White Paper

The Future Is Here: Multicloud for the Distributed World

redis

Table of Contents

Introduction

The shifts we have seen in society over recent years have changed the way we act as individuals and organizations. On-demand services, virtual and digital experiences, and cloud-first approaches have all grown; the way technology is built and works has morphed as well.

This shift means we need to rethink the architectures we use to build products—distributed applications and modular services become the norm when we're serving a global market. We can no longer rely on a single monolithic approach to technology. At an infrastructure level, this reduces our reliance on a single hyperscale cloud provider.

No longer can applications sit lazily at the center, waiting for us to interact with them. Today the norm is an architecture that is both broad and deep, relying on various vendors at the same time while having aspects that work in a central location, and others that work on the edge. Edge computing, IoT, multicloud, on-device processing—all terms that were mainly the purview of technologists only a few short years ago and now are the default for every industry in every country.

In this paper we aim to articulate this change more fully, explain the architectural impacts of the change, and posit why choosing a broad fabric to drive consistency of data and flexibility of architecture is now non-negotiable.

Pre-Covid-19: What got us to this point?

It seems strange to be writing, deep in the midst of the Covid-19 pandemic, of things that fundamentally changed organizational paradigms before Covid-19. Indeed, Covid-19 seems like such a turning point that one can forget how much the world had changed in the decade before any of us had even heard of Covid-19.

But change it had, and it is worth reminding ourselves of how far we had come following the Global Financial Collapse (GFC) to give us a sense of perspective and a baseline to measure subsequent changes.

Let us take a moment to look at the macro changes that we have witnessed in the past decade or so.

The Global Financial Collapse, capital constraints, and global competition

While caching is commonly used to improve application latency, a highly available and resilient cache can also help applications scale. Offloading responsibilities from the application's main logic to the caching layer frees up computing resources to process more incoming requests. Enterprise caching will provide a sub-millisecond response under an extremely high workload with high volumes of data across global deployments that require five-nines availability. Since it's an integral part of mission-critical application architectures, it has higher consistency, availability, performance, scalability, and geo-distribution capabilities than standard caching.

Caching has a proven track record of helping companies of all sizes improve application performance. In fact, many companies depend on caching to meet their customers' expectations and SLAs. Caching should be considered mission critical when you and your company rely on caching to meet key business goals or deliver key business services. When your caching is mission critical, enterprise-grade caching is required.

A demand for agility

Partly because of the aforementioned tectonic economic shifts, and partly because of a competitive landscape more challenging than ever before, organizations are increasingly focused on being more and more agile.

Enterprises are looking past their status quo and searching for ways to unshackle themselves from traditional large enterprise bureaucracy and processes to unlock innovation and hence help them to compete.

One of the best-known case studies of this move from process and construct to agility and flexibility is that of Proctor and Gamble. (P&G) set the aim of acquiring at least 50% of its innovations from outside the company. This isn't a reflection on innovation per se; rather, it is a tacit admission that it is the lack of agility, created by regimented and hierarchical workplaces, that often limits the growth and profitability opportunities for organizations.

This lack of agility is now well recognized, and companies globally are searching for ways to enable agility to occur across the organization. This drive for agility is well demonstrated by the unprecedented levels of individual employees and business units making their own technology decisions.

No longer will centralized command and control of enterprise IT be acceptable. Modern organizations have no option but to enable and empower a far more modular approach to IT. More on this later.

The rise of the millennials

Much has been written about the rise of the millenials, those people born around the 1980-2000 time frame. While there is some argument about the specific time period the term "millennials" relates to, what all researchers do agree on is that this group represents an incredibly large portion of the population in all markets across the globe.

Millennials have grown up with a very different understanding of the workplace, and the world in general. The first generation to grow up with the regular use of digital media and the global ease of communication it brings, this generation puts particular requirements on organizations as they enter the workforce. Millennials don't regard technology so much as a tool, but rather as an extension of themselves. As author Don Tapscott wrote in his book *Growing Up Digital: The Rise of the Net Generation*:

> Computers and other digital technologies, such as digital cameras, are commonplace to N-Gen members. They work with them at home, in school, and they use them for entertainment. Increasingly these technologies are connected to the internet... . Constantly surrounded by technology, today's kids are accustomed to its strong presence in their lives. Today's kids are so bathed in bits that they are no more intimidated by digital technology than a VCR or a toaster. And it is through their use of the digital media that N-Gen will develop and superimpose its culture on the rest of society. Boomers stand back. Already these kids are learning, playing, communicating, working, and creating communities very differently than their parents.

