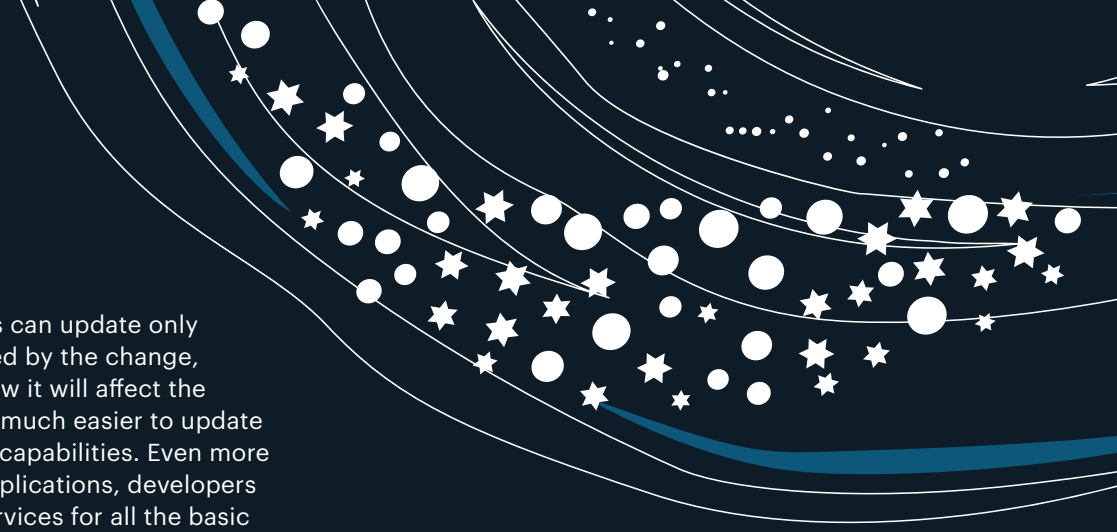

constraint goes away. Developers can update only the microservices directly affected by the change, and don't have to worry about how it will affect the system as a whole. That makes it much easier to update existing systems and create new capabilities. Even more important, when building new applications, developers can rely on cloud-based microservices for all the basic provisioning and functionality, speeding time-to-market by freeing them to focus on what's new and innovative, which drives competitive advantage.

When Amazon opened its online pharmacy in late 2020, for example, the company didn't have to go back to the drawing board to create a retail system for managing drug information and inventory and executing transactions. It already had the display and transaction systems in place, the developers just needed to add new microservices to verify that a doctor's prescription authorized the purchase.

Cloud + microservices + containers + serverless = cloud native

Importantly, simply moving existing monolithic applications to the cloud does not typically deliver the full benefits of cloud computing. While this kind of lift-and-shift cloud migration can give you access to more elastic, more reliable infrastructure, better security, and help avoid the expenses of building and maintaining your own data centers, it won't help you speed development and deliver a better customer experience. And it may not deliver all the possible economic advantages.

To fully optimize a cloud deployment requires going "cloud native"—a concept The New Stack describes this way: "Cloud-native technologies are used to develop applications built with services packaged in containers, deployed as microservices and managed on elastic infrastructure through agile DevOps processes and continuous delivery workflows."

You may have noticed that description cites the importance of containers. Originally popularized by Docker in 2013, containers emerged as a useful technology around the same time that cloud services and microservices were taking hold. As defined by Docker, "A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another."

Simply put, containers act as code-transfer vessels, making it easy to deploy applications on any host, with metadata that tells the new host what they need to properly start up and run the code inside. That enables developers to write code that can be packaged up to run in any location, provided the destination has a server with the operating system required to run the app.

Finally, there's one more critical component of a cloud-native approach: serverless, or Functions-as-a-Service (FaaS). It's called "serverless" because the user doesn't need to provision and configure a server—you get the compute cycles of the function code you want simply by calling for it. (Of course, the function code runs on a cloud server somewhere, but the user doesn't know or care where.)

Just as a microservice is a component of a larger application, serverless functions can provide components of a microservice. Serverless functions work with microservices for such things as granting a permission, executing a variable, or reacting to a software event. The cloud provider, such as Microsoft Azure Serverless or AWS Lambda, supplies the functions for a small per-use fee.

In that sense, serverless takes microservice architecture deeper into cloud-native territory. Serverless functions bring all the typical benefit of cloud computing—scalability, resiliency, data reliability—but users don't have to waste time inventing code already used by many different microservices.

Along with the other elements of cloud-native technology—serverless, containers, agile development practices, DevOps processes like continuous integration/continuous deployment (CI/CD), etc.—the combination of cloud computing and microservice architectures gives companies a future-proof approach to applications and infrastructure that maximizes business speed, innovation, and flexibility. That's why Markets and Markets Research predicts the cloud microservices market will grow more than 22% a year to $1.8 billion in 2023. More to the point, perhaps, modern enterprises are increasingly relying on this combination as the essential platform from which to launch their own technology moon shots.

Charles Babcock is a former editor-at-large at InformationWeek and author of Management Strategies for The Cloud Revolution.

redislabs
HOME OF REDIS

REDISEARCH 2.0 IS HERE!

Build modern applications with interactive search experiences

Get started: redislabs.com/modules/redis-search

Q+A WITH KELSEY HIGHTOWER

GOOGLE'S OPEN SOURCE ADVOCATE ON EVERYTHING CLOUD

BY FREDRIC PAUL

Google Cloud Principal Engineer Kelsey Hightower, recently described by *Protocol* as "one of the most prominent and respected faces in cloud computing and open-source software," talks about how 2020 changed the cloud, the rise of managed services, open source as a service, Kubernetes as a design principle, and the importance of a good story—plus how to whip up vegetarian chili.

How did 2020 affect the cloud?

Kelsey Hightower (KH): It depends on where you're coming from. When the pandemic hit, if you were unable to go into the office, but still wanted to be productive and collaborate, this cloud thing sure sounded like a great idea—the idea that you could log in and meet over Zoom, Microsoft Teams, or Google Meet and collaborate in real-time versus emailing back and forth. People got a chance to leverage those services as part of their core collaboration workflow, and for many organizations, those services were critical to business continuity. It's amazing how many companies rediscovered video conferencing in 2020.

If you invested in a lot of infrastructure that you could no longer use because of the pandemic, you're just sitting there holding the bag. This whole "scaling-to-zero" idea suddenly sounds really interesting because now you have a true forcing

function. It's no longer a theory. It's now a reality.

Were there significant cloud technology advancements during 2020?

KH: In 2020, you saw many providers say, "Hey, your favorite open-source project is now available as a managed service." This ranged from the major cloud providers to the teams behind those projects. Managed services were already on their roadmaps, but 2020 was a great forcing function to get those things shipped and delivered.

On the health front, we saw the rapid development of things like contact tracing. In a span of a few months, companies were able to collaborate and ship something that could be rolled out to the masses, thanks to the existing cloud and mobile infrastructure. So, not necessarily new technologies, but people were able to leverage existing ones that have reached maturity over the years.

What about the cloud is ripe for rediscovery?

KH: What I hear from people can be summed up as, "Make the new thing work the old way." If you look at the complexity of cloud, most of it has come from customers asking for things they're familiar with, like firewall rules, for example. The VPC (virtual private cloud) firewall rule concept attempts to replicate what you do in a traditional data center—restrict communication between IP addresses.

