kyndryl

Future of Compute For Financial Services

Future of Compute

Authors:
Jim Freeman – Chief Technology Officer (CTO), Kyndryl, ANZ
Chirag Arora – Head of Strategic Solutions, Kyndryl, ANZ

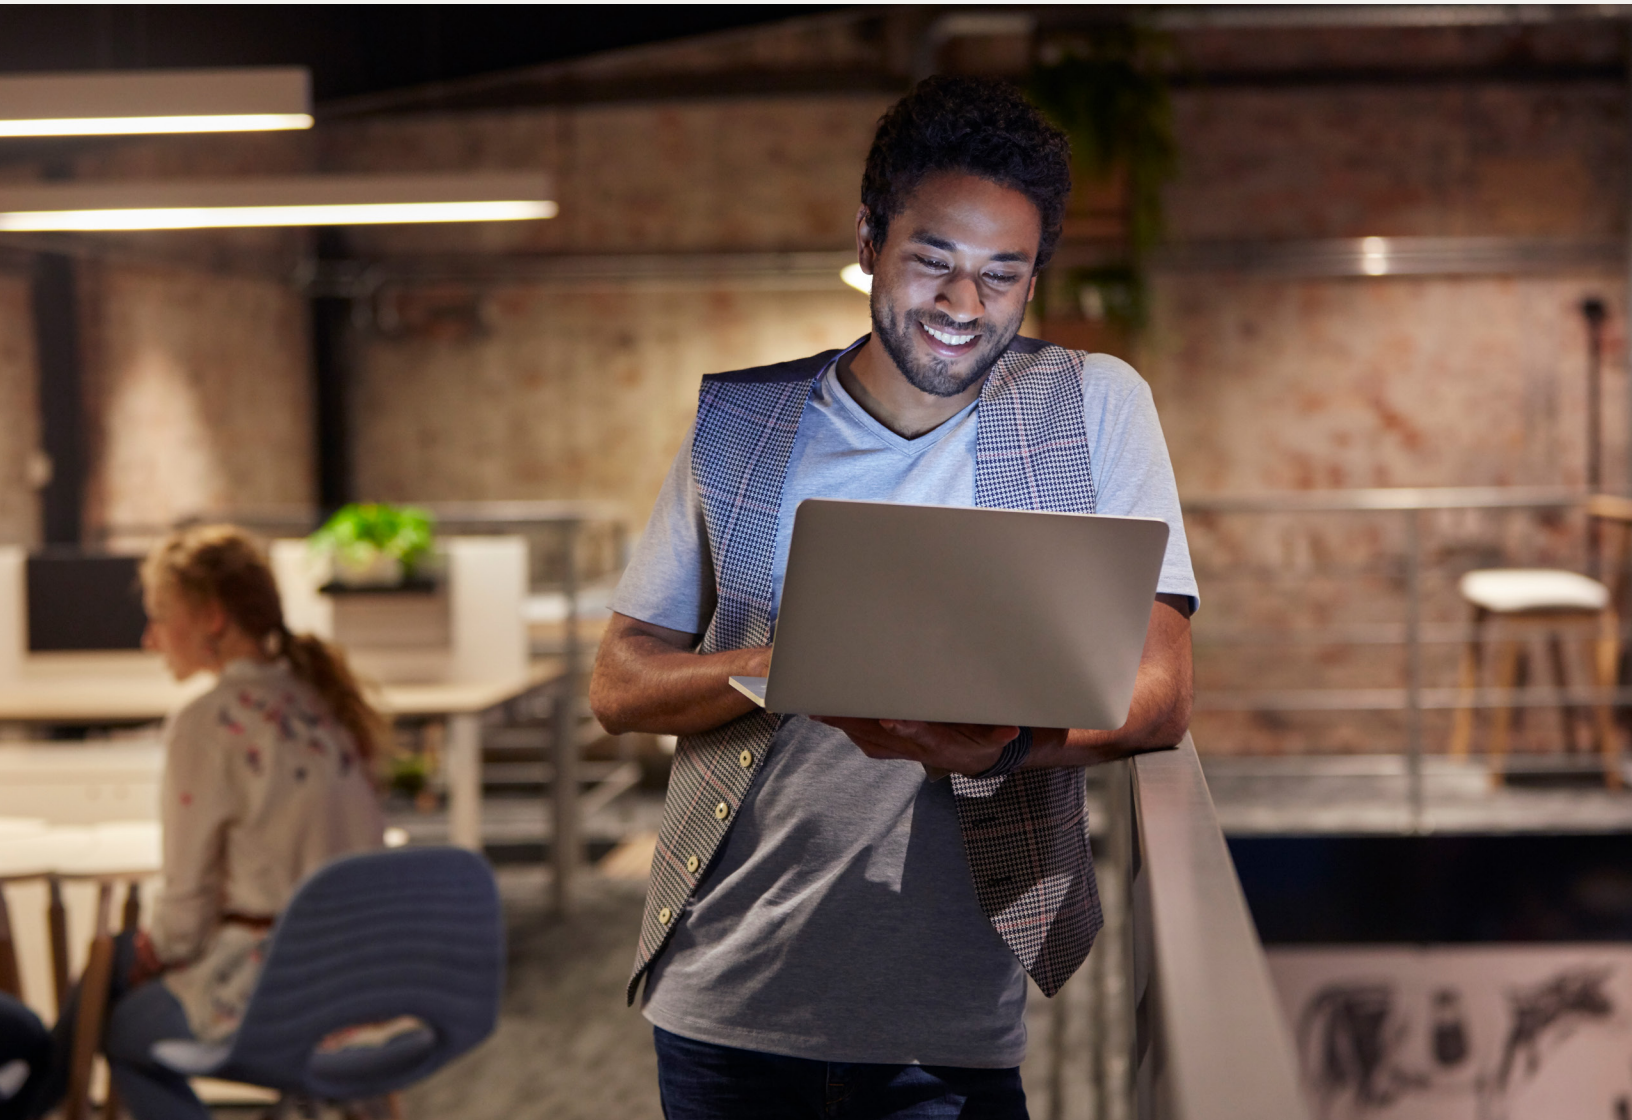

Contents

Summary

Applications have become integral to our daily lives, influencing how we work, communicate, learn, entertain, and even relax. These applications are the innovators of functions in the digital realm, embodying business logic, user interfaces, and data processing. While applications are innovators, computing infrastructure is the silent backbone supporting them. It includes servers, virtual machines, containers, and cloud resources. Just as a building's foundation remains hidden but crucial, compute infrastructure ensures applications run seamlessly.

The advancement in technology is driving enhanced and rich applications, which is inspiring many questions about the silent backbone of compute infrastructure—from "Is compute a commodity or enabler?" to "Is Cloud the answer to digital transformation?". But one in particular is complicated: "What is the future of compute?" If you go traditional on-premises, you may lose out on innovation and modernisation. However, going entirely on Cloud could increase your total ownership cost and risk profile.

This article provides practical strategies, considerations for approaching this question, and reasons to use it as a guide. It will provide a pathway to symbiosis: a harmonious dance between applications and compute infrastructure in the right environment. The article also provides perspective on what the future of compute would mean for a bank.

The silent backbone of computing infrastructure has been evolving recently due to various factors such as Artificial Intelligence and Machine Learning, Cybersecurity, High-performance computing, serverless computing, Edge computing, Quantum computing, and Containers. Most of this evolution has been possible because of cloud-related advanced technologies and their adoption, which offers several benefits to enterprises.

Cloud benefits for large enterprises include:

1. Flexibility, standardisation, and simplification of IT resources, as well as a marketplace for cloud-native software services that enable a composable enterprise.

2. Accelerating in realising transformation to a digital estate. Cloud services enable enterprises to develop and adopt digital assets across their value chain using automation, orchestration, and integration of services. Cloud services also provide data infrastructure that supports real-time access, scalability, and consolidation.

3. Non-disruptive innovation with insights. Cloud services also provide multiple technology choices and operating models that align with the enterprise's digital transformation priorities. Cloud also facilitates cross-industry collaboration and innovation, as well as data analytics and insights that fuel customer experience and competitive advantage.

4. Security, resilience, and compliance. Cloud providers ensure physical, information, and network security and compliance for their infrastructure and services. The cloud also offers resiliency and availability options, as well as identity, access, and control systems. Enterprises can inherit these security features and certifications by adopting cloud services.

5. Commercial efficiency: Cloud can reduce the capital and operational costs of IT for enterprises by eliminating the need for hardware, software, and data center facilities. Cloud also enables consumption-based pricing, real-time financial management, and cost savings for development, testing, and disaster recovery environments.

6. Environment and sustainability: Cloud improves sustainability and reduces IT carbon emissions by using renewable energy resources, power-efficient hardware, and surplus capacity and waste. Cloud also provides reporting and metrics on environmental and sustainability impacts for enterprises.

While Enterprise and Government organisations can experience numerous benefits of adopting the Cloud, it's essential to recognise that there may be a flow-on effect and associated risks related to a Cloud "only" approach for a complete ICT environment.

Key Findings

As organisations consider Cloud technologies and Cloud Service Providers (CSP) service models, they should be informed of commercial, technology, financial and operational risks and hidden costs that diminish the realisation of the Cloud benefits. The following highlights some of the unintended risks:

1. **Culture and operating model issues:** Adoption of cloud services requires a mindset shift for the enterprise, as well as how it creates a "two-speed operating model" between the cloud and on-premises environments. Enterprises often experience a corresponding loss of control and visibility when they use cloud services.

2. **Security and compliance concerns:** Cloud services increase the potential for data breaches, regulatory violations, and shared responsibility models that cloud adoption entails. Furthermore, enterprises may face reduced control over critical application availability and disaster recovery in cloud environments without careful landing zone designs.

3. **Economics of cloud:** Cloud adoption often results in unexpected costs, underutilized assets, and complex licensing models. Increasingly, on-premises operational expenditure offerings compete with public cloud providers by delivering lower unit costs.

4. **Environment, social and governance impact:** Cloud computing can contribute to global CO2 emissions and environmental impact. The effectiveness of cloud providers' renewable energy certificates in offsetting their carbon footprint continues to be debated.

5. **Data governance:** Any cloud adoption strategy should include adopting FinOps practices, DevSecOps culture, federated development processes, and well-architected frameworks. The risk of duplicate data and poor data hygiene is compounded by adopting cloud services. Hence, the importance of data integration and governance across multiple cloud environments is higher than in traditional estates.

The pace of change and technological advancements will continue to drive demand for innovation in computing infrastructure, further driving different business models and applications.

For example, Edge computing is a transformative approach that brings data processing, analysis, and storage closer to the location where it's needed, improving response times and security. Quantum computing represents a significant leap in processing power and opens exciting possibilities for scientific advancements and practical applications.

The combination of risks and benefits makes the hybrid cloud a compelling choice for enterprises looking to optimise their IT infrastructure for the present while preparing for future technological advancements. As technologies evolve and new business models develop, the future of computing may change, but in the foreseeable future, the hybrid cloud is our recommended approach to enterprises.