The catchy phrase relating to millennials (or "N-Gen members," as Taspcott called them) was that they grew up "bathed in bits." This factor alone indicates the expectations of this generation when they enter the workforce and the unique challenges that this generational shift introduces

Technology democratization

The need for agility, the demands of millenials, and technological improvements have led to a wholesale democratization of IT. If we think back to only a few decades ago, technology was in the hands of a very select few technicians who generally worked in cold, dark basements. The rise of the personal computer changed that and resulted in the ability of an organization to put a "computer on every desk."

This spreading of technology was just the very start of the trend, and it has been the growth of mobile access, smart devices, point applications, and cloud computing that has led to a point where data can be put in the hands of any employee, anywhere in the world, nearly instantaneously.

Cloud computing has also had an impact by moving the buying decision for technology products from centralized IT to individual business units. In the past, a technology purchase—be it hardware or software—generally required capital expenditure and hence a rigorous approval process. The rise of "as a Service" (aaS) business models means that organizations can acquire technology on a utility basis. A business unit has the ability to build itself a virtual data center, leverage a virtual development platform, or use applications just by using a cloud provider and a credit card.

All of these macro changes have resulted in a fundamental shift for organizations: a move to a different way of working and different requirements across all of their operations. While these changes came about before Covid-19, the pandemic has only accelerated and broadened their impact.

The Covid-19 impact

The Covid-19 pandemic is not yet over, but even at this point in its progression, it is clear that we live in a changed world. All of the shifts that have been detailed above have both broadened and deepened because of the pandemic.

Deepening the use of digital

We like to think of some companies as early adopters and others as followers or laggards. But this characterization ignores the nuance around digital adoption. If we think of the most "digitally savvy" companies, we often point to Silicon Valley-style organizations such as Google, Apple, and Amazon.

But even these organizations, as digitally invested as they are, previously relied on physical meetings, in-person conferences, and traditional sales approaches.

Suddenly, in a matter of weeks, these organizations had to move to leveraging digital tools for these activities. Whereas Amazon used to hold a user conference that attracted tens of thousands of attendees, suddenly they were putting on virtual events. While Apple relied, at least in part, on glitzy retail showrooms, they became a digital-only storefront. And whereas Google employees luxuriated in offices supplied with food, onsite massage, and laundry services, suddenly they had to work out what working from home actually looked like.

All of a sudden those digital tools that these early-adopter organizations used in parts of their business were expanded across every single operation and area.

Broadening the extent of changes

But it wasn't just the already digitally savvy companies that jumped on this journey. Prior to the arrival of the coronavirus, digital disruption was unquestionably on people's lips, but outside of those sectors directly looking to disrupt with technology, digital was seen as a nice addition and not something core to the business.

Covid-19 resulted in digital being adopted by a far greater range of organizations. Whereas telemedicine was, prior to the pandemic, something that is mainly seen in aspirational videos and science fiction movies, in the space of weeks we saw the norm move from in-person consultations with clinicians to virtual appointments.

A similar shift was seen in the grocery industry. While online grocery shopping had been something that only a limited few enjoyed (generally those with disposable income), Covid-19 made this practice the norm.

And while virtual meetings were certainly held before Covid-19, during the pandemic every organization moved to totally virtual models, leveraging Zoom and other similar platforms.

Suffice it to say that our pesky virus has resulted in a tectonic shift to digital delivery, a shift that fundamentally challenges traditional models of technology.

The challenges for technology in keeping up with new paradigms

Traditional technology approaches mirrored the conservatively paced and centralized model that traditional businesses worked under. As the shape of modern organizations has changed over the past decade, and as that change has accelerated in response to Covid-19, these traditional models have been found to be lacking.

The new requirements for technology are analogous to the new organizational models: A distributed workforce mirrors a distributed application. Acquiring talent on a project or skill-specific basis mirrors leveraging multiple tools from multiple vendors. Having employees work around the clock wherever they might be situated mirrors thinking of applications as "always on."