In the cloud, what you really want is a policy. You want a policy that says, "This application can talk to this application." You don't care what IP address the application gets because it could run anywhere, in any region, across the globe. Static firewall rules can never address that particular set of constraints. Once you start targeting IP addresses, that limits you to a single zone, or region, versus something that can work globally. If you want to be able to scale, and get the availability that comes with it, then you want higher level policy at the application layer.

In a data center, you say, "This IP can do these things." Well, any app can assume that IP and start abusing a security policy. IP addresses are a weak form of identity, and not suitable for building robust security policies and enforcing them. By saying, "Make the new thing in the cloud work like the thing in my data center," you lose out on all of those capabilities that can improve your security posture.

> **"If you invested in a lot of infrastructure that you could no longer use because of the pandemic, you're just sitting there holding the bag."**

Most people won't admit this, but you can obtain most compliance certifications while still storing usernames and passwords in plain text files. That's the status quo for most organizations. But in the cloud, we enable you to do away with things like static credentials. We provide tools including identity management, where a short-lived token is explicitly used to grant access for a particular service to another service. That's another example of a cloud advancement that goes under-appreciated because people prefer something that's familiar.

You mentioned the rise of managed services during 2020, why is that important?

KH: People use tools to complete tasks and solve problems. Most tools play a small role in a much larger system. Let's take Redis for example. You start by saying, "I need to store data. Redis supports the data models and access patterns that I wanna use, and it's fast." But you can't just run Redis in thin air. You're gonna need some machines, networking, and security layered up and down the stack. Oh, and you'll need to configure it all so it plays well together. That's just the day one concerns. Once you've deployed Redis you're now on the hook for ongoing security patches and upgrades. This is not what most people signed up for.

This is where cloud providers shine. Cloud providers not only provide ping, pipe, and power, they also provide a set of managed services on top and the expertise required to run them well. Going back to the Redis example, now you can leverage a service based on open standards, and focus on using Redis, not managing it.

From the enterprise point of view, they're saying, "I believe you can do a better job than I can because this is your primary focus." Instead of buying products and running them themselves, organizations are looking for something more, they want solutions. They're looking to buy expertise.

Is that a function of 2020 or was that progression inevitable?

KH: It was inevitable. Whether it's Redis, or firewalls, most people want to consume these things as services. Going forward, more people will re-platform on top of managed services, and delegate a subset of the responsibilities to a provider, like we do with public utilities. Most people don't know how to treat water, or manage the plumbing infrastructure necessary to distribute it, so we put our trust in public utility companies, and govern how they operate, so the rest of us can drink clean water right from the faucet.

What other cloud trends do you see for 2021 and beyond?

KH: Traditionally, privacy and security was an afterthought. You built your product for scalability and ease of use, then you layered on privacy and security when you needed it. But now, privacy and security is an end-to-end requirement and viewed as part of the overall user experience. People want to know what happens when they hand their data over to your platform. It's their data, and they want control

over it, and in some cases, they want to limit what service providers can do with it, or even view it at all.

So, some of the biggest trends we're seeing are around how customers share and manage access to data. Most companies generate and collect a lot of data but struggle when it comes to sharing that data with others, even within the same organization.

The question then becomes, how do you establish trust when you don't necessarily want to trust an entire organization, and ensure data is used only for the purpose intended? Well, a new approach requires we limit access to data to specific versions of software, software that has been audited and signed, and move away from trusting organizations, but rather the software operating on the data. This would give us the ability to revoke access to shared data in ways that were not possible a few years ago, and given the trajectory of cloud technologies, especially around trusted computing, the ability to do this is now possible today.

"Cloud providers not only provide ping, pipe, and power, they also provide a set of managed services on top."

We'd be remiss if we didn't talk with you about Kubernetes. How is its role evolving?

KH: We're learning that you can apply Kubernetes' design principles, which we call the Kubernetes Resource Model (KRM), to other things. Imagine a world where you can define a CI/CD pipeline through a set of declarative configurations. Imagine if Redis had a KRM interface. You could do something like, "Hey, Redis control plane. I

GETTING TO KNOW KELSEY HIGHTOWER

How has 2020 changed how you live and work?

KH: There's something special that happens when people get together to share ideas and experiences. That's how memories are formed: I don't remember the last successful code that compiled on my laptop, but I do remember the last time I had dinner at my favourite Ethiopian restaurant.

2020 was a really tough year for many people, so I committed more time for office hours. I jumped on Twitter and said, "Send me a direct message, we can hop on a Google Meet and chat about whatever you want." I attracted many up-and-coming developers as well as those struggling to work from home. I found that by doing it virtually, I was able to connect with many more people than normal. I think that part is what I enjoyed the most.

And I'll be honest—as someone who travels probably too much, being home with my family was a welcome change of pace. Now I have time for the simple things, like grabbing a burrito with my daughter during our lunch breaks, and catching up on new music.

Is there anything specific that you did with your family that you'd like to share?

KH: I'm really big into financial freedom, minimalism, and living without debt, especially consumer debt. Just like tech can be very empowering for people building things, it's equally empowering when people have control over their financial situation. During the pandemic I was fortunate to have time to invest in my 13-year-old daughter's financial education. I finally got her a debit card and taught her how to use it. I also taught her about fraud prevention, budgeting, and how to think about saving money. She's enjoying watching her allowance go into her bank account and the power that comes with having full access to it.

Later on, we're going to pick a few stocks based on the products and services she consumes as a way to get her into investing, and understanding the companies behind the products she uses. As a young lady, she has reached a level of independence and maturity that I didn't reach until much later in life. That was very special to me.

Since many folks have been eating at home a lot more in 2020—do you have any favorite dishes or recipes you'd like to share?

KH: I actually learned how to dice an onion properly and now the pieces look like squares and not little jigsaw puzzle pieces. Also, I'm vegetarian and I found a way to make chili that people seem to like, using Impossible Burger [plant-based meat].

You need a can of black beans, a can of kidney beans, a small can of yellow corn, and two cans of diced tomatoes with a little green chili in them. Brown half a diced onion with a tablespoon of olive oil, and toss in a whole pack of Impossible Meat, which browns just like ground beef. Dump in a pack of the chili seasoning that you get at the grocery store. Once it cooks for about 10 minutes, start adding the drained beans, diced tomatoes, and corn. From there, you just taste it and add a little hot sauce or whatever else you need to give it that spice you want. In about 15 to 20 minutes you should have some chili that would even fool your carnivore friends.

want a five-node cluster across these zones, and I want this much memory allocated to each cluster member," and then Redis would go and provision itself. That's what we mean by applying the KRM model to other systems.

If you explore the internals of Kubernetes you'd discover that Kubernetes adheres to its own design principles. It's not just one big binary that you install and run. It's actually a set of components that leverage the KRM. Kubernetes is built on the fundamentals of config management, which is built on the fundamentals of promise theory. There are resource definitions, and chunks of code that run in a loop, that work together to turn resource definitions, also known as promises, into running applications that match the desired state.

It turns out you can manage more than containerized workloads using such a model. Now storage appliances and network fabrics are fair game, and this realization is pushing Kubernetes to evolve into a universal control plane.

Kubernetes is a prominent open source project. Do you have any thoughts on the state of open source software?

KH: The bar has risen dramatically since I got involved in the late 2000s. I've grown to see open source as a way to serialize our ideas, make them tangible, and share them with people. Those people form communities and create feedback loops that help us improve upon the initial idea. That's the power of open source, it's the ultimate collaboration model for serializing innovation.