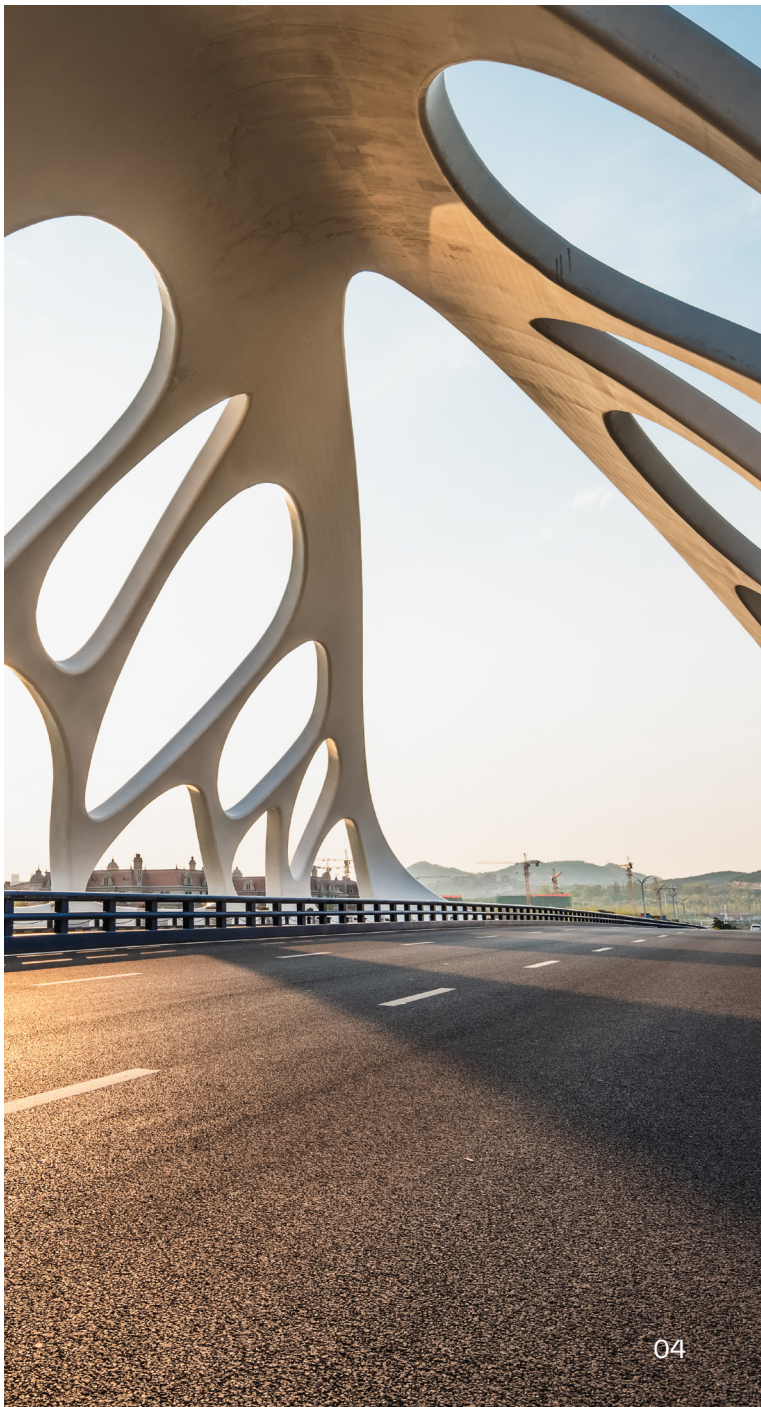

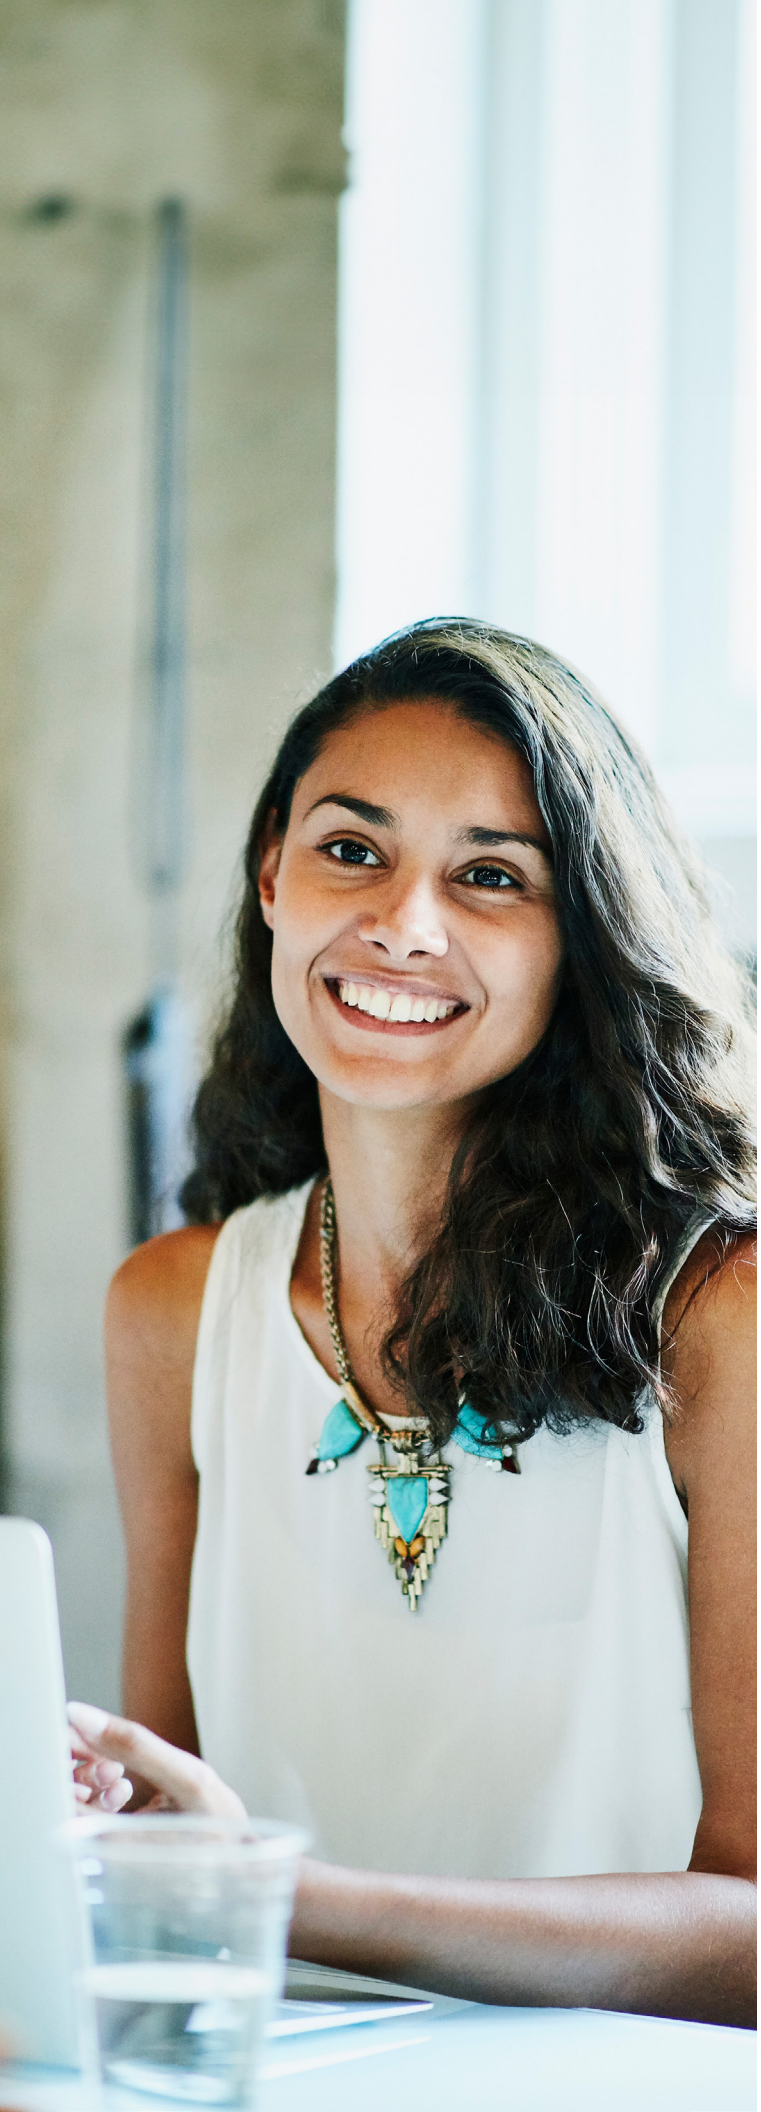

Why Lean into the Disruption?

As early as 1999, Salesforce became a popular example of successfully using cloud computing Software as a Service (SaaS) model. This SaaS model of Salesforce provided Customer Relationship Management (CRM) function as an outcome to enterprises while eliminating software installations, updates, security patching, and hardware and infrastructure maintenance. In 2007, companies such as AirBnB demonstrated that applications built using cloud technologies can support enterprise-grade businesses. The transformation in consumer payments services, particularly in markets like India and China, has demonstrated the potential of cloud-based digital financial products. By contrast, Federal Reserve in the United States points to the high cost of wire transfers in traditional banking. This has clear implications for core system modernisation.

The strategic planning and design of mission-critical applications on cloud computing platforms can lead to significantly faster delivery times, all while maintaining confidentiality and enhancing resiliency and performance. Moreover, the use of standardized, commodity cloud infrastructure has made it incredibly easy for application programmers to provision and manage their infrastructure, thanks to the highly automated, self-service catalogues and APIs built by vendors.

However, enterprises were still heavily invested in traditional, labour-based infrastructure management paradigms, which, by comparison, became even more of an impediment to application programmers. Furthermore, all their applications were designed to run on traditional enterprise infrastructures, and information security and risk management teams were attempting to apply legacy controls to new infrastructure.

The COVID-19 pandemic has accelerated the need for digital transformation, pushing many enterprises to leapfrog their strategies by 5 years. In the absence of a comprehensive digital transformation plan and with reduced response time, many enterprises initiated this transformation by modernizing their infrastructure, either by migrating applications to the cloud or adopting SaaS offerings from cloud providers. The cloud not only promised to address issues such as technical debt and capital expenditure but also provided access to innovative technologies, opening up new opportunities. These innovative technologies have significantly reduced barriers to creating new business models and have made the execution of new ideas born in the cloud much easier.

Over the last 5-10 years, the world has seen successful businesses born in the cloud, with the purpose of addressing problems and creating new opportunities to improve the way we live, work, play or collaborate.

What value did adoption of born-in-the-cloud practices promise for the Enterprise?

Traditionally, an enterprise's Information Technology (IT) department has been following business demands in a reactive mode, becoming a bottleneck for the enterprise to enable business growth. Maintaining a balance between predicting growth, spare infrastructure capacity, supply chain duration, data centre space and capital expenditure is always a challenge. In addition, many legacy applications are monolithic and designed for vertical scaling. Easier access to the technology stack, with many open-source capabilities, has provided entrepreneurs with flexibility and agility to test their ideas and succeed or fail fast.

Flexibility with Simplification enabling Composable Enterprise

Cloud promises to solve many common challenges across large enterprises, irrespective of their industry vertical. The scale of the data centers in cloud dwarf enterprise data centres, which means that resources such as floor space, power, storage, and compute are virtually unbounded in the cloud. This unlimited capacity to consume flexibly on demand at scale, removed capex, supply chain, and spare capacity concerns while serving unpredictable business growth.

In addition to basic technology components of computing, storage and network infrastructure, the cloud platform fosters a sense of community and collaboration. It provides standard processes, operating models, and tools, creating a common ground for all users. Collaboration is materially improved when standard processes, vocabulary, and roles are in place. The wholesale adoption of standard cloud tooling, configuration, and shared responsibility models has lifted collaboration to a new level, making the audience feel more connected and part of a larger community.

The cloud's user-friendly interface and simplified provisioning and management of cloud services further enhance flexibility and standardisation. Standardising components and managing Application Programmable Interfaces (APIs) makes the technology more accessible and less daunting. These APIs permit application developers to use automated tooling to support their application development life cycle, wherein those tools dynamically configure and automatically deploy the infrastructure the developers need, making the audience feel more at ease and less overwhelmed by the technology.

On top of its offerings, the cloud provides a marketplace that hosts cloud-native software services built on its platform by third-party developers. These software services are integrated with the cloud platform's tools and processes, further enhancing standardisation for enterprises leveraging these services. A robust marketplace for software services means that application programmers can consume and re-use micro-services for standard capabilities. They are now able to build entire solution using reusable micro-services to provide integrated digital experience and enable a composable enterprise. One notable example is Stripe which offers pay-as-you go credit card processing including assumption of all regulatory compliance and liability for protection of PCI Data Security Standard.