This new distributed and modular approach toward technology was predicted nearly a decade ago. In what can be seen as a prescient post on The Composable Enterprise™, Jonathan Murray, former CTO of Warner Music Group, wrote about what technology would have to look like in order to deliver what organizations need in terms of speed and agility. Based on his crystal ball gazing and the initial changes he was seeing, Murray described the Composable Enterprise this way:

> Business functions, processes, organizations, supplier relationships, and technology need to be seen as building blocks that can be reconfigured as needed to address the changing competitive landscape.
>
> This new Component Operating Model (COM) requires a 'Lego brick' approach to designing and implementing processes and the organizations that support them. Implementing a COM-based approach will have profound impacts on the structure of organizations and the nature of work.
>
> Business designs based on COM will create significant stress for traditional IT infrastructures and organizations. Our current IT services were built to serve a static—and often functionally siloed—operating model. IT needs to become much more dynamically adaptable to keep pace with the speed of business today.

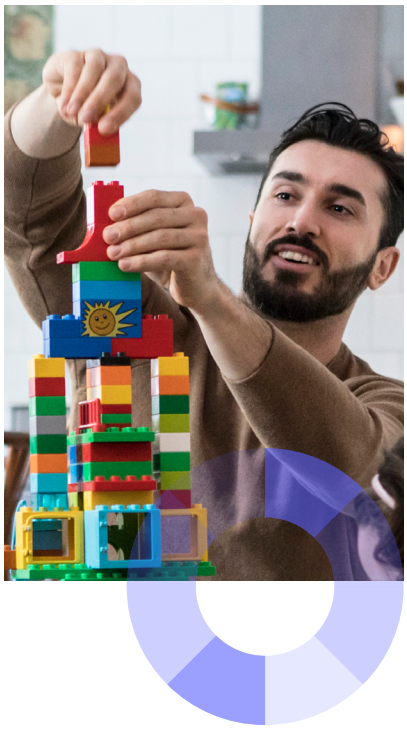

A new Component Architecture Model (CAM) approach to IT infrastructure, applications, and services will be required to ensure that IT can deliver what the business needs. The time between identifying a business need and delivering the required IT solution needs to become hours and days rather than months and years.

Murray couldn't have foreseen Covid-19 and its impact on organizations, but his predictions have certainly come to pass.

Covid-19 has only accelerated an existing shift in the way infrastructure is used and applications are built. With the rise of containers, microservices architectures, discrete modular application tools, and the like, keeping an application functioning and ensuring it functions well means juggling dozens of services, regions, geographies, service providers, and more.

Initially this move to composability was primarily seen in how applications are built—modular services, widespread use of APIs, and new infrastructure paradigms such as containers and serverless, for example. Increasingly, however, we need to think about multiple approaches to the infrastructure that these services sit upon. This will result in a far more heterogeneous cloud approach than before, with distributed clouds and multi-geography/ multi-vendor models being the norm.

So it is worth taking a minute to reflect upon the benefits that multicloud can bring.

Multicloud: The benefits of choice and flexibility

We are moving into a period when organizations will have to balance two opposing forces: on the one hand, the simplicity and control that comes from leveraging a single monolithic stack from a single vendor; on the other, the flexibility and agility that comes from accepting a more complex model and embracing composability.

At first glance, it seems counterintuitive that any organization would intentionally increase its infrastructural complexity by leveraging multiple independent clouds. Such a move could arguably be said to increase risk, result in a more complex operating model, and put additional strain on administrative functions.

While composability, a component model, widespread APIs, and better interoperability mean that we can leverage multicloud, what are the advantages that mean that we will do so?

It is worth reflecting on the advantages that a multicloud strategy can bring to those organizations that leverage one:

- → **Increased uptime**
- → **Governance, risk, and compliance**
- → **Optimizing ROI through cost control**
- → **Reducing vendor lock-in**
- → **Flexibility to compose at will**

Increased uptime

Every time a cloud vendor has an outage, the headlines scream about the risks of relying on the cloud and the need to plan for outages. Oftentimes these overly simplistic headlines do little to increase the maturity of the conversation; rather, they perpetuate fear and uncertainty among those who don't have full access to information.

As has been documented in countless articles over the years, modern organizations need to plan for failure. Planning for failure was initially considered simply through the lens of relying on more than a single region or zone for cloud infrastructure—Netflix is possibly the best-known example of this kind of single-vendor, widespread distribution.

But as has been seen on occasion over the years, relying on a single vendor, albeit one with a widely distributed infrastructure, doesn't give the highest levels of security in the event of failures. Under the unlikely, but still possible, event of an outage affecting a vendor's global infrastructure, leveraging a multi-vendor approach will help to lessen the impacts of these outages and keep an organization's applications live.