Today we live in the world of managed services, which brings along a new set of expectations, so the question we have to ask is, how does open source accommodate the people who are more interested in using software, not downloading, installing, and managing it? Let's Encrypt is a real-world answer to that question. Let's Encrypt is a nonprofit, backed by a healthy open source foundation, focused

on securing the internet by democratizing access to critical PKI infrastructure—think HTTPS and that little lock you see in your web browser when you're dealing with a secure connection.

For example, if you were to visit takemycreditcard.com, you want to do that over a secure connection, and that requires the website owner acquire and configure a SSL/TLS certificate. Just a few years ago it was common to pay hundreds of dollars for a reputable company to give you an SSL/TLS certificate. You would essentially send them a certificate request and after a few days you'd get an email that read, "Your certificate has been approved, and here are instructions on how to retrieve it." Then it's on you to configure your web servers to use the certificate, and as an added bonus, you got the privilege of doing that every year. Of course, most of us would forget, and our certificates would expire, and force everyone to scramble to get it renewed before the CEO found out.

"The cloud made it easier to provision infrastructure, but it did not eliminate the toil of managing it."

Let's Encrypt comes around and says, "We're gonna automate this end-to-end. We're gonna give out SSL/TLS certificates for free. As a nonprofit, backed by an open-source project, we'll enable others to run our software, and become their own providers, but on day one we're gonna offer a fully managed service so you don't have to." Now we have letsencrypt.org, where most software—even Kubernetes—has the ability to generate industry standard SSL/TLS certificates automatically and keep them up-to-date in the background at no cost to you.

We've gone from free software to free service. That's going to push the bar for what it means to have a successful open-source project, especially for big, popular open-source projects backed by foundations and large organizations.

What are you working on in 2021?

KH: Democratizing infrastructure. It's the original promise of the cloud, but now it's time to move up the stack, and help liberate people from the massive liabilities that come with managing infrastructure, without locking them in. Managed, but open, is the key to democratizing this stuff. The cloud made it easier to provision infrastructure, but it did not eliminate the toil of managing it. In 2021, I'll continue working to figure out how to give people a reliable platform to run their applications, and manage their data, without the overhead of low-level system administration—moving in that serverless direction for as many layers of the stack as possible. Cloud Run and Cloud Spanner are two big areas of focus for me this year.

You're known for being an inspiring figure in the tech world. How do you motivate people to care about their work?

KH: I learned how to tell the story of why I'm doing what I'm doing. If you pick Kubernetes because you read a blog post, and now you want to install Kubernetes, that's not very compelling or inspiring. That's irresponsible and possibly dangerous.

When you have a story, you're able to walk into a room and say, "Hey, our system has gone down 30 times in the last month, because we didn't have the ability to track resource utilization, or restart applications that fail. After installing Kubernetes, I noticed it's keeping our applications running better than our current tool, and we went from 30 to 2 outages as a result."

That's the kind of inspiration most people want. Clarity around a specific problem and a possible solution that is showing early signs of promise.

REDISCOVERED IN THE CL

How the cloud has helped us rediscover, reevaluate, and reengage with influential industries, companies, technologies, and products

BY DAVID NEEDLE

To paraphrase Visa's classic ad slogan, cloud computing technology is everywhere you want to be. Cloud infrastructure and cloud-native software are bringing speed, scalability, flexibility, and cost-effectiveness that's changing the trajectory of technologies and brands large and small in almost every industry. And while the cloud is often seen as enabling a new generation of digital-first companies, it has also helped the world rediscover the power and importance of many existing companies, products, technologies, and industries.

This rediscovery isn't due to the cloud alone, of course. Combining cloud computing's unique benefits with parallel advancements in areas such as user experience, mobility, artificial intelligence, and big data are driving the digital transformation of everything from pizza delivery to supersonic aircraft.

It's even happening in some of the most traditional businesses. "Look at a company like H&R Block, their entire pitch was they offered a tax preparer who could sit with you in a physical office and share a financial relationship," Rebecca Wettemann, Principal Analyst at research firm Valoir, said in an interview. Now, she notes, H&R Block serves clients remotely via the cloud, using cloud-based features like screen sharing and video chat to mimic the rapport developed in an office visit.

From startups to legacy, brands and products, to entire industries, here are five examples of the cloud sparking a rediscovery of what is possible—and what is profitable. The litmus test for inclusion: How has the cloud made us look at this industry, product, or technology differently?

Creative collaboration: real-time brainstorming

Cloud computing gives designers a giant virtual whiteboard to scope out ideas and collaborate with partners around the world to bring the next great inventions to life. One of the companies leading the charge is Figma, which offers a browser-based service where designers collaborate and share reusable components (more than 800 plug-ins), styles, and tools from anywhere they can access a web browser.

As with other cloud-based or Software-as-a-Service (SaaS) offerings, Figma users don't have to download software to their computer to view a file. Designers can collaborate and share a Figma file with a simple hyperlink.

Kimberly-Clark, the company behind such well-known brands as Kleenex, Scott, and Huggies, uses Figma to transform the design process by consolidating all of its tools on a single platform. "With Figma, everything is just there," noted Kimberly-Clark UX Manager Andy Ford in a Figma blog post. "It becomes the single source of truth .... Feedback meetings transformed into collaborative real-time working sessions with our various stakeholders. Everyone started to participate in the making process and could see design updates come to life in front of their eyes."

WHY IT MATTERS: Creative professionals are no longer limited to tools designed for a particular operating system or physical device, and can collaborate in real-time from anywhere.

CLOUD

Manufacturing and industrial design: leveling the playing field

"What does cloud computing have to do with how we fly through actual clouds?"

Blake Scholl, Founder and CEO of Boom Supersonic—which is using cloud-based design tools and procedures to speed the design of a new, faster-than-sound airliner—answered his own question in a keynote at AWS re:Invent 2020: The cloud is "the key to a new era of aircraft design that levels the playing field in aerospace," he said, "saving Boom years of schedule and millions of dollars" by letting the startup company test more designs more quickly. This kind of power was previously available only to governments and huge enterprises, not tiny startups.

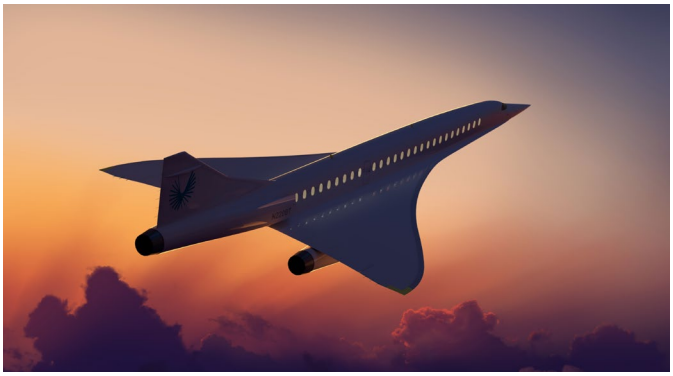

Boom uses AWS' cloud-based high-performance computing (HPC) capability to accelerate the design and construction of its XB-1 demonstrator plane and the commercial Overture airliner to follow.

Boom said the XB-1 project has already generated more than 525 terabytes of design and test data that the company is transferring over to AWS.

Looking ahead, Scholl told CNBC that Boom will use AWS to "coordinate the entire supply chain and how that happens digitally." As a result, Boom will know where every single part is in the supply chain as it goes into the factory for assembly and can trace each part on the aircraft.

Denver-based Boom says it has 30 of its aircraft on pre-order and hopes to launch its first flight in 2021, with commercial service starting before the end of the decade.

WHY IT MATTERS: Industrial design in the cloud is helping startups challenge giant, entrenched industry leaders, speeding the pace of innovation.