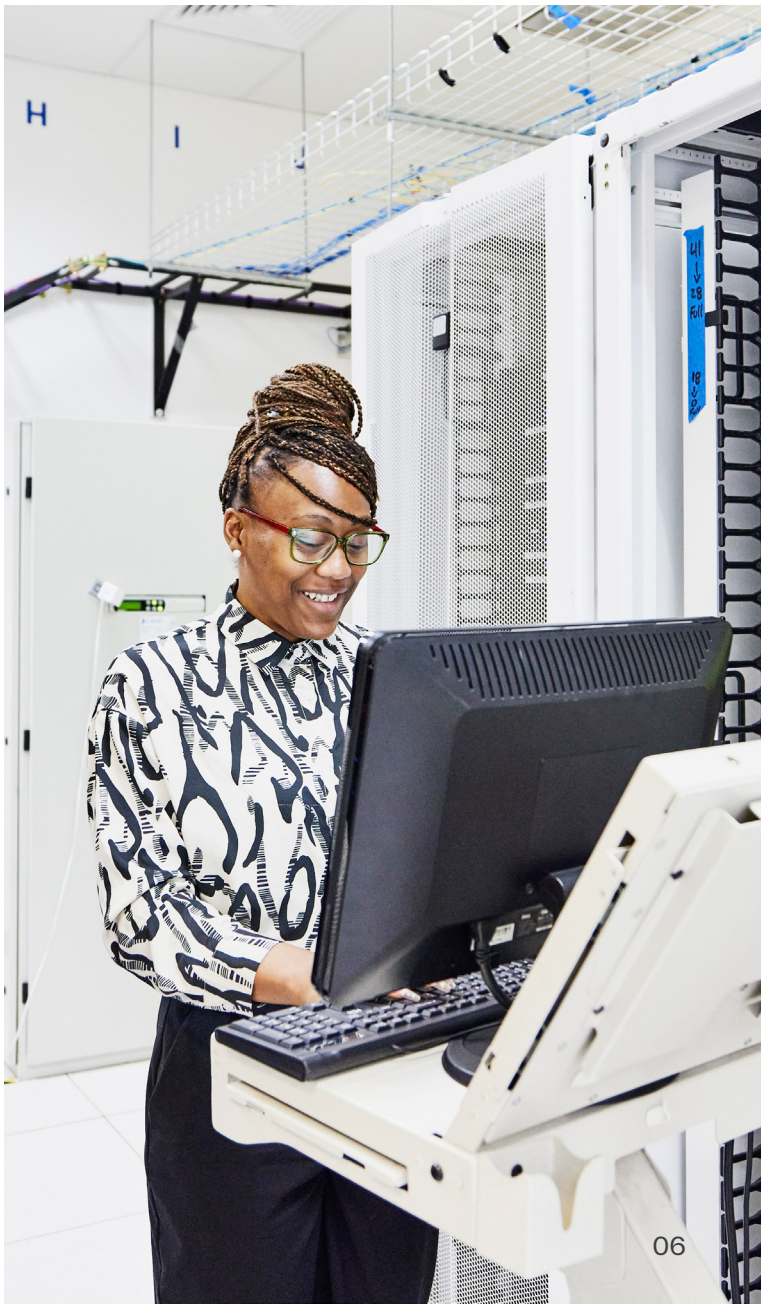

Accelerating Digital State

The ease and speed of access to cloud services have been accelerating the development and adoption of digital assets for an enterprise across its entire value chain. The functions across human resources, travel, IT operations, finance, marketing, sales, delivery, or customer service business units use cloud-based digital SaaS offerings such as Workday, SAP Concur, Salesforce, and ServiceNow.

Cloud services are available at any time, from anywhere, using any device, without long ordering and installation times. Moving away from manual effort and legacy technologies, mindsets, and processes to composable services and automation of application life cycle tooling, such as Infrastructure as Code (IaC) and Continuous Integration / Continuous Delivery (CI/CD), has increased agility and speed for application development.

As infrastructure configuration and management are achieved through automated catalogue fulfilment or API calls, the applications teams have the requisite skills to perform all application development and maintenance activities on cloud platform – including infrastructure. By reducing the time to value for the business applications combined with underlying integration with management and monitoring tools, the cloud has reduced friction between IT development and operations teams, enabling agile and DevOps practices. These DevOps practices continue to identify patterns of activities that require manual effort and use cloud services to digitise these activities using code.

Data is an enabler for digital transformation, providing context and meaning to the information presented to customers. As a result, the ability to access and process data is critical as a foundation for applications in a digital state. Cloud provides infrastructure for data that allows real-time and flexible access and scalable performance while enabling data consolidation for an enterprise in a single cloud, eliminating data silo problems.

Non-Disruptive Innovation with Insights

Cloud showed a pathway to achieving digital innovation for customers with different technology enablers available on the same platform. The myriad of services available on the cloud builds a vision for many enterprises to leverage technology for new products and services and new business models to create better customer experiences and competitive advantage.

Through its native services and third-party applications in the marketplace, as well as different service models of IaaS, PaaS and SaaS, the cloud provides multiple technology choices and operating models. Enterprises can align their digital transformation priorities, as per their business strategy, with the right operating and service models and technology services.

As enterprises from adjacent industry segments were present on the same platform, cloud reduces the barrier to integrate their services and to innovate with much higher capital efficiency. This digital innovation introduces genuine transformation for servicing customers, gaining insights, and enriching experiences. By building digital engagement and consumption channels for customers, enterprises can also gain insights into their customers' behaviour, choices, and decisions, through the data collected from these digital channels. Using cloud services such as data lakes and data warehouse, enterprises collect data at scale from multiple engagement points, perform analytics and apply insights, fuelling innovation.

When it comes to data in the cloud, security and risk have been key considerations for enterprises, especially given the rising cyber security threat environment.

Security, Resilience and Compliance

Confidentiality, Integrity, and Availability are fundamental tenets of security that enterprises consider, and cloud providers have incorporated these principles into the cloud. Cloud service providers provide and maintain physical, information, and network security and compliance for global, regional, and availability zones' infrastructure that hosts cloud services. This includes authorised access management for physical access in the data centers.

Resilience is an essential aspect of data security that every enterprise aspires for their services. Cloud provides resiliency, by default, for their physical infrastructure, with options to design resiliency using other cloud services. If allowed, enterprise data is stored redundantly on multiple devices across geographically separated facilities to achieve resiliency outcomes. Due to geographical presence and scale using many data centers, the cloud has a big advantage in providing service availability to enterprises.

Based on the industry, enterprises are required to comply with regulatory standards for their business services in accordance with local, national, and international laws. Compliance may also apply to cloud usage, including data residency and privacy. Cloud providers have many security standards and compliance certifications built into the platform and environment due to the design, processes and controls applied. Enterprises have the advantage of inheriting these services by adopting cloud services.

The cloud's security advantage is also evident in the different identity, access, and control systems deployed in a cloud environment. These control systems, such as automated Distributed Denial of Service (DDoS) attack protection, multi-factor authentication services, key vaults, conditional access policies and role-based access control, are available on demand and integrated into the environment.

Commercial Efficiency

One of the challenges with the traditional IT environment in large enterprises has been the capital investment required to procure, install, design, and configure hardware and software assets. In addition, many enterprises have capex or opex costs for datacenter facilities and maintenance of hardware and software assets. With infrastructure being available in the cloud, enterprises don't need to procure, deploy, and maintain hardware and software assets. Datacenter facilities, removing the need for capital expenditure while reducing the cost of labour for initial deployment or refresh of hardware as part of lifecycle management. The labour cost is further reduced by eliminating the need to maintain the currency of hardware and software assets.

For an on-premises IT environment, enterprises purchase enough surplus capacity to cover annual peaks in workload (tax time, holiday shopping, or other industry-specific business cycles) and factor growth for the next few years. This surplus capacity results in underutilisation, higher capex and opex costs, and waste of resources such as power and space. Cloud removes these inefficiencies of surplus capacity by providing resources on-demand and reduces cost with a consumption-based pricing model, where enterprises pay for what is used.

Further, the cloud reduces technical debt for datacenter facilities, hardware, and related maintenance, including currency, while being more cost-effective. In addition, the cloud enables real-time financial management by empowering finance teams and supporting CIO organizations to provide real-time feedback on burn rate, removing retroactive cost management at the end of month and quarter.

By providing automation, orchestration and integration of services, the cloud removes manual effort and tasks required to configure, integrate, test, monitor and manage the environment. As a result, enterprises can save on the professional effort required to perform these activities. With different service models of IaaS, SaaS and PaaS and related shared responsibility models, enterprises further reduce costs related to skills required to perform traditional on-premises environment-related tasks.

Enterprises maintain additional environments for development, testing, non-production, innovation lab, and disaster recovery. These environments are required for different purposes, such as simulating production environments, fixing software bugs, validating the quality of services, system integration testing, user testing, piloting new solutions and disaster recovery to maintain business continuity. All these environments incur hardware, facilities, software, network, servers, storage, and associated labour costs. By using the cloud for these environments, enterprises can save associated costs for that environment.

Environment and Sustainability

Technology-based solutions are a key driver in improving sustainability and reducing carbon emissions in the IT environment. Data Centres are responsible for 2% of the world's carbon footprint, and this number is steadily growing.

According to 451 Research, US data centers consumed about 268 terawatt-hours (TWh) of energy in 2019. By adopting the cloud and migrating to an on-premises environment, enterprises reduce energy consumption and carbon emissions related to data centers, servers, storage, and other related infrastructure.

Due to the on-demand availability of resources and automated scaling solutions, the cloud removes the need to maintain surplus capacity, resulting in improved sustainability. Access to innovative technologies combined with data intelligence drives resource management and waste reduction solutions.

Cloud providers and their suppliers are utilising renewable energy resources, where possible, for data centers and related construction materials. In addition, cloud providers are investing in more power-efficient chips and hardware to reduce energy consumption and related carbon emissions. Many cloud providers, such as Amazon, Microsoft and Google, have also set a carbon-free goal for their operations in the next 3 to 8 years.

As enterprises understand environmental and sustainability risks, recording and reporting on them, including defining metrics, is becoming critical. Many cloud providers enable enterprises to report on their environment and sustainability metrics by providing carbon usage data. For example, Microsoft Cloud for Sustainability helps enterprises measure, understand, and take charge of their carbon emissions, set sustainability goals and take measurable action.

Depending on strategic, operational, technical, and economic factors, enterprises have varying expectations of the benefits and value attained from embracing the cloud. The above list is not exhaustive of all the values the cloud promises but identifies the most common values that an enterprise with legacy applications and an IT environment expects.

The state of realisation

Rapid development of cloud and related technologies has pushed technology industry to new heights, reduced the technology barrier for many start-ups and accelerated business continuity during pandemic. Based on PwC US Cloud Business Survey, despite the promises of cloud, 53% of the companies are yet to realise substantial value from their cloud investments. The value realisation gap of 53% is dependent upon the business outcomes that enterprises, in the survey, were expecting to achieve through cloud. This section will assess the state of realisation of promises made by cloud, as mentioned in previous section for enterprises with legacy applications and IT environment.

Flexibility with Simplification enabling Composable Enterprise

Enterprises have a challenge of changing infrastructure and application resources dynamically, with business demand, in their on-premises IT environment, without any capital expenditure. Cloud has addressed that challenge based on consumption model and automation tools to scale resources up/down based on usage. Bi-modal change management allows for rapid deployment of changes in the cloud, while continuing to manage legacy environments, creating a holistic governance model that can be audited for compliance.

Reduction of friction in application development dividends continue to grow, but the pace of improvement is being slowed by the increasing dependency on traditional, monolithic applications needing to be modernised. The business case for modernising many of these monolithic applications is not viable from cost, benefits, and risk analysis. While data virtualization and API encapsulation are starting to shrink the perimeter of the monoliths, enterprises maintain non-native technology stack on cloud same or similar as on-premises to ease migration, improve interoperability, and management. This limits complete value realisation of flexibility and simplification by not benefiting from cloud native services.