This is even more important given geographically dispersed applications that rely upon consistent data being available to every user at all times. It is important to remember that a composable model still requires consistent data across the stack—thus a multicloud approach can be the underlying factor that provides for a consistent and always-available data layer that applications can rely upon.

Governance, risk, and compliance

We live in a hugely complex compliance and regulatory environment—SOX, HIPAA, and GDPR are but three of the regulatory frameworks that organizations need to be cognizant of when delivering data to consumers.

These various regulations, and the many others like them, mean that it is unlikely that a single vendor will meet the various requirements that modern organizations must follow.

Organizations delivering critical governmental services, for example, may have to leverage a specialist cloud located in-country. Often, this will necessitate using a third-party vendor. While the large hyperscale cloud vendors are certainly increasing their footprint and rolling out in more and more regions, they are still unable to tick every single geographical box, and therefore multicloud becomes a requirement for these compliance scenarios.

In addition, from a governance perspective, the organization will quite possibly see multicloud as a way of reducing their risk profile and ameliorating any single-vendor dependency issues that they may face.

Optimizing ROI through cost control

We could dedicate an entire paper, or several, to the nuances of cloud cost optimization. Suffice it to say that there is huge complexity around the cost of cloud infrastructure, and different vendors have different strengths and weaknesses when it comes to their pricing models.

As such, savvy CFOs are realizing that cherry-picking what gets hosted with whom is a smart move. Instead of enjoying the more economically priced aspects of a particular vendor, and simply putting up with their less attractively priced ones, organizations can leverage the best and most economically priced individual services from different vendors to give themselves the most "bang for the buck."

While we have long advocated for a "value" rather than a "cost" lens through which to view cloud infrastructure, the fact is that it makes total sense to cherry-pick where possible.

Again, the benefit of a multicloud strategy is that it offers the consistency that modern applications rely upon: an always-available and always-consistent data layer that can be achieved even within the context of an optimized cost model.

Reducing vendor lock-in

Hyperscale clouds have long been considered the Hotel California of the technology world—you can check out any time, but you can never leave. While unquestionably they deliver huge value to their customers, they do so within the context of a vice-like grip.

Vendors know full well that once a customer is deeply invested in that particular vendor's platform, moving is an expensive, complex, and highly stressful proposition.

As such, there is always the risk that a vendor will revert to a "bait and switch" approach, where customers are drawn in with great pricing, only to be faced with a very different situation once they get more reliant on said platform.

It is the reality of the situation that, when it comes to ensuring vendors know that customers have options, a multicloud approach is, all things being equal, more likely to deliver increased pressure on vendors to remain economically competitive.

Flexibility to compose at will

Whether multicloud forces organizations to think in a composable manner or multicloud enables a composable approach, the fact of the matter is that by leveraging a multicloud footprint, it is more likely that the services and applications running on top of that infrastructure will demonstrate high levels of composability.

Organizations that follow a default approach of multicloud will, by definition, consider their applications from a perspective of modular building blocks. This in turn results in applications that eschew the constraints of monolithic approaches: they're more readily altered, they're more flexible, and they're more likely to meet the dynamic requirements of the future.

Delivering speed in a multicloud world

We've made a strong case for why multicloud should be the default approach toward infrastructure, so the question needs to be asked: Why do organizations avoid multicloud? For many, it comes down to concerns about speed and complexity, as well as the limitations of traditional technologies.

Applications do not sit in isolation—they use and create data, and it is the provision, capturing, and analysis of that data that often creates barriers toward a multicloud approach. Applications need to fulfill two seemingly contrary requirements: their data needs to be consistent across clouds and across geographies, and access to that data needs to display as low latency as is possible.

And this is where considering distributed and multicloud applications within the context of a data layer that delivers those traits—traits that are at first mutually exclusive—comes in.

A walk down CAP lane

This paper will inevitably be used by those aspiring to build globally-distributed applications that are hosted in a multicloud context. They will also be searching for that holy grail: applications that display performance akin to locally located ones. It is therefore worth looking at some constraints around distributed data layers.

Some 20 years ago, computer scientist Eric Brewer developed the CAP theorem, which relates to distributed applications, specifically the data those applications create and consume.