Sports stats: not just for the pros any more

Pro scouts and agents might have been the initial market for statistics and in-depth analysis of how players perform, but the cloud changed all that. Now real-time analytics are a broadcast staple and even casual fans have instant access to advanced stats.

Major League Baseball is leading the way using cloud technology to calculate, track, personalize and share stats in real time. After moving to AWS in 2014, MLB Advanced Media switched to Google Cloud in 2020 as part of a move to offer fans more personalized features. Now, for example, prominently featured clips and data will be those most relevant to an individual fan (e.g., based on their favorite team) as calculated by Google Cloud-powered algorithms. And there's plenty of data to work with: MLB can use Google Cloud's data warehousing technology to process and store more than 100 data points from every pitch in every game.

The NBA, meanwhile, is upping its cloud game via

a multiyear agreement to use Microsoft's Azure cloud that started with the 2020-2021 season. As part of the deal, a new streaming platform leverages AI and machine learning from Microsoft to learn fan preferences and recommend content based on a user's favorite team and players. If a player reaches a significant milestone during a game, for example, viewers are immediately presented with video clips of previous players to reach that milestone. The system will also give fans a lot more options, including an alternate audio feed of the crowd sounds (when crowds are allowed back in the arenas) without broadcast commentary.

According to the NBA, the cloud system will make it easier for fans to search its archives and integrate existing NBA products, services, and merchandise. If the system detects a person's interest in, say, Giannis Antetokounmpo, the league's reigning MVP, it will present content and offers highlighting the Milwaukee Bucks superstar. "Our goal, working with Microsoft, is to create customized content that allows fans—whether they are in an NBA arena or watching from anywhere around the world—to immerse themselves in all aspects of the game and engage directly with our teams and players," NBA Commissioner Adam Silver said in a statement.

This data is being used online right now, and the deal to move more of the NBA's systems into the cloud means this sort of detailed info will be more accessible than ever. For example, the league's player stats pages have an API that fans are already using to create their own shot charts and other analysis tools. This is critical as younger fans often prefer to follow the league on social media and online rather than watch full games on TV.

WHY IT MATTERS: Cloud computing has democratized advanced sports analytics, making an unprecedented amount of statistics and analysis available to fans and pros alike. This kind of engagement is critical for sports leagues dealing with fans—especially younger ones, "watching a lot less conventional television," Silver noted in an interview with Stratechery's Ben Thompson.

Media and gaming: distributed production and better streaming

From streaming replacing broadcasting to virtual movie making and more affordable spectacular special effects—not to mention the transformation of video games to deeply interactive and widely shared experiences—cloud computing is revolutionizing the entertainment industry in a wide variety of ways. Paul Cheesbrough, CTO and President of UK media company Fox Corp., told viewers of an AWS re:Invent 2020 keynote that the way media companies stream and use video content "has radically changed in the last five years in the cloud." Increased

cloud capacity is now letting distributed production teams create and deliver uncompressed video—which offers a vastly improved viewing experience—from and through the cloud, with full redundancy. "It enables us to produce live events with less latency, increase reliability and [do it] more efficiently," noted Cheesbrough.

Cloud computing has also helped startups and established studios rediscover the different ways movies and games can be produced and distributed—improving the quality, the economic model, and the subscription value of the delivery services. Universal Filmed Entertainment Group (UFEG) has started a multi-year process of moving its film and TV production from in-house servers to Microsoft's Azure cloud platform. In 2020, UFEG CTO Michael Wise told *Variety* the company wants to take advantage of the platform's "hyper-scale" storage and compute capabilities for studios to tap into on an as-needed basis, predicting the move "will unlock a new way to make movies in a way we haven't been able to do before."

Working with Microsoft, UFEG wants to extend DreamWorks Animation's proprietary production platform to include live-action content that is housed in the cloud, making it easier for partners to connect with them in a more open, standards-based way. A major film can involve dozens of third-party partners and anything that streamlines the process figures to be a huge time-saver.

Smaller production studios are also moving to the cloud. *Like After Death*, a darkly comic interactive "choose your own adventure" film, was produced by Heredia Vision and new media company All Together Now. "With the help of cloud-based software we can build a stage, broadcast a live show internationally, and provide more seats than anything we could do in person," Jessica Ryan, CEO of All Together Now, said in an email.

The scalability of cloud computing also helped create Massively Multiplayer Online Games (MMOGs) and helps enable a wide variety of MMOG features. These

fantasy and role-playing adventures, like *World of Warcraft, Final Fantasy, RuneScape*, and many others, change the gaming experience from a solitary pursuit to one shared with hundreds of thousands of simultaneous players. And while die-hard gamers can spend thousands of dollars tricking out their systems with high-end graphic cards and other performance add-ons, a robust cloud and edge infrastructure is essential to providing a fast and level playing field no matter where players are located.

And don't forget mobile gaming—mobile esports games platform Skillz uses the cloud to host more than 30 million players competing in over 3.5 million daily tournaments.

WHY IT MATTERS: The cloud is remaking both production and distribution of media and gaming, smashing barriers to entry for content creators while creating more choices for consumers.

Electronic healthcare records (EHR): more personalized care

The days when your physician based your care on detailed, personal knowledge of your medical history have largely passed, but that's not as bad as it sounds. Medical records in the cloud promise to help bridge that knowledge gap, eliminating inefficient, hard-to-share paper records while delivering smarter, more personalized treatments and better health outcomes.

But there are still growing pains—like the fact that medical records are stored in a wide variety of formats and are still not all online, making them notoriously hard to pull together. The holy grail is to digitize all these disparate records to build a holistic view of the patient.

One promising effort is the DrChrono platform, an Electronic Health Records (EHR) service with more than 24 million patients under provider care. The DrChrono app functions as a gateway to a backend cloud infrastructure that lets doctors see lab work, a picture of a patient's face, or their medical history, and even write a prescription right from their phone.

AWS and Rackspace provide the primary infrastructure for all of DrChrono's DevOps and business operations. "It's a huge benefit to us because we never have to touch any servers, Amazon and Rackspace manage all that," DrChrono Co-Founder and COO Daniel Kivatinos said in an interview. "If there's an issue with one of their servers, they have redundancy so they can bring another one on quickly and customers never know there was an issue."

As DrChrono and other medical providers grow their reliance on cloud services, they can take advantage of the compliance foundation built by the cloud providers. And there's another huge potential benefit: unlocking what has historically been proprietary databases on cancer patients and other diseases housed in on-premises data centers. Kivatinos says that it is far easier to securely access anonymized data in the cloud than extract something from an on-premises, proprietary database. For example, a real-time analysis of COVID-19 patients in a cloud-based EHR system might show that something they're eating or drinking reduces or eliminates symptoms of the virus.

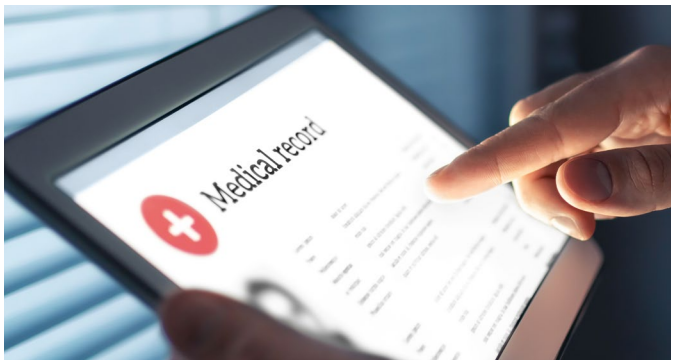

"We launched as a cloud service in 2009 when others were shipping on CDs you had to install," said Kivatinos. That was important, he added, because "doctors are not sedentary creatures—they go home, they travel." But if there's an emergency with a patient at another location, critical information stored on disk might not be available.