Through standardisation, enterprises have improved collaboration, increasing velocity of application development; team members can quickly understand and, therefore, quickly augment functionality in sprints with little learning curve from sprint to sprint or as they move from squad to squad. Cross-business development is enabled through microservices development and governance of API gateway services allowing for rapid development, sharing, and updating code for reusability across business units.

Accelerating Digital State

Enterprises have realised selective benefits of accelerating digital transformation by adopting cloud, especially during pandemic phase. Many of these digital transformations involved migration to SaaS platforms to provide like-for-like functions as on-premises software, but on consumption basis. Similarly, many applications were moved to cloud using lift-n-shift model to improve access from distributed cloud points-of-presence rather than enterprise centralised data centres. However, due to unplanned emergency response during pandemic, the end-to-end target digital state wasn't aligned to business strategy for many enterprises, resulting in limited benefit.

In IT Operations, Infrastructure as code (IaC) continues to deliver value off premises but, as vendors enable configuration and management APIs for infrastructure, enterprises have realised IaC benefits with their on-premises estates as well. However, the lack of on-premises availability of PaaS services especially around database, data pipelines and ML models has left the old and slow proprietary software stack as the barrier to accelerating digital state of enterprises. In addition, many applications need to be re-written to become cloud native, to accelerate end-to-end value chain.

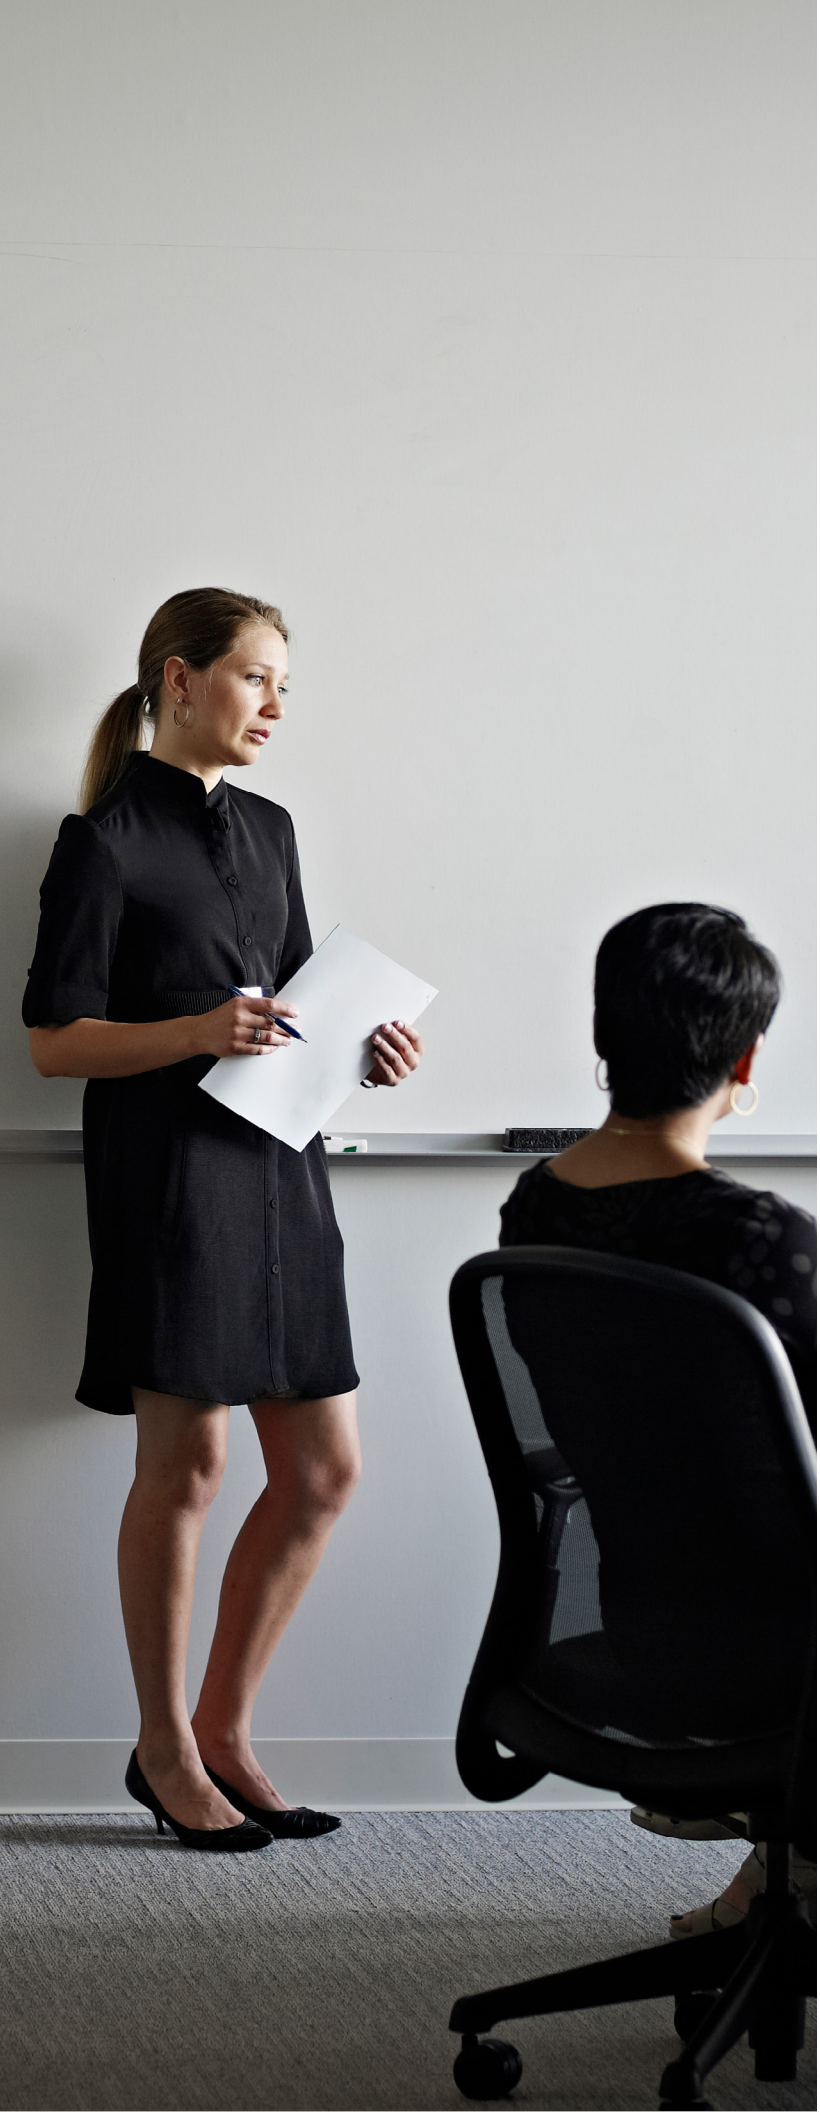

Data is key enabler for transformation, and enterprises have realised benefits of cloud in standardisation, consolidation, and integration of data on the same cloud platform. However, the multi-cloud adoption and increase in SaaS applications by enterprises is resulting into fragmentation of data, creating new data silos in cloud environments. While data silos have existed in the on-premises environment, ease of cloud consumption along with many SaaS applications is creating cloud silos resulting into information bottlenecks.

Non-Disruptive Innovation with Insights

Different cloud technologies combined with flexibility, scale, and speed to continuously plan, code, build, test, integrate, deliver, deploy, operate and monitor has accelerated innovative solutions and new business models. Many industries continue to benefit from innovation technologies available through cloud platforms. As per 451 Research Cloud Price Index, the big 3 Hyperscalers Google, Amazon, and Microsoft have added 300,000 net new services in Q2 2022 to their Google Cloud Platform (GCP), Amazon Web Services (AWS), and Azure cloud platforms.

Non-disruptive innovation continues in operational and delivery methods, based on different low code/no code technologies, and service and shared responsibility models available from cloud providers. Enterprises continue to upskill and redeploy in-house staff to higher business value services as they migrate on-premises services to PaaS and SaaS cloud offerings. This innovation further extends into new financial consumption models by many enterprises for their services as key underlying IT cost is not a fixed cost.

Enterprises are utilising multiple cloud-based solutions, such as API management, data governance and service mesh, to drive incremental innovation, and co-create offerings with customers and partners. Cross industry innovation using data driven insight is increasing especially between government and private sector. For example, Australian mining sector has advanced technology solutions using data and analytics, robotics and automation, edge computing, and off-grid and energy storage technologies that can be applied to defence and agriculture industries.

Innovation benefits are being realised by many enterprises especially where a parallel digital native of an enterprise is created to build new digital offerings. The innovation benefits are diluted where new digital product is dependent upon legacy applications, systems, and data to extract value.

Security, Resilience and Compliance

Traditional applications residing on on-premises infrastructure assume that the on-premises estate will protect them from vulnerabilities. Enterprises are now learning from their Chief Information Security Officers (CISOs) and security organizations that they should assume bad actors have already infiltrated on-premises environments, and they must shift their thinking

about the ‘safety’ of on-prem vs cloud. Over the last 4 years, native hyperscaler offerings in threat assessment, vulnerability monitoring, security in DevSecOps and remediation automation have improved the security posture of applications on cloud. Enterprises have also revised their standards and controls for being more cloud aware and more security functions are being delivered as code.

Enterprises are realising benefits of large geographical presence of cloud, combined with automation technologies, for improving disaster recovery and resiliency capabilities of their IT environment, while reducing fixed cost of data centre and other infrastructure. Resiliency capabilities are built-in the PaaS and SaaS cloud applications, improving security and reliability posture for many enterprises IT environment. However, where enterprises have legacy applications and hybrid environments, the resiliency benefits of cloud are not being completely realised.

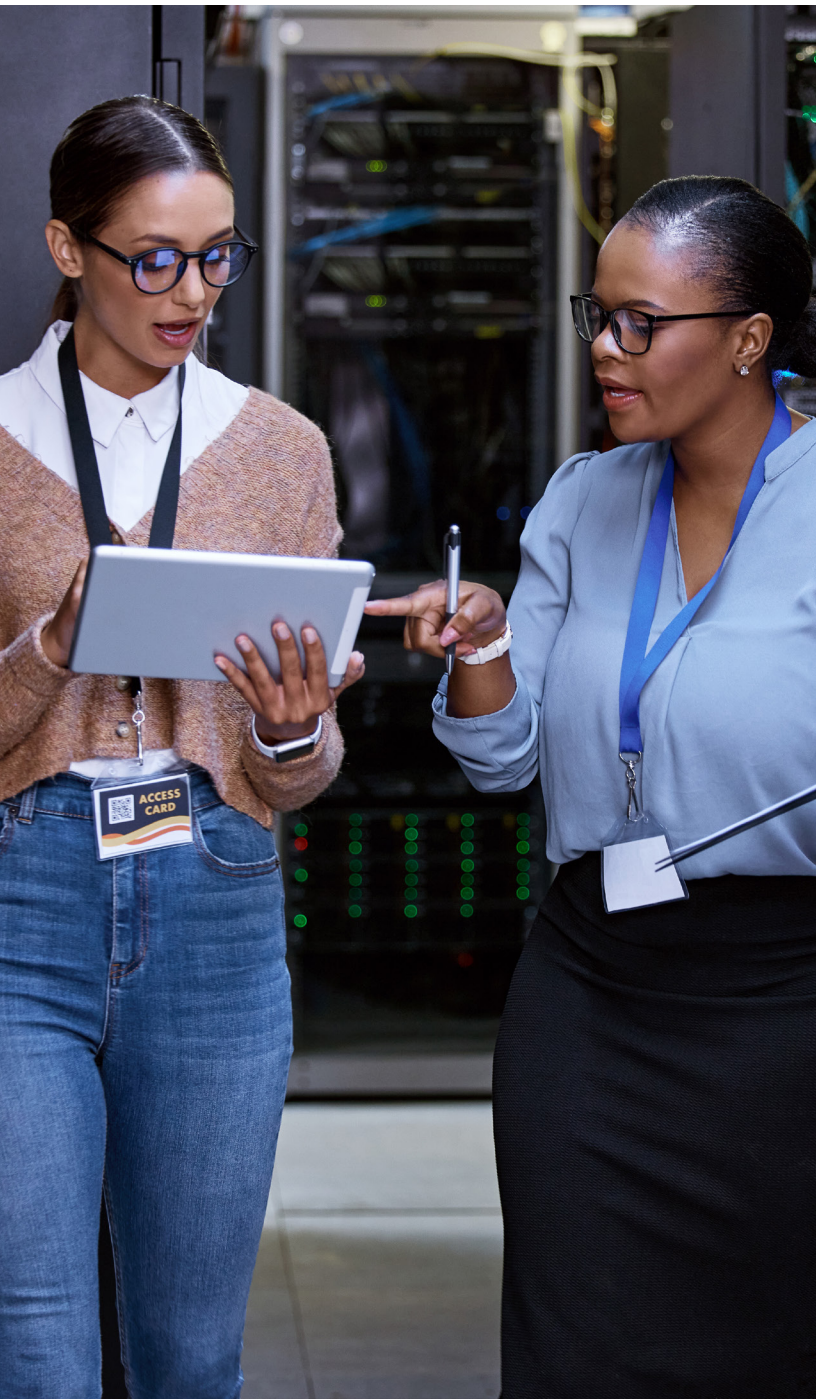

The self-service capability and low entry cost of cloud computing has left several corporations in breach of their regulatory framework policies. However, these regulatory framework policies are being revised across industries to become more cloud friendly. Bi-modal operations and change management are included in policy updates to ensure governance, controls, and processes are implemented across the new hybrid solutions.