The CAP theorem, in the simplest terms, asserts that any networked, shared-data system can have only two of three desirable properties: consistency (C) equivalent to having a single up-to-date copy of the data; high availability (A) of that data (for updates); and tolerance to network partitions (P).

And given that, in those early days when a quest for speed was all-important, the CAP theorem meant that the approaches most likely to give the fastest speeds and application availability (network partitions and high availability) would also result in data inconsistency.

In the decades since the CAP theorem's introduction, however, new approaches to handling distributed systems have been developed that allow that theoretically impossible feat: data consistency, availability, and partition tolerance. The rise of new data approaches means that we can have low latency without forgoing data consistency.

While this isn't the place for a definitive technical treatise, it is important for those with the responsibility for their organizations' applications to understand the rudiments of how modern applications work.

As noted, in a world where applications are, by necessity, distributed, there will be multiple nodes included in many individual applications. Organizations have no choice but to build these applications while leveraging a data layer that empowers and enables this distributed and multicloud approach.

Enter Redis Enterprise: A unified multicloud cache and database

Just as Redis was at the cutting edge of the global shift toward real-time applications, Redis Enterprise is leading the way as a multicloud data layer. Redis Enterprise is the only database that can deliver the flexibility, reduced risk, and ROI of multicloud without sacrificing speed or data consistency.

Redis delivers unprecedented speed by storing data in lightning-fast DRAM to provide sub-millisecond speed at any scale. Redis Enterprise's Active-Active Geo Distribution provides the ability to read and write to multiple nodes simultaneously, enabling Redis to distribute the same dataset with consistency across geographies, deployment environments, and clouds.

Because of its ability to deliver on the promise of multicloud, Redis Enterprise has been widely adopted by companies of all sizes. Redis Enterprise is commonly used as an enterprise-grade cache to add speed at scale while reducing load on primary databases. Redis is also widely used as a primary database to power modern real-time applications. As a primary database, Redis Enterprise provides unprecedented flexibility with support for a variety of popular data models.

Redis Enterprise can be deployed on premises or in AWS, Azure, and Google Cloud.

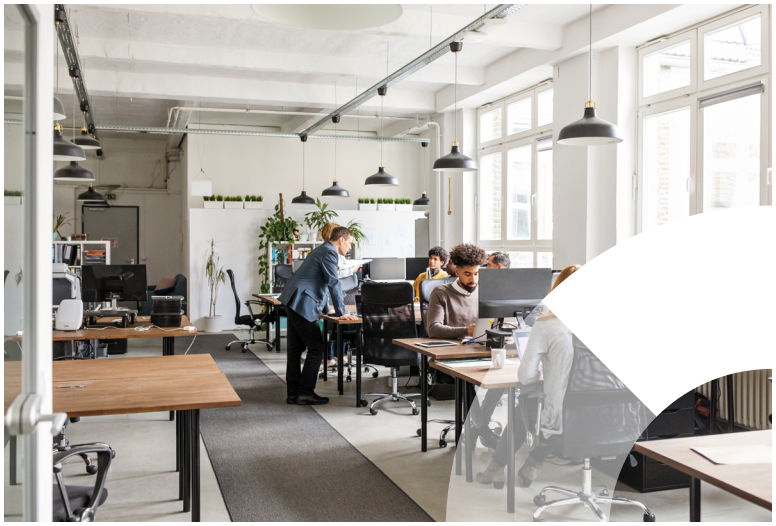

Summary

The last decade has seen a huge shift, and that shift has been greatly accelerated by the Covid-19 pandemic. "Digital-first" is no longer a luxury; rather, it is a core necessity and this approach to digital-first has cemented in place a "multi-everything" approach toward technology. Multi-service, multicloud, and multi-geography have all become the new normal.

But this new normal introduces some challenges to those working within its constraints: how to ensure speed and distribution, security, compliance, stability, and flexibility—not to mention the best price and best performance.

We believe that an approach that leverages multicloud—and building blocks that enable the use of multicloud—is this new normal, and a way of operating that all organizations need to understand and leverage.