Kivatinos recalls a White House healthcare meeting with then Vice President Joe Biden who was frustrated by the difficulty of getting medical records sent from one hospital to another when his son was battling brain cancer. "Biden is very passionate on the subject," Kivatinos said, predicting that with Biden in the Oval Office, "you'll see the cloud embraced way more than it is today."

WHY IT MATTERS: EHR in the cloud promises more informed, personalized, and efficient health care, less time spent on redundant paperwork, and ultimately better health outcomes.

A BRIEF HISTORY

From punch cards to the cloud, a visual timeline of database development

BY HALEY KIM

DBMS

1960s: Charles W. Bachman develops the **first Database Management System (DBMS) in 1960.** His Integrated Data Store, a network model, is created to catalog materials needed for the Apollo program's Saturn V moon rocket. The other key DBMS model of the decade, the hierarchical IMS from IBM, is created for the Apollo spacecraft and General Electric's manufacturing operations.

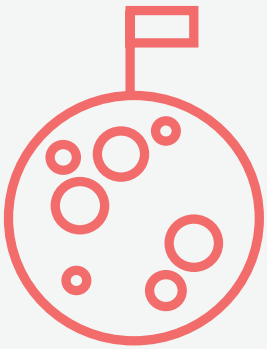

SQL

1970s: IBM computer scientist Edgar Codd outlines new ways to **create and organize databases in "A Relational Model of Data for Large Shared Data Banks".** Fueled by Codd's paper, Michael Stonebraker and Eugene Wong create INGRES (Interactive Graphics and Retrieval System), which uses the QUEL query language—forming one of the key capabilities of a modern database, **the ability to query data.** IBM counters with SQL (Structured Query Language) in 1974.

Punchcards

1880s: Spurred by the difficulties of collecting data for the 1880 US Census, Herman Hollerith creates the **punch card ("Hollerith card"), key sorter, and tabulator unit record machine.** His company eventually merges with three others to form International Business Machines (IBM).

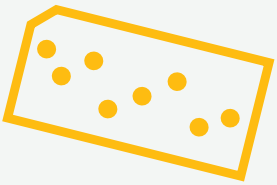

RDBMS

1980s: SQL becomes the de facto database standard. **Relational Database Management Systems (RDBMS),** which store and process structured data, rule the market.

OF DATABASES

NoSQL

2000s: As the internet continues to grow, the first NoSQL database as we know them today arrives in 2009. **NoSQL databases become popular due to their speed working with large amounts of unstructured data.** Four key types of NoSQL databases emerge: graph, key-value, document store, and column store.

ML and IoT

2010s: Enterprises move mission-critical database apps to the cloud in search of operational simplicity, speed, elasticity, and agility. **Machine learning and the Internet of Things generate more demand for data storage and analysis**, spurring development of cloud databases.

The Internet

1990s: The rise of the internet powers growth in the database industry. The RDBMS model, designed to manage the data of a single organization, was unprepared to handle data from waves of internet users. As performance declines and maintenance increases, developers look for a new data model.

The Cloud

2020s: The future of the database market is in the cloud: **by 2022, 75% of all databases will be migrated or deployed in the cloud** and only 5% will go back to on premises, according to Gartner.[1]

[1]GARTNER, "2020 MAGIC QUADRANT FOR CLOUD DATABASE MANAGEMENT SYSTEMS," DONALD FEINBERG, MERV ADRIAN, RICK GREENWALD, ADAM RONTHAL, HENRY COOK, 23 NOVEMBER 2020`

MEET THE OPEN SOURCE REDIS CORE TEAM

A roundtable discussion with the five people most responsible for the future of open source Redis

BY FREDRIC PAUL AND HALEY KIM

OSS Redis Team

Yossi Gottlieb

CHIEF ARCHITECT
REDIS LABS

Oran Agra

SENIOR SOFTWARE ARCHITECT
REDIS LABS

Itamar Haber

TECHNOLOGY EVANGELIST
REDIS LABS

Madelyn Olson

SENIOR SOFTWARE
DEVELOPMENT ENGINEER
AMAZON WEB SERVICES

Zhao Zhao

SENIOR ENGINEER
ALIBABA

In June 2020, Redis creator Salvatore Sanfilippo stepped back from his role as maintainer of the open source Redis project. In his place, a new community-driven governing structure took shape, led by Redis Labs' Chief Architect Yossi Gottlieb, Senior Software Architect Oran Agra, and Technology Evangelist Itamar Haber, along with Amazon Web Services Senior Software Development Engineer Madelyn Olson and Alibaba Senior Engineer Zhao Zhao.

We caught up with this five-member core team to find out how the new light governance model has been working, what's next for open source Redis, and much more.

How is the new light governance model working?

Itamar Haber: We've given a lot of thought to making Redis' new model serve the needs of the community and the developers involved in the project. While it may be too early to say for sure, it looks to me like the new model is already proving to be effective. The number of contributors to the project is steadily growing, as are the overall number of contributions. There's more interaction between developers and it definitely feels like that momentum is increasing. There's still a lot to improve, but in my opinion that's an amazing beginning.

Oran Agra: From a technical perspective, we can each merge simple things into Redis with a review from one core team member, but major decisions that affect backwards compatibility, architecture, or liability for the future (like new commands), require a full consensus or at least a majority and no objections.

How is the team working together? Did you all know each other before joining the team?

Zhao Zhao: We have many ways to work together, mostly through GitHub. Slack is also a frequently used tool, and we hold a Zoom conference every two weeks. I met most of the team at RedisConf19, it was like a meeting of netizens.

Itamar Haber: Almost all of the work is done in the source code

repository. We make it a point to discuss everything there and in the open, so we can get the community's perspective as well as have an immutable record for posterity.

Have there been any issues with geography and not being able to meet in person during the pandemic?

Madelyn Olson: One benefit of being globally distributed is that we have almost round-the-clock coverage of responding to Redis issues and pull requests. The downside is that it's really hard to find a time everyone can sync up. Usually someone is up late into the night, and right now that someone is unfortunately me.

Zhao Zhao: We come from three different countries, so it's not easy to find a good meeting time when everyone is working. At first the time was 11 p.m. in China, so I worked as an owl for a while. Then we changed to 3 p.m. in China, thanks to the other core team members' understanding.

How is the new model influencing the evolution of Redis?

Yossi Gottlieb: Time will tell, but for now I think that being able to share the burden of maintenance and have more people actively involved increases the attention span. So, smaller features, fixes, or cleanups that would slip under the radar previously can now be handled. Of course, attention and capacity are still scarce and we still need to continuously prioritize and reject things.

Madelyn Olson: The new group of maintainers all bring unique perspectives to Redis, so we'll see new features because we have a new group of people. The current set of maintainers also all have a background in managing Redis in cloud providers, so we'll bring all that knowledge to harden Redis further.

> "Openness isn't just about the source code, but also in the thought and decision-making processes in the project."
> —Itamar Haber

Itamar Haber: I think I speak for everyone in the core team when I say that the new model will help drive the evolution of Redis by focusing on being more open. Openness isn't just about the source code, but also in the thought and decision-making processes in the project. By opening up and encouraging the community to participate, we hope to not only improve Redis itself but also grow the community of developers around it.

Did Salvatore pass along any Redis wisdom to you all?