In addition, the shared responsibility model for security and compliance between customer and cloud service provider varies based on the cloud service model of IaaS, PaaS and SaaS. In absence of detailed understanding of demarcation boundaries with a cloud provider, enterprises are exposed to security and compliance risk.

The data residency compliance for many industries and / or enterprises, especially with critical infrastructure as identified by local regulatory authorities, is limiting some of the resiliency and security benefits available from cloud. Some of the cloud providers, such as Microsoft Cloud for Sovereignty, are creating an environment inside of Microsoft’s cloud services that meets the compliance and security options needed in the public sector through a particular emphasis on data residency, sovereign controls, governance, and transparency.

Commercial Efficiency

Cloud’s usage-based pricing for infrastructure, platform, and software, continue to be available to clients, providing scalability and flexibility benefits. This has increased the rate at which fail-fast practices are being adopted to move resources to more successful ideas. However, it does require changing financial forecast models when moving from capital planning to expense planning to set and manage IT budgets.

Because of different billing models, enterprises often face unexpected fees in terms of data storage, network egress. In addition, traditional cost management tools are being ineffective at cloud cost management. Hyperscaler native and third-party offering have brought Financial Operations (FinOps) practices to cloud infrastructure of all flavours. Predictive models around cost and management have brought financial management into the modern age. Consumptive cost management is different from capital acquisition-based cost models in timing and frequency, and shifts are in progress to make the process nimbler.

The substantial increase of consumption during the pandemic was offset by reductions in costs elsewhere in the enterprise. Finance teams are learning to pro-actively report and manage expenses in the cloud more effectively, ensuring CIOs and technology teams are aware of their consumption and can adjust real-time to ensure they stay within budgetary commitments.

Data centres are likely the largest stranded assets as workloads move off-prem, and more datacentre floor space is freed up, which deliver no cost advantage to enterprises that own data centres. For those enterprises that lease space, typically lease adjustments for reduced consumption occur on an annual basis (if not longer). Some hyperscalers are offering to purchase customer data centres as incentives to accelerate cloud adoption.

Any application migration – be it to cloud or just to a new version of infrastructure – will carry with it duplicate costs for running the old and the new estates. In many organizations, the 'bubble cost' can be smoothed out or offset, with the ability to depreciate new assets, because of re-writing or refactoring applications to run in the cloud. Also, typically a large percentage of applications are retired through application rationalization efforts prior to moving to cloud.

Environment and Sustainability

The environmental benefits have been realised by enterprises by reducing carbon footprint of their organisation, measuring, and recording current footprint, and reporting to the stakeholders. While these enterprises have reduced their carbon footprint by migrating to cloud, and utilising software code-based solutions compared to physical appliance, these enterprises are dependent on the cloud provider's sustainability efforts to further reduce emissions.

Many organizations are using cloud-based offerings from a variety of providers. This not only can create more management complexity, but also make it challenging in terms of monitoring the sustainability efforts at the various cloud service providers. The providers have different commitments toward sustainability, and this affects how they create and operate their cloud infrastructures.

There are other barriers for enterprises to achieve sustainability value of cloud. These include a lack of integration with existing systems and data, cybersecurity and privacy issues, governance challenges, and a lack of alignment/clarity on roles and responsibilities relating to cloud ownership. To help overcome these and other challenges, enterprises are creating a governance program to reap maximum benefits related to environment and sustainability from cloud. In addition, few cloud management platforms feature artificial intelligence and machine learning capabilities to better manage workload costs and sustainability initiatives.

Most of the value realized from adoption of born-on-the-cloud principles have been from new application development projects or treatment of applications can tolerate differences between an on-prem infrastructure architecture and that of a Hyperscaler. Enterprises find themselves in the next phase of deployment and progress is materially hampered by applications that cannot be migrated to Hyperscaler infrastructure architectures.

What are the unintended consequences of cloud adoption?

The speed of adoption of cloud by enterprises, for gaining operational efficiencies and benefits identified in above sections, has been accelerated by pandemic and hybrid mode of working. In many instances, this adoption has been based on technology and consumption model, to achieve intended immediate benefits, without considering other factors such as overall business strategy, architecture, operating model, and regulatory compliance.

Like with any new technology adoption, enterprises must be aware of and address unforeseen consequences and unexpected risks associated with cloud adoption. These consequences and risks might apply more to specific industry verticals due to political, economic, social, technical, environmental and location factors.

The following sections will identify some of the unintended consequences of cloud adoption for the enterprises with legacy applications environment.

Culture

People and processes are key to success of many new technology transformation programs. Similarly, moving to cloud requires mindset shift for an enterprise including IT operations, application development, testing, risk, security, and other functional business units. The "no need to turn off" mindset from on-premises environment, where developers' don't have to shut down environment when not in use, results into bill shock and uncontrolled costs in cloud environment, due to utility-based consumption model. Adoption of FinOps practices is central to a consumptive management model to gain visibility into cloud costs, find value in cloud investment and maintain financial governance.

Migration to cloud environment is a multi-year journey for large enterprises as they assess the applications and identify their path to cloud. Many large enterprises decide on maintaining their on-premises environment due to various factors such as security, latency, regulatory compliance, and high project cost to rearchitect legacy applications. While the cloud comes with its own automation and orchestration tools on the platform, lack of programmability, automation, Development and Operations (DevOps) tools and integration for on-premises environment has created an agility gap between cloud and on-premises environment in terms of speed to innovate and deliver new business solutions. This agility gap is creating friction in the enterprise and for their customers for different "order to cash" experience based on underlying platform, on-premises vs cloud, for offerings.

Tech research giant, Gartner, states that 83% of all data migration projects fail and that more than 50% of migrations exceed their budget. The approach of "we are moving everything to one cloud" has evolved to a more calibrated approach tied to the benefits of innovation. Beyond the benefits of "fail fast fail often", DevSecOps culture when properly adopted drives a culture of value ordered application backlog. Pure migration for migrations sake will give way to more rebuilding for value and a bigger percentage of legacy being sunset.

Operating environment

Adoption of cloud has created complexity in operating models and decreased the level of control that enterprises' CIO have on the environment while their accountability for complete IT environment is maintained. This loss of control is a result of multiple factors such as ease of purchase of cloud services by different business units a.k.a. shadow IT and / or cloud service provider managing different components of technology stack depending on IaaS, PaaS, or SaaS service. A repeat of the shadow IT sprawl of the 90s, with ease of use, high function and self-service nature of cloud computing sees corporations facing the same challenges they faced when personal computing mushroomed.

This continues to be a significant problem that most corporations suffer despite focused effort to maintain control of data, security, cost, and compliance of cloud computing solutions. In a world of consumerization of IT, enterprises have been forced to embrace participation of employee devices and business unit developed applications. As federated security in DevSecOps and federated data provisioning processes take hold in a data driven enterprise, federated processes will replace legacy policing processes with control through visibility.

The processes, governance, operational level agreements, service levels and service delivery framework of the current on-premises environment has further created "two speed operating model". The speed of cloud-native application development meant that they were often waiting for the traditional application and infrastructure teams to process change requests. Enterprises typically modified those traditional application backlogs to include API capabilities for data and function that can be used by the cloud-native teams to consume legacy functionality in a self-service fashion thereby decoupling the synchronous release dependency between the cloud speed and traditional speed development processes. The problems of legacy applications can often be traced to the bolting on of new

functionality and process to a monolith core without the loose coupling required for evolution of various services. The reduction in legacy perimeter is just starting to show results.

The cloud providers have defined shared responsibility models based on cloud service type, IaaS, PaaS or SaaS, within their cloud environment. This shared model combined with different service levels for similar services in multi-cloud environment has resulted into complex operating models. This includes lack of customized service level for different IT services, which might require enterprise to choose closest acceptable service level, including those related to application availability and disaster recovery.

Security and Compliance

Early experience with cloud lead to several data breaches, mostly through a lack of understanding by cloud consumers of the radical difference in cloud security model. Similarly, data sequestration capabilities often permitted sensitive data to be transmitted and/or stored outside of the country. IT teams and cloud providers are now much more sensitive to these issues, but the shadow IT teams may not be.

Several liabilities of moving to cloud were identified very early principally as the radical difference in traditional, on-prem security models and those of off-prem public cloud. As such, regulators and CISOs were on alert to be certain that existing security and resiliency postures were at least not compromised.

However, the low friction associated with cloud consumption meant that an enterprise could find themselves critically dependent on cloud capabilities with no mitigation of the vulnerabilities – be they cost, security, resiliency, or performance.

Due to shared responsibility model of different services, enterprises have reduced control on critical application availability and disaster recovery. A cloud SaaS application is pre-designed for availability and disaster recovery as part of service offering, and offloads infrastructure and application management to the cloud provider. Apart from application data, the cloud provider is responsible for everything else, from maintaining the server hardware and software to managing user access and security, storing data, implementing upgrades and patches and more.

Economics of Cloud

Despite demand-based consumption, flexibility and scalability benefits, lack of adoption of public cloud has driven new business models by Original Equipment Manufacturers (OEMs) vendors. Many OEMs have started providing consumption-based IaaS or storage-as-a-service cloud business models for enterprises hosted in on-premises data centres.

This is a welcome development providing relief on the cost side, but the legacy data centre architectures were never built with a triple availability zone per region model making it hard to match the elasticity of public cloud for fast growing services. In addition, many on-premises OEM vendors are collaborating with hyperscalers such as Azure to provide same technology stack in on-premises environment.

Furthermore, as enterprises move to cloud, there are several underutilised, stranded assets that will yield decreasing value while continuing to be depreciated. This impacts overall business case and Total Cost of Ownership (TCO) analysis for adopting cloud for enterprises. In addition, with respect to software licensing many software vendors permitted novation of traditional, on-prem licenses to cloud estates. However, several of them permit novation only to their own clouds. As applications move off servers, their utilisation diminishes, but the depreciation doesn't if the enterprise has purchased (or leased) equipment. The advent of cloud has driven traditional hardware vendors to offer consumption based commercial models to mitigate this dilemma in no small part to maintain their business by selling on-prem equipment.

For a public cloud, the financial risks are mainly related to the variable nature of costs—that is, running up the cost of using public cloud due to poor planning and requirements from the business. As an analogy, think electricity. A consumer may use an electric heater all day during winter without knowing the cost incurred until the electricity bill shows up at the end of the month. Managing cloud costs needs a level of focus, skill, and tools that were not required in the past.