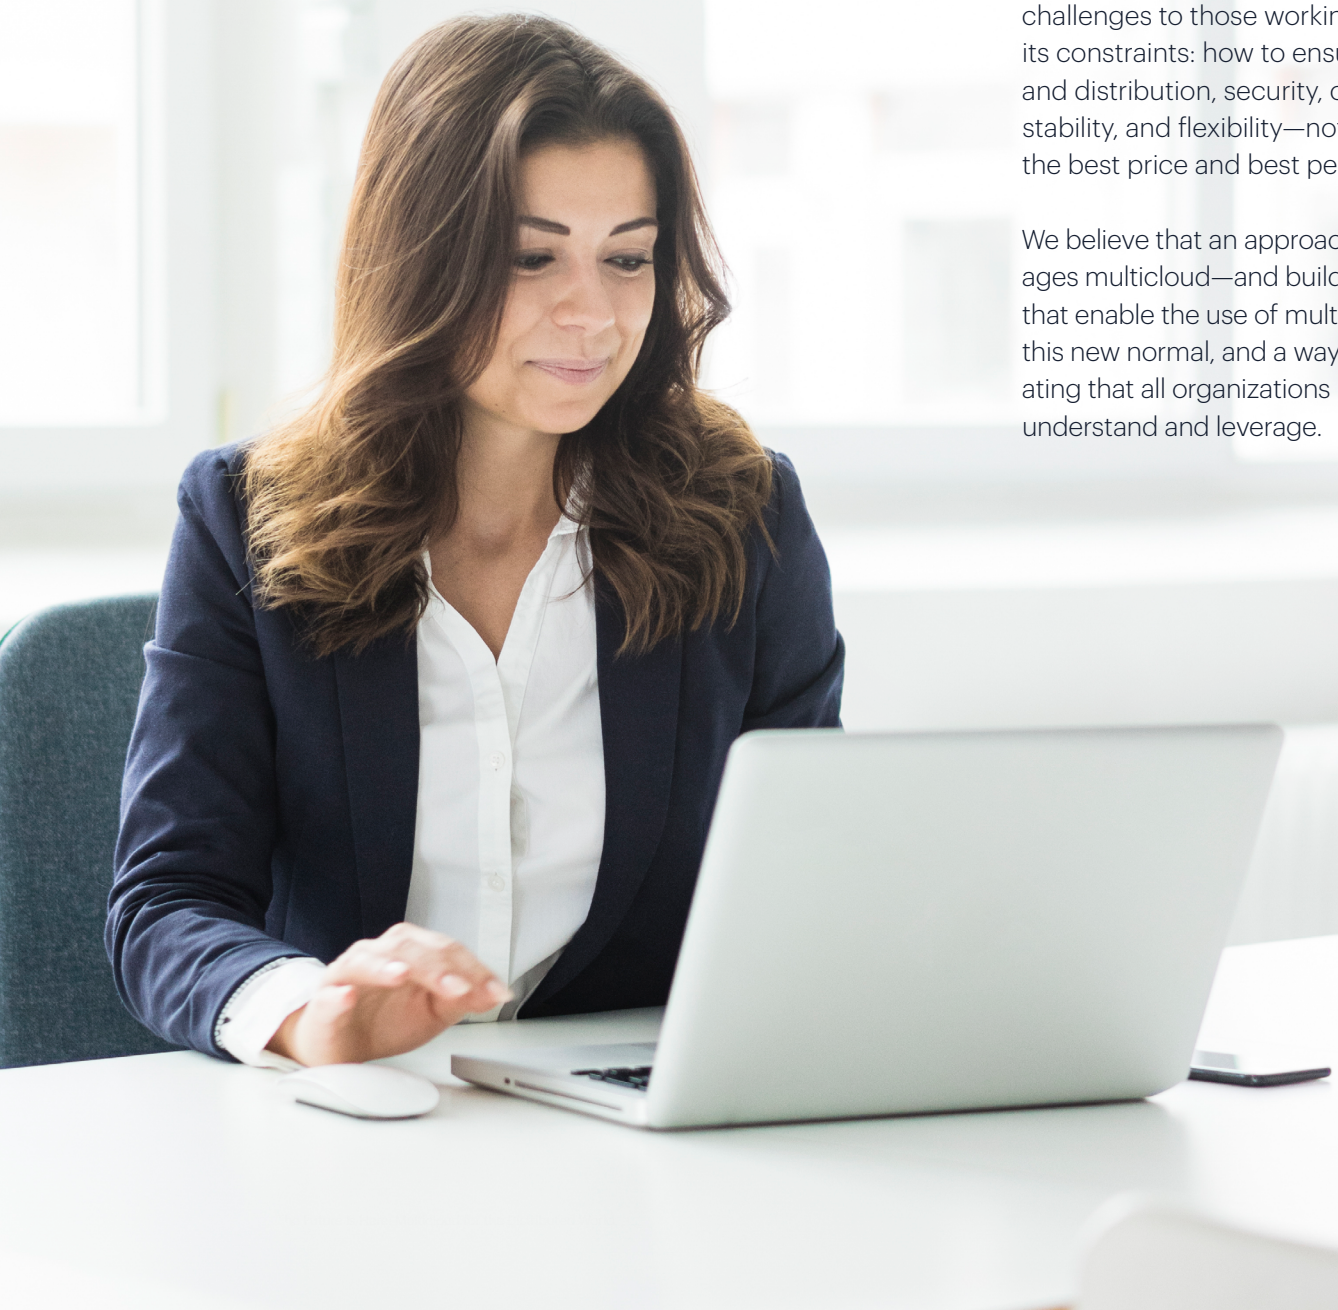

Learn more

Want to further explore how Redis Enterprise can deliver on the promise of multicloud? Check out the resources below to learn more about the value of Redis Enterprise as a multicloud cache and database.

Enterprise Caching: Strategies for Caching at Scale
E-book

Business today requires applications that are fast, always available, consistent, and reliable. Your users expect it. Your business depends on it.

And just as businesses grow, so do data size and cache complexity. Basic caches just don't cut it at enterprise scale. An enterprise-grade cache needs to bring more than speed. It needs to provide flawless performance at any scale, anywhere, at all times.

Read the E-book to learn things like:
- The most effective methods for scaling Redis
- How to overcome increased complexity to maintain cache consistency
- How to cache effectively across multiple clouds and on-prem
- And how you can do all of the above – with Redis Enterprise

Download now

Strategic Data Flexibility
White Paper

Databases used to be just for storing and retrieving static data. As data sets and use cases grew, this basic function became more complex, leading cloud service providers to advocate for purpose-built databases for each use case or functionality set. However, this approach to database models adds overhead and complexity.

While purpose-built databases meet a certain need, they fall short when you need to deliver high-performance, robust applications that meet the real-time requirements of today's customers. Redis Enterprise presents a modern approach to data strategy, providing a multi-model database that adapts to your business needs rather than adapting your business needs to the technology.

Read to learn:
- The pros and cons of the various approaches to data strategy
- How Redis is architected to provide you with ultimate data flexibility
- Why Redis Enterprise is the best choice to power your modern applications

Download now

Top 5 Benefits of Adopting a Cloud Migration-Proof Data Layer
E-book