Yossi Gottlieb: I think it'd be more accurate to say Redis has his wisdom embedded in it, as in his approach to defining problems and coming up with solutions. For example, oftentimes the real challenge is to avoid complex solutions, and the way to get there is to define the problem, what we're trying to solve, and—not in any way less important—what we're *not* trying to solve.

Oran Agra: One point I learned from Salvatore is it's better to refuse some features and accept some limitations or even inefficiencies in the software than have it be overly complex, with a tendency for bugs and difficult to maintain. The second point is to always think about how something is actually going to be used. If there is maybe a better possibility (even if it's "mathematically" the wrong way) to achieve something that's almost as good as perfect but at a lower cost (in performance, memory, or code complexity) then that's probably a better solution.

Zhao Zhao: I really like Salvatore's development philosophy for Redis: simple is better, using simple methods to achieve complex functions, which can enhance the maintainability of the code, allow more people to participate, and reduce the occurrence of bugs. The code should be self-explanatory, so that it can be understood without too many comments.

Itamar Haber: His spirit lingers and his teachings echo in our ears. Whenever I need to make a Redis-related decision, I ask myself WWAD (what would antirez do)? Salvatore taught us (and I mean "us" in the broadest sense) a lot, from the beauty of code to the joys of simplicity. However, in my opinion, the most important lessons are the tasks that he left us that need to be addressed immediately. These aren't new features or bug fixes, but rather listening to and addressing the needs of the already-huge and ever-growing community of Redis users.

What Redis developments are you working on right now?

Yossi Gottlieb: We've been finalizing Redis 6.2 and looking at what last-minute features and improvements

> "Redis is evolving all the time, and so does the way users see it and how they use it."
> —Yossi Gottlieb

can still make it in the release. Redis 6.2 includes a number of new commands and command improvements that fill some gaps, and some core improvements that may be less visible to most users.

Oran Agra: I've just finished a project I was working on for a few months, which was to make Redis immune to crashes due to corrupt payloads provided with the RESTORE command (#7807).

Zhao Zhao: Mainly Redis-related open source projects, like redis-shake, a tool for synchronizing data between two Redis instances. And some interesting open source modules, like RediSearch, RedisJSON, RedisBloom, etc.

Madelyn Olson: I'm currently working on an improvement to failover, so that it's easier and has less downtime than the current implementation. I'm hoping to spend my time moving forward innovating in cluster support, delivering several incremental features that people have been asking for, but I don't want to over-promise anything!

What do you see as OSS Redis core team's biggest challenges going forward?

Yossi Gottlieb: Redis is evolving all the time, and so does the way users see it and how they use it. It started as a "data structure server," popularized as a cache, and now more and more users use it beyond that for various real-time applications. I think the biggest challenge for Redis (and thus also the core team) is to continue to support this evolution.

Oran Agra: We have so many ideas for improvement and not enough time to get to it all. I think our main challenge is to increase our development throughput by recruiting dedicated community members to take a more active part in developing big features that take months to complete.

Madelyn Olson: Redis has had a large impact on the world, and there is an expectation Redis will continue to be ground-breaking and innovative. It's a daunting expectation to live up to, but I think with our shared knowledge and experience we'll be more than up to the challenge.

Itamar Haber: The first big challenge is increasing the involvement of capable developers in the project. Ideally, I'd like to see the core team double in size so we'd have enough attention and resources to make serious progress. Because core team members aren't born, but rather step up from the community, it goes without saying that the community itself needs to grow.

I think that the next big challenge is making people more aware of what Redis can do. This may sound funny, with Redis being the most popular data management backend, but I often feel that most of its users are familiar with only a small part of its nature and capabilities. This is a challenge for the core team because we lack feedback from users, both about existing features as well as the direction we'll take Redis in the future.

Perhaps the most significant challenge is dealing with the technological legacy in the project. Redis, like any other software, accumulated technical debt and some aspects of the project need major TLC. One example is Redis Sentinel, which—in a nutshell—provides high-availability and discoverability to single-instance Redis deployments. Sentinel has been a part of the project for a long time, and while it has served its purpose well, hindsight is perfect and lessons were learned. As much as I'd like to see Sentinel get more attention, the challenge is balancing that with the rest of the project's requirements.

How would you characterize the state of open source software right now?

Yossi Gottlieb: Open source software and, more broadly, the open source development model, have emerged as winners. We need to remember that in the past, many users were suspicious about open source software and needed to be convinced it's as good, and usually better, than proprietary products.

But the models of distributing and using software also changed. Software-as-a-Service (SaaS) is becoming ubiquitous, especially with big cloud vendors, shifting the center of gravity, and I believe the ecosystem is seeking a new equilibrium.

Madelyn Olson: I think open source is evolving rapidly because the way people view and contribute to open source projects has changed. A lot more individuals and companies are building entire stacks on top of open source software and are contributing back to the projects since they are invested in their success. This includes companies like AWS, my employer, which has recognized the importance of open source software and actively contributes to a wide range of products, including allowing me to dedicate time to the Redis project. I think the new governance model for Redis is a good example of how a wide range of interests have come together to improve an open source project.

Itamar Haber: As long as humans write software, open source will remain. When we reach the singularity and AI writes software that we can't understand, that will probably be the end. Until then, open source software is here to stay, even if only as a channel for human creativity.

> "I think the new governance model for Redis is a good example of how a wide range of interests have come together to improve an open source project."
>
> — *Madelyn Olson*

THE TIME HAS COM

REAL-TIM

RETAI

What it takes to create the instant, seamless, omnichannel retail experience customers demand

ME FOR

E

L

BY AARON SUN

During the last decade, widespread access to fast, reliable internet access, the ascendance of smartphones and broadband mobile connectivity, plus the evolution of innovative online services have combined to create the era of real-time everything. Amazon and other leading retailers and direct sellers have transformed consumers' retail expectations to include fast and flawless online and omnichannel shopping experiences, backed by efficient and reliable fulfilment. Salesforce, for example, has found that consumers also demand personalization: almost three quarters expect companies to have connected processes and understand how individual customers use their products or services.

But a great shopping experience doesn't mean a thing if it's not delivered in real time. In this era of instant gratification, nine out of ten US consumers will abandon a retailer's website if it is too slow, according to Retail Systems Research (RSR). And it's not just that—almost three out of five US consumers would leave a poorly performing website to buy from a competitor, one in five would never return, and one in seven would vent their frustration on social media.

At the same time, providing a competitive online and omnichannel retail experience has never been more complex. A successful real-time retail strategy must provide a consistent real-time view of inventory, managing updates from stores and enterprise systems to give customers and staff a clear, accurate view of stock availability. It must support supply-chain optimization to enable more efficient fulfilment, cost management, and planning. It also needs to be resilient and scalable to satisfy consumers' expectations and manage periods of increased demand.

Clearly, the right technology is critical to delivering real-time retail, yet retailers often tend to overlook a crucial step when developing a real-time retail technology stack: investing in the data layer.

HOW RETAILERS CAN DELIVER REAL-TIME RETAIL

A shiny new website or user interface facelift can highlight a retailer's digital transformation, but these are surface-level changes. To truly deliver a stellar customer experience, retailers must build a data layer capable of supporting real-time operations. This requires not only fast database performance to keep retail applications responsive, but also capabilities that make data more accessible to customers. The data layer must also support the dynamic scalability required to cope with periods of increased demand without having to pay for those resources when they're not needed. Here are three best practices organizations can adopt to deliver retail-time retail:

1. Embrace the need for speed

Organizations should start by asking whether their existing data layer offers the low latency and high throughput that can deliver the instant gratification customers demand in the era of real-time retail. The solution might be as simple as adding a cache to speed up page-load times, or supporting use cases like full-text search, graph processing, and recommendations with faster-performing databases. It can also be helpful to explore paradigms such as event-driven architectures, where data processing is triggered in response to changes in software state.