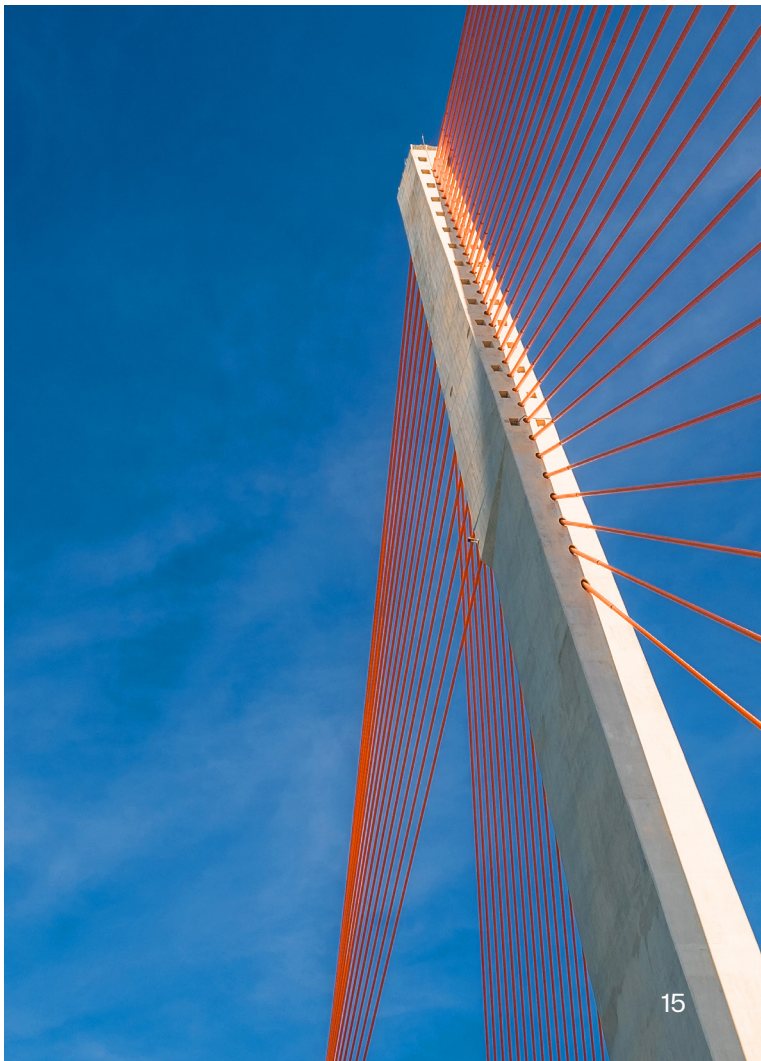

Environment, Social and Governance

Increased use of information technology continues to grow global environmental impact. Everyday part of our lives that we increasingly rely upon — our online communications — also represents a growing segment of energy use and, in turn, environmental impact. Digital transformation and move to digital business, including shift to online communications, video streaming, and online games, may seem like a greener way, but our digital technology use was responsible for about 4% of global CO2 emissions in 2019 , with this figure continuing to increase.

According to the Paris-based non-profit, The Shift Project, global emissions from cloud computing range from 2.5% to 3.7% of all global greenhouse gas emissions, thereby exceeding emissions from commercial flights (about 2.4%) and other existential activities that fuel our global economy.

The big three hyper-scalers, Microsoft Azure, Amazon Web Services (AWS) and Google Cloud Platform (GCP), who provide two-thirds of world's cloud computing continue to rely on fossil fuels for energy production. Many cloud providers utilise financial instruments such as Renewable Energy Certificates (REC) to offset carbon emissions on paper. While this may result in a good statistics and reporting outcomes, however, the carbon emissions and footprint are high.

Risk Management

Enterprises are adopting cloud at an accelerating pace to create digital business, driving operational efficiencies, and customer experience. But as they adopt cloud technologies and related operating models, enterprises often overlook different types of technology risks. New cloud technologies and related models have changed the way new risks are generated, which lays out new requirements and challenges. These risks, beyond IT systems, are financial, reputational, operational, resiliency, compliance and regulatory.

In case of hybrid and multi-cloud environment, the technology risks are further challenged with following:

- Consistent financial and operational governance and policies.
- Unified experience of monitoring and managing the entire environment.
- Reuse of assets.
- Single abstraction layer for automation across complex IT estate.
- Innovative business outcomes using technologies and applications across different cloud environments.

These challenges, while not exhaustive, either directly or indirectly result into technology risks for an enterprise, hence increasing risk profile.

The absence of consistent financial and operational governance and policies in a hybrid cloud environment result into inconsistent configurations, security posture, access controls, recovery policies, and compliance gaps. Enterprises are required to define multiple policies for same outcome specific to the cloud environment. This not only creates administrative overhead but also technical debt.

This technical debt is increasing because of development of multiple software assets, separate for each environment. Examples of these software assets are:

- Infrastructure as Code (IaC) automation assets, which help reducing time to market for application developers by removing any manual infrastructure build and enabling self-service.
- Multiple APIs to enable inter-application communication to address different use cases.

The hybrid and multi cloud environment results into lack of unified observability and asset management, creates operational inefficiencies, performance issues and longer recovery times in case of an incident or problem. This challenge is further exacerbated with lacking talent and skills development for IT department.

Cloud provides multiple technologies and services across different domains such as network, security, computing, data analytics and API integration. Absence of cloud specific skills, lack of well-architected and/or well-managed cloud environment(s), may create cyber, and resiliency risks, which further creates other risks such as compliance, regulatory, reputational and data privacy.

Data has become key for driving value for enterprises. Multiple cloud environment, private or public, cause data silos, which further inhibit applying artificial intelligence to identify actionable insights. Lack of these insights is detrimental for achieving business outcomes for enterprises. In addition to data silos, multiple cloud environments are also creating complexity due to lack of consistent compliance and governance framework.

The foreseeable future of compute: Private or Public or Hybrid?

Application discovery exercises to determine how to realize the next level of born-on-the-cloud value for the entire application estate are always characterised by high cost and material risk. Because of these factors, financial return has never justified a wholesale move to cloud. However, if the firm views technical innovation as giving it strategic advantage, the move to cloud becomes an imperative.

Hyperscaler platforms are and will continue to be the home of all innovation. Some of that innovation may make its way into on-prem products but only if the vendors deem that there is material financial return. Key differences between on-prem infrastructure and a hybrid estate where both traditional and Hyperscaler infrastructure are used are most acutely felt when looking through the enterprise business lens:

1. Data gravity is the phenomenon of increasingly poor performance, as more data an application uses and longer the data path is between the application and that data. The use of cloud in a hybrid environment, in all but a few cases, introduces significantly longer data paths. Applications performing marketing analytics, image processing, machine learning, and web crawling are classically data intensive. Typically, the data used by these applications are stored on traditional mainframes or other vertically scaled platforms. The value of an application of these types that is lodged in a Hyperscaler is radically diminished for these types of application if the data use remains on-prem. In fact, the latency in data retrieval may prove to be untenable. The pain would also be compounded by costs of data ingress and egress fees charged by the Hyperscalers.

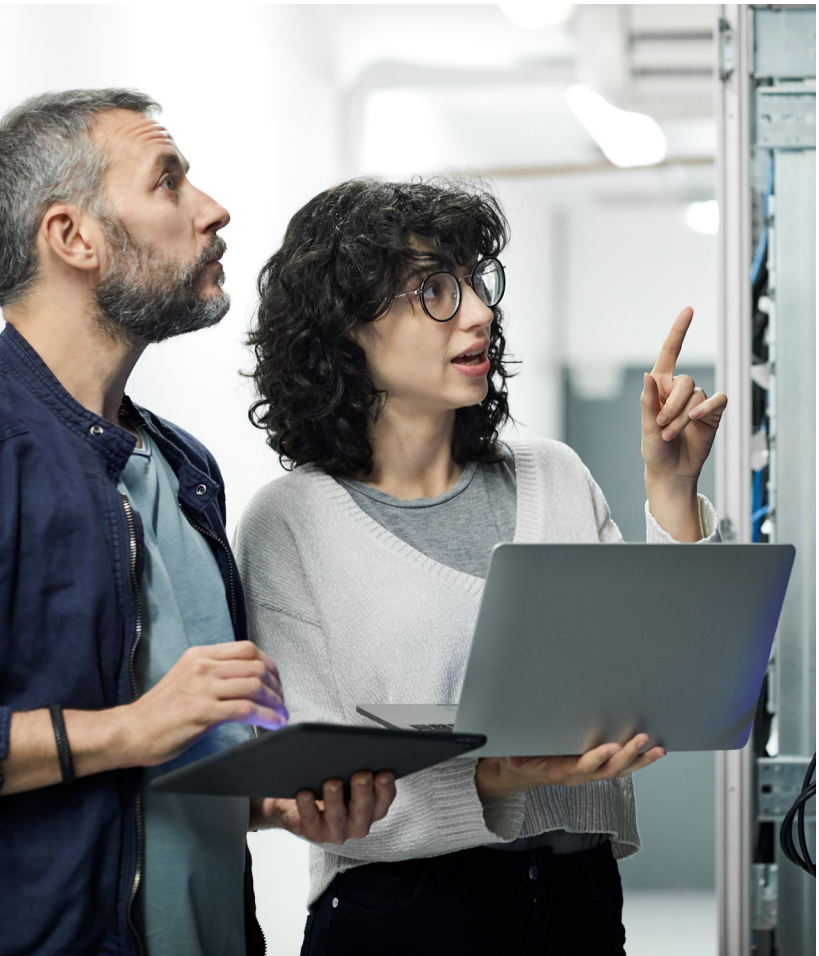

2. Security on prem is based on protecting access to the network running inside the enterprise data centre through firewall perimeter security. With effective perimeter security, it can be assumed that anyone who has successfully logged onto the network is authorised to access several/all assets; only authenticated users are using the network in the data centre. However, cloud computing networks are shared by design; all users have access to all assets that do not have individual protection. As such, a 'zero-trust' principle for cloud assets dictates that every asset must be protected by a firewall and all data be encrypted in motion and at rest. Note that stronger intrusion detection and prevention is encouraged with using cloud assets as the volume of data housed by the Hyperscalers make them a rich target for cyber threats; one vulnerability in the hyperscaler security permits access to hundreds or thousands of clients' data making them attractive to cyber attackers. In addition, cloud assets typically have multiple access means including the ability for a novice employee to quickly spin up virtual machines that are particularly vulnerable. Because of this increased threat of attack on cloud assets, additional security measures are frequently employed including multi-factor authentication, isolated cyber vaults for data, monitoring of the cloud account and unexpected behaviour within the account. Note, these techniques are increasingly being adopted for traditional, on-prem estates as the threat of cyber-attack and the recognition that even authorised employee access is a material vulnerability.

3. CPU intensive applications, unlike data intensive ones, are traditionally run on-prem on highly parallel chipsets to increase aggregate throughput for problems found in securities pricing financial portfolio optimization, scientific computing, Computer-aided Engineering (CAE), biosciences, and Electronic Design Automation (EDA). Increasingly, Hyperscalers are offering these kinds of specialised services through use of GPUs. However, for AI based applications, cloud-based computing may be cost prohibitive. According to OpenAI, the training process of Chat GPT-3 required 3.2 million USD in computing resources alone. This cost was incurred from running the model on 285,000 processor cores and 10,000 graphics cards, equivalent to about 800 petaflops of processing power. As such, the problem has been neutralised from a compute perspective; the data gravity problem likely still exists, however.

4. Inter-application communication speed may become intolerable if one application's partner application response becomes intolerably slow. Moving one of the partner applications to the cloud and leaving the other on-prem will undoubtable increase response time – potentially cancelling the transaction for tripping for the application time-out value or the missing performance expectation of the end-user. This occurs frequently enough that it is common practice to group applications into tranches that should be moved to the cloud together as a cohort.