As the global cloud adoption wave maintains its momentum, organizations continue to turn to the cloud for its promise of scalability, performance, cost-savings, and rapid innovation. In order to truly realize this value, the data layer needs to navigate migrations and hybrid or multi-cloud deployments fraught with challenges.

Download this Redis e-book to learn how a migration-proof data layer can ease your transition to the cloud while unifying your data—and extend the benefits of cloud to hybrid and multi-cloud environments.

Top reasons to download:
- Identify key risks associated with migrating your data to the cloud
- Uncover key strategies to help prevent data migration challenges as your organization moves to the cloud
- Learn how a migration proof data layer can help realize the true value that the cloud promises

Download now

About Ben

Ben Kepes is a technology analyst, commentator, and consultant. Over the past decade and a half, he has built up a significant following as a globally recognized subject matter expert in the areas of cloud computing, enterprise technology, and digital transformation. Ben's commentary has been widely published in such outlets as *Forbes*, *Wired*, and *The Guardian*, and he has been invited to speak at a wide range of technology, business, and general interest conferences.

About Redis

Businesses are more digital today than ever before. They need to build, deploy and run real-time services in order to stay ahead of the curve. The notion of real-time is not just a nice-to-have anymore. It's an expectation. It is what sets a merely good user experience apart from a great one. A real-time data layer is a critical enabler in creating those real-time experiences.

Redis makes apps faster, by creating a data layer for a real-time world. We are the driving force behind Redis Open Source, the world's most loved in-memory database, and commercial provider of Redis Enterprise, a real-time data platform.

Redis Enterprise powers real-time services for over 8,500+ organizations globally. It builds upon the unmatched simplicity and speed of Open-Source Redis along with an enterprise-grade data platform that offers robustness of modern data models, management, automation, performance, and resiliency to deploy and run modern applications at any scale from anywhere on the planet.