2. Do more for your customers with real-time data

Faster performance for mobile apps and websites is only the beginning. Retailers can make the user experience better by ensuring that basic steps in the customer journey—showing products in stock, searching through previously bought items, reserving items for pickup—are convenient, user-friendly, and efficient.

For example, displaying whether an item is in stock across different physical stores requires a real-time view of inventory that can accurately resolve conflicting updates from multiple regions. Successfully optimizing the digital supply chain can also result in more efficient order fulfillment, cost management, and forecasting across the business. Using the latest real-time data ensures that communications with customers related to order status or stock availability are not just timely, but also accurate, and can provide useful insights to drive marketing and promotional activity.

3. Expect the unexpected

A real-time retailer's omnichannel and supply chain systems must also be able to scale up when required to meet increased demand around predictable, major events of the retailer's year, such as Black Friday and Cyber Monday, as well as special events such as the release of limited-edition items. The RSR study revealed that two-thirds of US consumers surveyed experienced a slow or broken website during Black Friday 2019, and about half of this group abandoned those websites to buy elsewhere.

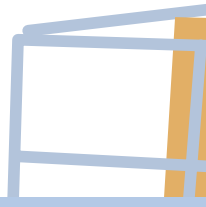

6 Qualities of a Real-Time Retailer

1. Instant gratification
Implement real-time, personalized services, based on consistent customer data across all points of contact.

2. We hear you
Make each step of the customer shopping experience highly responsive, user-friendly, and efficient.

3. Stock up
Enable a consistent view of stock availability across brick-and-mortar stores, warehouses, partners, and online channels.

But retailers know that you can't always prepare for traffic spikes in advance. So retailer websites and fulfillment systems must also be able to scale up to meet consumers' expectations during unpredictable surges in demand. Think online influencers endorsing products or unexpected external events, such as when demand for bicycles and home improvement products soared during COVID-19 lockdowns. These surges can also be the consequence of unexpected natural disasters like hurricanes or fires.

For a retailer to succeed in creating a real-time online or omnichannel retail operation, it must meet *all* of the consumer expectations present in the era of real-time everything. The factors that determine the success or failure of real-time retail all depend to some extent on the retailer's ability to use customer and product data effectively and consistently across all service channels. A world-class data layer alone is not sufficient to deliver market-leading real-time services, but a sub-standard data layer will pose huge challenges for retailers trying to develop and implement effective and compelling real-time offerings.

IT'S NOT TOO LATE... YET

Delivering real-time retail requires responsive, personalized, and convenient customer-facing processes. Customers want insight into your inventory in real-time, but they couldn't care less about the technology required to make it happen. Although many retailers have prioritized other real-time technologies, a highly performant data layer can make an enormous difference and help drive long-term customer engagement. Develop your strategy, explore test cases, and start deploying the data layer technology required to make it happen—or watch your customers leave to shop elsewhere.

EDITOR'S NOTE:

This article was adapted from a Redis Labs white paper: Retail in the Era of Real-Time Everything. Go to RedisLabs.com/retail-wp to download the full white paper.

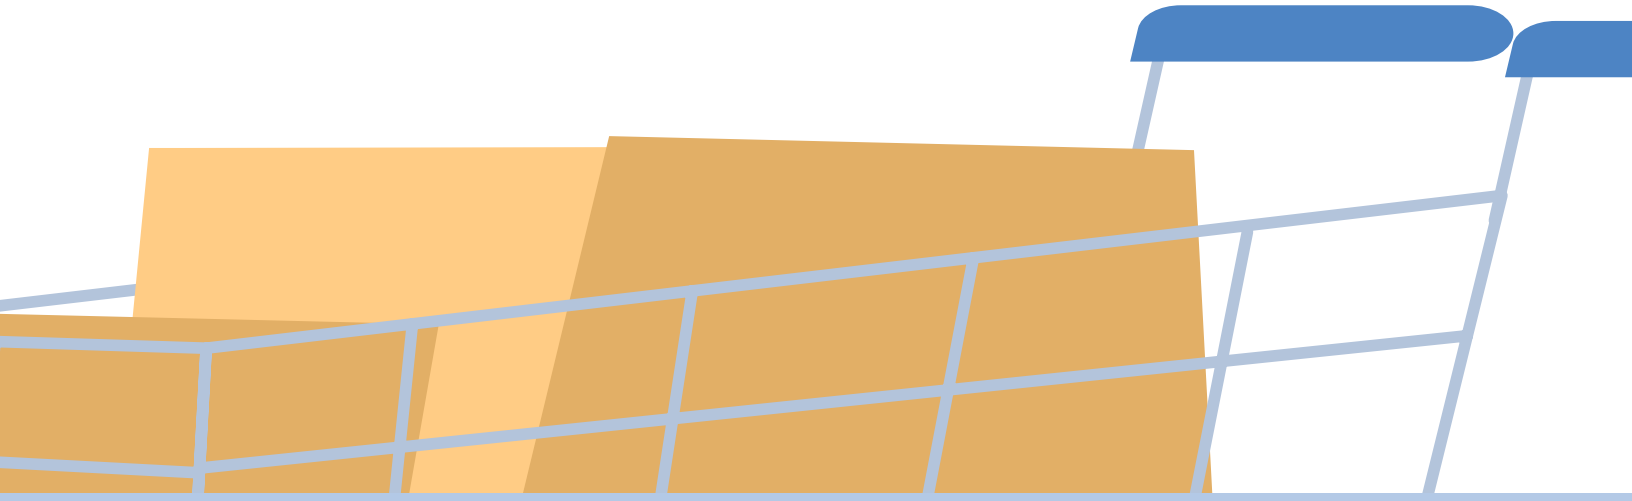

Successful real-time retailers need a half-dozen specific capabilities:

4. See clearly
Optimize supply chain management with end-to-end visibility, beefing up cost management and forecasting capabilities.

5. Last mile
Install fast, flexible, and responsive fulfillment processes that can quickly adjust or adapt to changing logistics.

6. Be flexible
Deploy online systems capable of scaling up or down in response to predictable and unpredictable fluctuations in demand.

NEXT GEN REDIS LABS

From DevOps to sales, meet some of Redis Labs' youngest members

BY HALEY KIM

Joining a startup is an adventure for anyone at any stage of their career, but especially so when it's your first or second full-time job. As the first person hired for my role as Content Producer at Redis Labs, I've been given the chance to shape my passions to business initiatives and launch first-time projects. But my story is hardly unique: these six young Redis Labs professionals represent the next generation of tech leaders, innovators, and builders.

Nir Avrahamov

Solution Architect
Office: Mountain View, Calif.

Nir Avrahamov knew he enjoyed the problem-solving and creativity that comes with computer science, but he didn't think coding all day would be very appealing. In the summer before his senior year of college, he chanced upon a developer-relations engineering internship. He loved it, and was hired by his employer's solutions engineering team for his first full-time job.

Now he's a Solution Architect at Redis Labs, where he specializes in helping prospects and customers find the right solution for their business needs. That means a lot of teamwork—he works closely with his Account Executive and Sales Development Representative (SDR) on multiple accounts, while also collaborating closely with his Solutions Architects team. As the youngest member of that team, he gets plenty of help and guidance from his more experienced coworkers. That help includes both in work and outside activities—he recently picked up surfing from one of his teammates.