5. Data consistency for a traditional environment is rarely a problem as all modern database management systems, transaction managers and queuing middleware ensure all copies of a customer balance, for example, are updated concurrently. This operation requires all systems recording changes to the data to wait for all its peers to complete the update. When a portion of the application is moved to the cloud, the problem of inter-application communication speed also occurs, only between the transaction middleware instead of the applications. Middleware timeouts are very likely to occur because of increased completion time so may need to be lengthened. However, increased tolerance for network latency will result in unacceptably long delays in end user interactions.

6. Proprietary programming languages, middleware and operating systems are no longer selected in favour of Open-Source alternatives thereby avoiding vendor lock in. However, virtually all enterprises have a large percentage of their estate technically dependent on these proprietary technologies. The key issue is that these proprietary technologies are not available on the most common Hyperscalers. As such, the vendor lock in remains as a material inhibitor to moving applications to the cloud.

7. Total cost of ownership (TCO) considerations for enterprises, with large number of legacy applications performing critical business functions, may not demonstrate value for money. The value for money may include key parameters such as risk, cost, innovation, security, and speed of change. The legacy applications may have to be re-factored i.e. modify code, to suit the cloud environment, which may introduce significant cost and risk. The application may be migrated as is, i.e. rehosted, which may introduce other costs and complexities of remote access tools, such as Citrix, for applications. These may be required to maintain lower latency for applications with client-server architecture.

These considerations prevent extremely large portions of an enterprise application estate to be easily moved to cloud. The industry has responded with application discovery methodologies and supporting tooling to assist in building a migration plan that will assume that a portion of the estate can

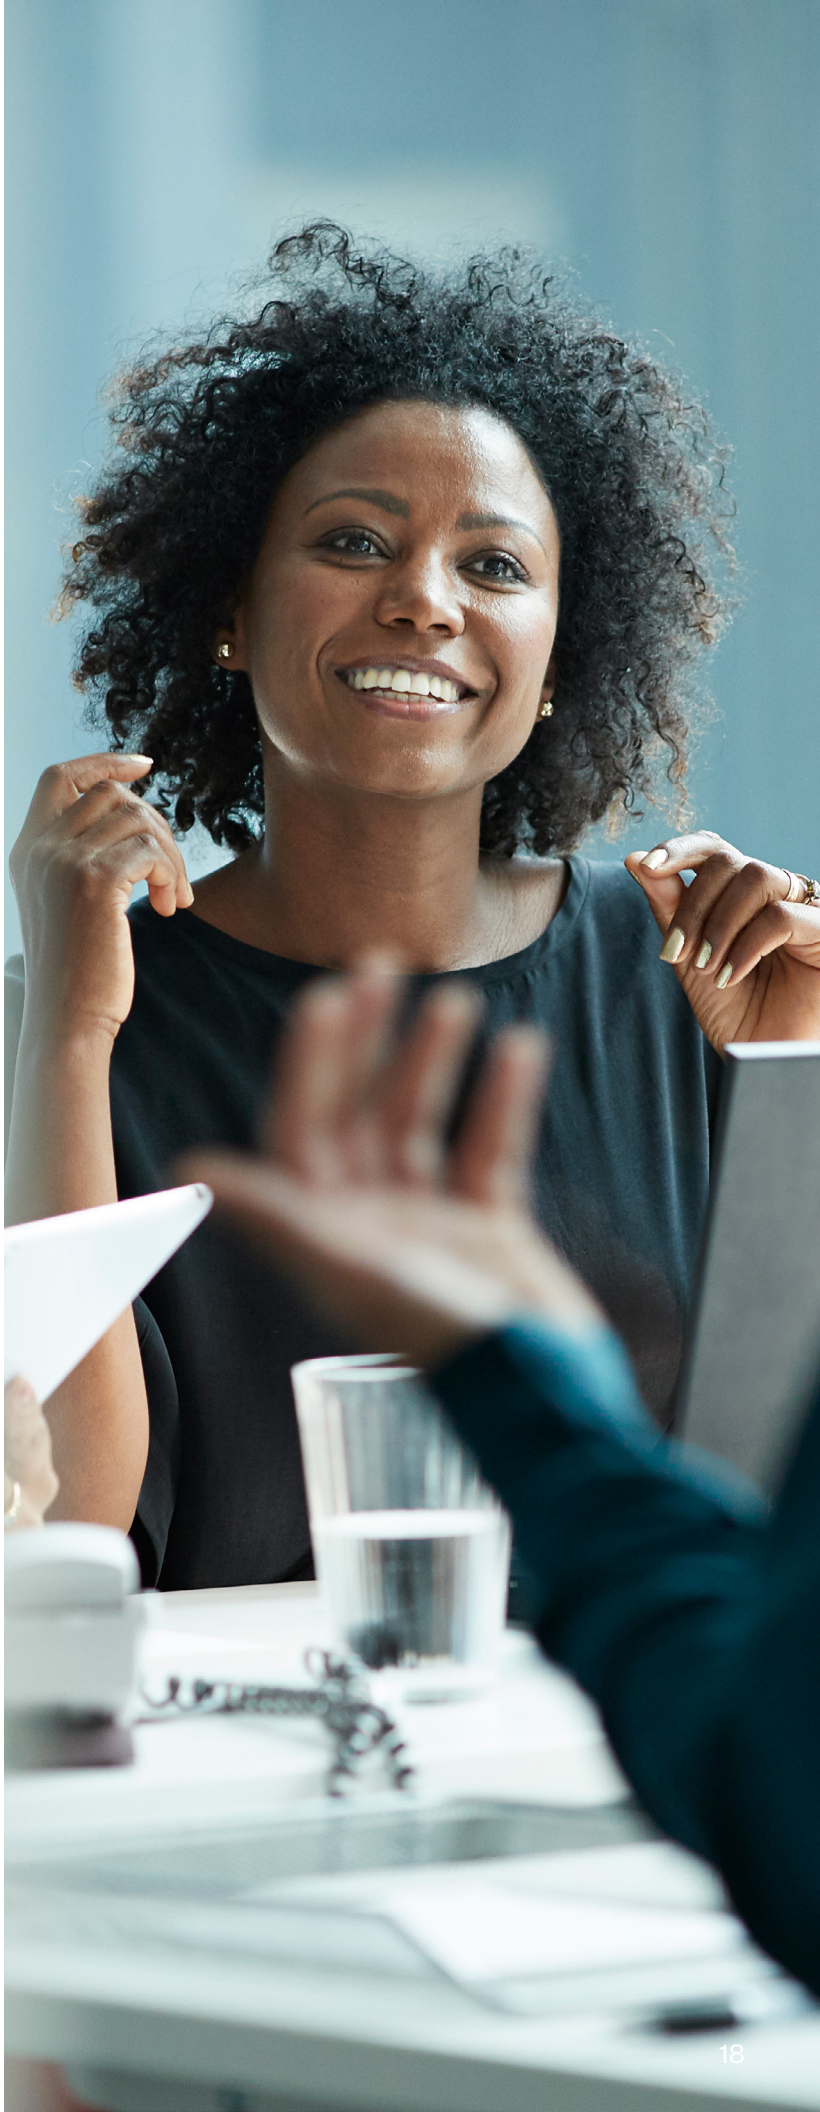

be moved to the cloud with relatively little effort, a portion can be moved with material effort and a final portion cannot be practically moved because cost or risk factors fail any internal rate of return hurdles.

Organically, over time, the enterprise may retire or re-write this final portion, but that is not likely to happen for several years or decades. This reality means that enterprises will be running in a hybrid estate for the foreseeable future – some of the estate will continue to run on-prem (albeit modified to cooperate with the cloud estate), some will be running in a hyperscaler, and some will be running in a private cloud that delivers the hyperscaler benefits of Opex consumption pricing, self-service through extensive automation, but with the security that comes from isolated compute, storage, network and data centre. and automation.

The percentage that will run on each of those three options will depend on the demand for innovation, the value that a digital customer experience will improve top- and bottom-line growth, the budget, the risk profile, the ability for the organization to accommodate a material change in its operating model and the approaches it takes to mitigate the issues outlined above. Candidate approaches are outlined in the next section.

On-prem applications rely on several on-prem infrastructure capabilities that are either not present or materially different than those of Hyperscalers. These technical differences have several remediation tactics that form the basis of application discovery. The born-on-the-cloud value proposition is limited for these legacy applications and different remediation techniques have emerged as best practice. This widespread model assesses large portions of an application portfolio and determine one of the following dispositions for each of those applications.

1. **Rehost a.k.a. Lift and Shift:** Perform migration of application, with no change in code, from on-premises to cloud infrastructure.
2. **Retire:** Identify lack of application's need for the enterprise, resulting into decommissioning of the application and negate its migration to cloud.
3. **Retain:** Maintain application on-premises due to strategic alignment and / or other restrictions such as regulations.
4. **Re-platform:** Make applications compatible with cloud environment using emulation or other approaches, without modifying application architecture.
5. **Re-architect a.k.a. Refactor:** Modify application architecture and code to maximise benefits from cloud environment.
6. **Replace:** Identify application and/or service that can replace functions performed by on-premises legacy application

By adopting wider context of strategy, financial, skills, risk and technology, enterprises will combine above dispositions when planning for future roadmap.