Tal Ayalon

DevOps Developer
Office: Tel Aviv

Tal Ayalon knew he wanted to be a software developer since he was 15 years old, when his dad bought him a programming book about C. He focused on computer science in high school and then joined the Israel Defense Forces (IDF) as a DevOps Engineer, spending four years learning multiple technologies. He was thrilled to accept his first job at Redis Labs after leaving the army in November 2020—Redis was already his favorite database!

DevOps, which combines development and operations teams, encourages team members to work across an application's entire lifecycle (instead of just development or operations), with the goal of speeding development processes and software releases. Still new to Redis Labs, Tal has focused on onboarding and learning about the company, and he's excited to start working on Redis Enterprise. Eventually, he hopes to mentor new employees. His advice for other young professionals? "Because you're young, you need to prove you're excited to join the workforce—so work on projects and technology you're passionate about!"

Chi-Lan Tran

Partner Marketing Manager
Office: Mountain View, Calif.

San Francisco Bay Area native Chi-Lan Tran knew she wanted to build a career in the tech industry. After studying media studies and communications at the University of California, Berkeley, Chi-Lan found her calling in marketing. She began her career in event marketing, focusing on developer and user conferences, before shifting her focus to partner marketing at Google Cloud. She landed at Redis Labs in late 2020 as a Partner Marketing Manager, where she works cross-functionally with the marketing and partner team to align on joint marketing strategies with Microsoft Azure, AWS, and Google Cloud.

Chi-Lan loves working at Redis Labs because she gets to work with other passionate, driven team members who have created a strong team culture, which was apparent even during her remote onboarding. While her puppy Zuko might be her only in-person teammate right now, she's looking forward to building her career in marketing at Redis Labs, with plans to eventually earn an MBA.

Alex Ette

Corporate Sales Representative
Office: London

Alex Ette knew he had an interest in persuasive language after studying English in university, so after graduating he looked for a job where he could flex his communication skills. He found his calling in sales, first at Red Gate Software and Okta before landing at Redis Labs as an SDR in January 2019. He spent his time progressing leads and setting meetings for his Account Manager while working on his next step in his sales career.

That move came in December 2020, when Alex was promoted to Corporate Sales Representative. Now on the other side of the sales process, he's the one taking meetings and closing deals. Working in sales means talking to all kinds of people, from those who love Redis to others who know almost nothing about it, and figuring out how to communicate and provide what the prospective customer is looking for. It's a challenge Alex looks forward to every day.

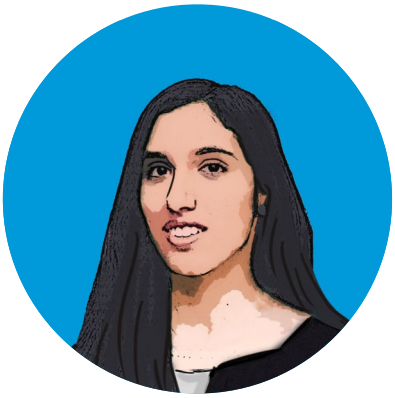

Ortal Golzar

Test Automation Developer
Office: Tel Aviv

For Ortal Golzar, a good day includes singing, guitar-playing, and learning something new about automation development. After an introduction to computer science in high school, Ortal spent two years in the IDF as a QA and Automation Developer before becoming a Test Automation Developer at Redis Labs in March 2020.

At Redis Labs, Ortal spends her days working on different projects on the Management Platform scrum and running automation cycles, verifying there aren't any bugs or issues with the code. She's still just as passionate about automation development as she was when she was in the army, and she's hoping to learn more about API automation, UI automation, and cluster automation as she progresses through her career. Her advice for finding a job in tech? "Study what you like most, work on your own projects, and try to learn something new every day."

Phoenix Moomaw

Manager, Sales Systems
Office: Mountain View, Calif.

At one of his first jobs, Phoenix Moomaw was encouraged by his boss to drop out of school. It was a tempting offer to dive headfirst into a career in tech, but Phoenix ultimately decided to finish his degree in business management. It all worked out in the end—he's now working with that same manager at Redis Labs. As Manager of Sales Systems, Phoenix spends his time speaking with different Redis Labs stakeholders on the internal platforms they use (like Salesforce) and how they can be improved for better productivity. That also means spending many hours tweaking user interfaces and building out systems.

Sales operations affects many different departments and employees, and understanding how his work contributes to the success of the entire company is one of Phoenix's favorite aspects of his job. Building productivity and a great work environment in a company is one of his passions, and he's taking his management ambitions to the next level this year by building out his team.

History is made by Davids **challenging** Goliaths

Gartner, Magic Quadrant for Cloud Database Management Systems,

November 2020. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

Learn more at
redislabs.com/gartner

The "Real" Meaning of Common Cloud Terms

John_Doe_CEO

@Team! Heads up: the Board says we need to be "cloud-native" by the end of the quarter. As you know, I don't have time to dig around Google and there are just too many cloud terms to learn. Make me a list and explain these to me like I just dropped in from outer space. Thx.

Jane_Smith_CMO

Sure boss! We did hours of research to make you this glossary of common cloud terms to help get you up to speed:

- **Caching:** That place where all the hot insights hang out until somebody flushes them by mistake.
- **Cloud backup:** When a cloud falls, you have to put it back up better than before.
- **Cloud bursting:** An automated technique for running underperforming software simultaneously in-house and in the cloud. Used to make poor user experiences infinitely scalable. Note: Has nothing to do with an overfilled cloud exploding.
- **Cloud-native:** Startups engineered for sky-high stock market valuations by rejecting every IT advance made before 2005.
- **Cloud storage:** The place where your individual socks go before the entire pair disappears for good.
- **Elasticity:** What happens to the bodies of cloud architects who embrace yoga.
- **Hybrid cloud:** The love child of a data center and a Prius.
- **Instances:** Dates you went out with a few times who you'd rather forget about.
- **Lift-and-shift:** What Amazon warehouse workers do.
- **Microservices:** Tiny gestures that reveal how much you care for your customers in the cloud.
- **Multicloud:** If you want to make it to cloud 9, you also need clouds 1 through 8.
- **Multitenancy:** Pre-pandemic, six software engineers sharing an apartment in Silicon Valley. Now, six software engineers sharing a condo on the beach in Cozumel.
- **Partly cloudy:** A meteorologist's way to say "hybrid cloud."
- **Private cloud:** If there were such a thing, we wouldn't know about it, would we? But we do, so....
- **Serverless functions**: Responsibilities of your job that do not serve a valid purpose. Alternatively, a business dinner (remember them?) with a buffet instead of waiters.
- **SLA:** A contract between you and your cloud provider that defines how screwed you are when the cloud goes down. For entertainment value only.
- **Software-as-a-Service (SaaS):** An application program that feels remorseful about its technological privilege and tries to make up for it by dedicating itself to the good of all software-kind.
- **Virtual private cloud:** A public cloud that pretends to be private.
- **Workloads:** What many tech companies ask their employees to scale up while WFH.

Scale with Redis Enterprise Cloud as a fully managed service on Google Cloud Marketplace today!

Run • Scale • Manage

redislabs
HOME OF REDIS

Google Cloud

redislabs.com/google

Future-Proof Your Cloud Strategy

Redis Enterprise Cloud

Unify hybrid cloud and multicloud deployments

Consistently deliver instant experiences

Scale your business at light-speed

redislabs.com/cloud