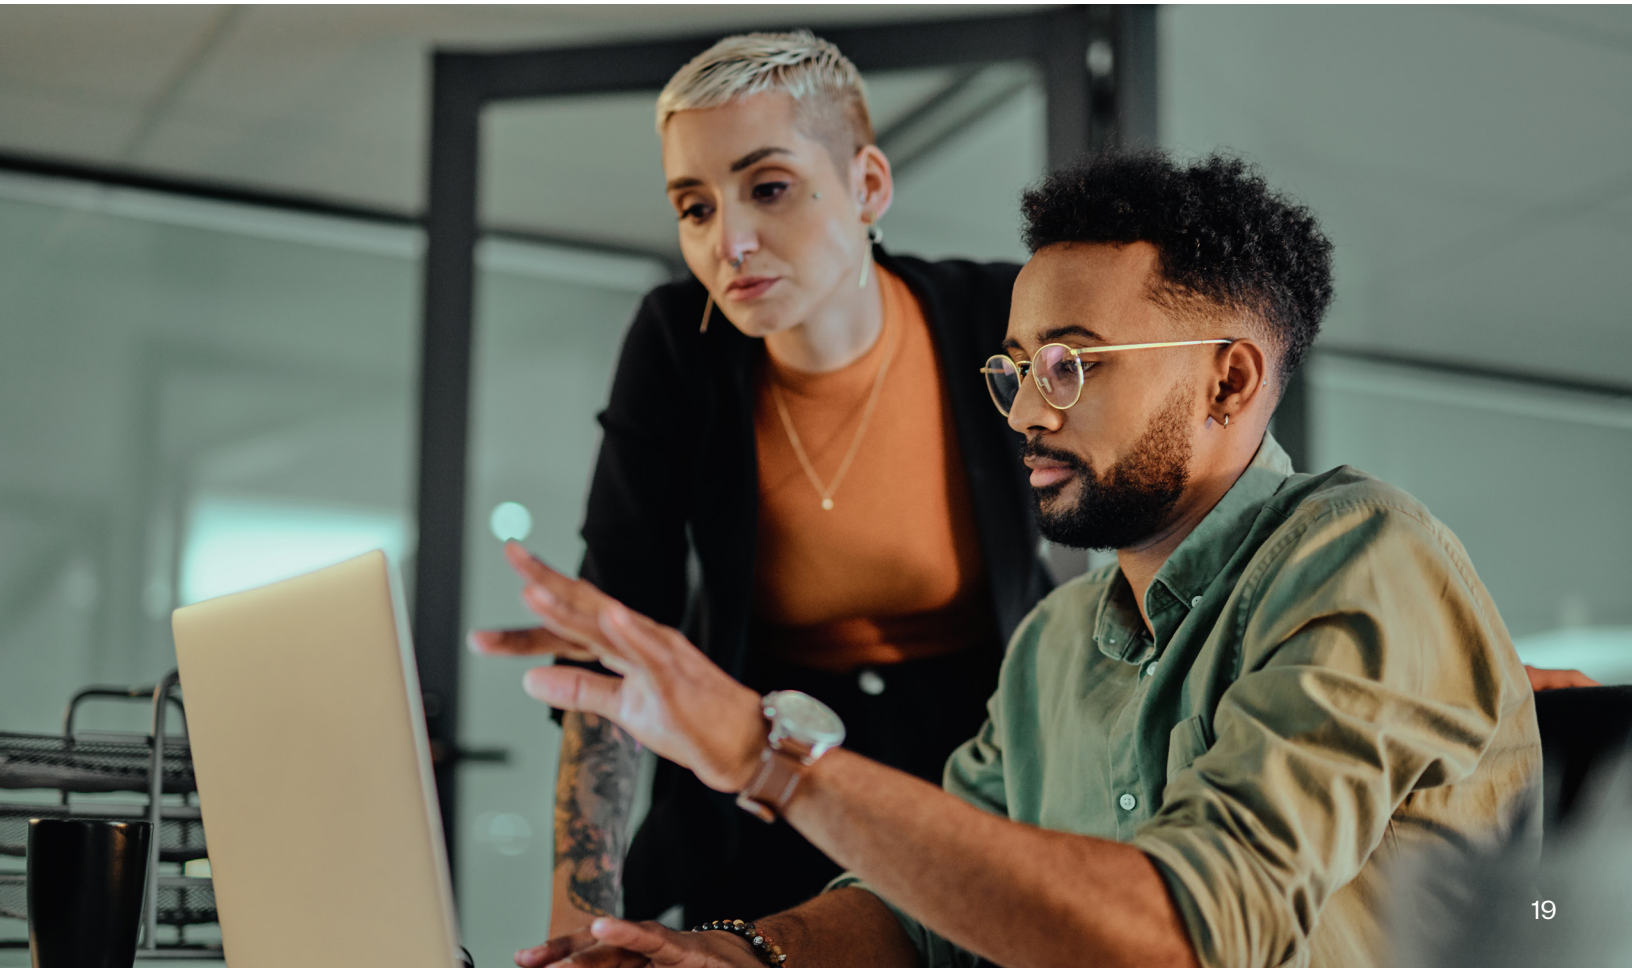

What is the decision matrix?

While many large enterprises had a defined strategy and plan to adopt different cloud services, but Application discovery exercises typically address the differences between on- and off-prem infrastructure capabilities as outlined in the following table.

Attribute	Remedial Treatment
Data gravity	i. Establish a strong meta-data management team to outline and deliver a strategy for the enterprise data inventory and life cycle management. ii. Adopt eventual consistency application architecture for all but the most time critical applications and start distributing data to where it may be needed. iii. Establish a federated query architecture and associated APIs to permit queries to be run close to the data iv. Consider moving large legacy systems to cloud adjacent data centres
Data sovereignty	i. Identify the regulatory compliance, data residency, risk management and resiliency requirements for business applications based on the functions performed by the application. ii. Establish and mature security governance framework and process end-to-end for entire IT estate, including on-premises private cloud and off-premises public cloud. iii. Perform effective security risk assessment of public cloud(s) especially for business-critical applications, including support resources, shared responsibility models and regulatory compliances. iv. Understand applications ability to support cloud native architecture attributes such as modularity, programmability, and resiliency
Talent and Skills	i. Adopt a shift left approach in designing cloud architecture and systems that will easily adapt to changing customer, business, and compliance requirements. ii. Build automation into the environment and application code, to enable orchestration of infrastructure including security, availability, and scalability. iii. Transform organisation culture across engineering and operations teams to Agile and DevOps to reduce time to value by increasing automation and integration.
Compute intensive	i. Understand on-premises compute usage patterns and metrics for applications, identified for cloud, based on daily and peak business seasons. ii. Identify the appropriate compute series type and hardware chipset in the cloud, based on function performed by applications. For example, memory optimised compute for heavy in-memory workloads such as SAP HANA or high-performance computing for financial applications or storage optimised compute for low latency and high throughput applications iii. Consider PaaS or SaaS application alternatives for on-premises applications, where possible, to leverage native cloud capabilities while being cost effective, in comparison to lift and shift using IaaS.

Data intensity	i. Understand the status-quo of data system for data intensive applications and, establish target state data system strategy and architecture based on use-case. The data system may include components such as database for data storage, search index for enabling search capabilities, cache to accelerate data reads, stream processing for message flow and/or batch processing for crunching bulk data. ii. Define use-cases for data intensive business applications and categorise based on underlying function such as On-Line Analytical Processing (OLAP) if they are performing analysis and/or On-Line Transaction Processing (OLTP) if they are performing transactions. iii. Design data storage and querying cloud capability based on data model type such as the relational model, the document model, or the graph model. iv. Consider combining multiple cloud services as part of data system architecture to achieve the reliability, scalability and maintainability requirements for applications and supporting infrastructure.
Data consistency	i. Establish data consistency strategy and requirements for applications based on status-quo, use cases and situations. ii. Define target state cloud strategy for applications based on data consistency, performance, cost, and availability factors. For example, strong consistency may be the status-quo for application in on-premises data centres but strong consistency in cloud geo-replicated services, that span multiple data centres, will generally result in higher cost, lower performance, and reduced availability for reads or writes or both. iii. Ensure target cloud architecture addresses the desired data consistency, whether strong or eventual or somewhere in between, and associated performance and availability guarantees.
Inter-application communication speed	i. Document inter-application dependencies and message flows required to address business use case, especially where applications are across hybrid environment. ii. Establish target state underlying infrastructure strategy considering applications' communication characteristics, such as latency and packet loss, and its ability to adapt to varying network conditions, such as TCP congestion and bandwidth control. iii. Define API and mesh strategy for inter-applications communication, depending on RESTful and general request/reply to interactions or asynchronous event-based interactions. Applications based on container based microservices may have service mesh architecture that provides reliable, fault tolerant and load balanced communication. However, event mesh will be more appropriate for communication between legacy on-premises and cloud native applications
Legacy hardware, programming languages and middleware	i. Establish a squad to define strategy and architecture for applications that are targeted for cloud migration and run on legacy non-x86 hardware or use programming languages and middleware that are not available on cloud. ii. Identify and define short- and long-term migration options based on value for money including business benefits, technical advantages, time duration and cost. iii. For short-term, consider different emulation options for legacy hardware and middleware, available with cloud providers, to avoid changing application code and migrate to cloud with short project duration and minimal cost. iv. For long-term, understand the relevance of application in alignment with business strategy and define strategy considering legacy automated refactoring, rewrite of application code and/or alternate cloud service options performing same function

Assessing the enterprise cloud strategy

There are several benefits outlined above – as well as challenges; the question turns to arming decision makers and senior stakeholders with the instruments to assess the efficacy of the plans that the technology leadership is recommending. Key considerations in that assessment include:

1. **Leadership team** – does the leadership team have seasoned professionals from both the application and infrastructure endorsing the plan? Does the team have a mix of leaders who have managed all elements of traditional application and infrastructure life-cycle management as well as talent experienced with the full stack of native cloud development? Are they incentivised to cooperate through joint remuneration schemes that support the KPIs outlined below? If not in-house, do they have access to that talent in the form of a trusted advisory board or consultant?

2. **Regulatory** – what is the extent to which the current regulatory framework has been adjusted to accommodate the proposed strategy? Is it minor modifications to accommodate new technologies but the same primary and secondary controls can be used? Or does the proposed strategy require material re-work of the framework. Have other financial institutions secured regulatory approval with similar solutions?

3. **Adequate cost** – is there sufficient funding to get through at least the build out of the fixed and one release of an MVP? Are the innovation benefits being adequately assess and added to the hard benefits?

4. **Risk** – is the execution plan composed of incremental deliverables or Minimum Viable Products that permit incremental investment and return. How much of the initial phases deliver common capabilities that will be used in all subsequent phases?

5. **Fallback plan** – what are the alternatives if a target platform is unable to deliver the requisite capabilities? Is the solution abstracted to a point that it could be built on another cloud provider? On existing, on-premises technology? On a private cloud that may run on prem or on a hyperscaler?

6. **Organisation change** - Is there a comprehensive organisational change plan? Are the roles, responsibilities, deliverables and KPIs well understood by the teams operating in new ways of working? Does that org change plan include reshaping of the roles of executive management – CFO, COO, CIO will all need to manage in a hybrid environment while the strategy is executed and the projects being run in the new way necessarily have different methodologies, deliverables and KPIs. The senior management team will need to be educated on best practices for assessment of progress, risk, and success. Is the construction of the agile squads suited to the objectives? Squads responsible for delivery of functions to end clients should have participation from all disciplines to adequately prioritise trade-offs between functional and non-functional requirements. Failure to do so will lead to a surprise accumulation of technical debt as funding is typically limited to fulfilling new functional requirements; it is often assumed that native cloud application development places the onus of technical debt on the hyperscaler, there are other classes of technical debt that are introduced and often missed. Most common examples include the need to perform application testing when Hyperscalers upgrade, patching of cloud components is generally the responsibility of the consumer, CI/CD automation scripts to accelerate application development are themselves code that must be maintained.

Organising for success

The success of cloud adoption depends on many factors beyond technology services. The journey to success starts with understanding and assessing Political, Economic, Strategic, Technological, Environmental, Operational and Regulatory impacts of adopting cloud. The following sections highlight considerations for an enterprise to plan for successful cloud adoption.

1. Operational support model - Agile squads to cover the full application stack.

The advent of programmable infrastructure components means the infrastructure teams managing those components can move at pace with the application development teams. As such, the product squad and its leader can now assess and prioritise investments in non-functional requirements.

Before the advent of programmable infrastructure, the product manager would negotiate capabilities in the infrastructure for her application. Typically, the funding for the changes in infrastructure are covered in the product development costs, but the ongoing run costs are allocated to the infrastructure budget. As infrastructure requirements change (typically because of new releases, increased regulatory requirements or a desire to improve cyber resiliency), most firms find their run budgets are underfunded and technical debt accumulates. The infrastructure teams often try to negotiate budget relief from the product areas but typically these non-functional enhancements are prioritised lower than are the product's functional requirements.

As proportion of infrastructure on public or private cloud increases, most of the burden of technical currency is shifted to the public cloud providers (and, increasingly to the private cloud providers). These shifts to cloud are typically accompanied with usage metering at the product thereby eliminating the pooling of infrastructure costs for the entire estate and recovery of infrastructure costs through relatively fixed allocations to the product managers.

While the product managers can now be billed much more precisely for their infrastructure consumption, there is typically a material increase in their ability to tailor infrastructure to support both their functional and non-functional requirements through automation programming of that infrastructure. This includes augmenting the product squad to include SREs. Having an SRE on the product squad elevates the consideration of non-functional requirements (including technical debt) to the same level as functional requirements so the accumulation of technical debt is at a minimum understood and consciously address by the product manager rather than the debt being a moral hazard.

2. Guard-rails to mitigate unintended consequences of migration to cloud.

As outlined above, there are several unexpected risk areas that can materially reduce the expected value of cloud computing. As such, care should be taken to ensure that the custodians of cost, risk, performance, client experience, regulatory compliance, purchasing, resiliency, and line of business leaders are educated on how the risks in their area of responsibility can be quantified and common techniques for mitigation:

A. The Finance Function is often the victims of unexpected fees – either for unfunded projects or budget over-runs for authorised projects. The ease of use of cloud compute is due in large part to the absence of the need for centralised IT support to provision compute. Individual employees can use a credit card (personal or corporate) to initiate service with a Hyperscalers. Similarly, unlike on-prem, capacity-priced compute, cloud compute is consumption based and it requires a change in mindset to turn off un-used compute. Failure to do so means that charges are accruing without associated real use. Finally, items that are covered in an on-prem, capacity-priced compute are not necessarily covered in a cloud deployment. Fees to transmit data from the on-prem estate to the cloud and back (ingress and egress charges) are often missed when planning a cloud deployment – particularly when an application lodged in the cloud has dependencies on manipulation of large quantities of on-prem data.

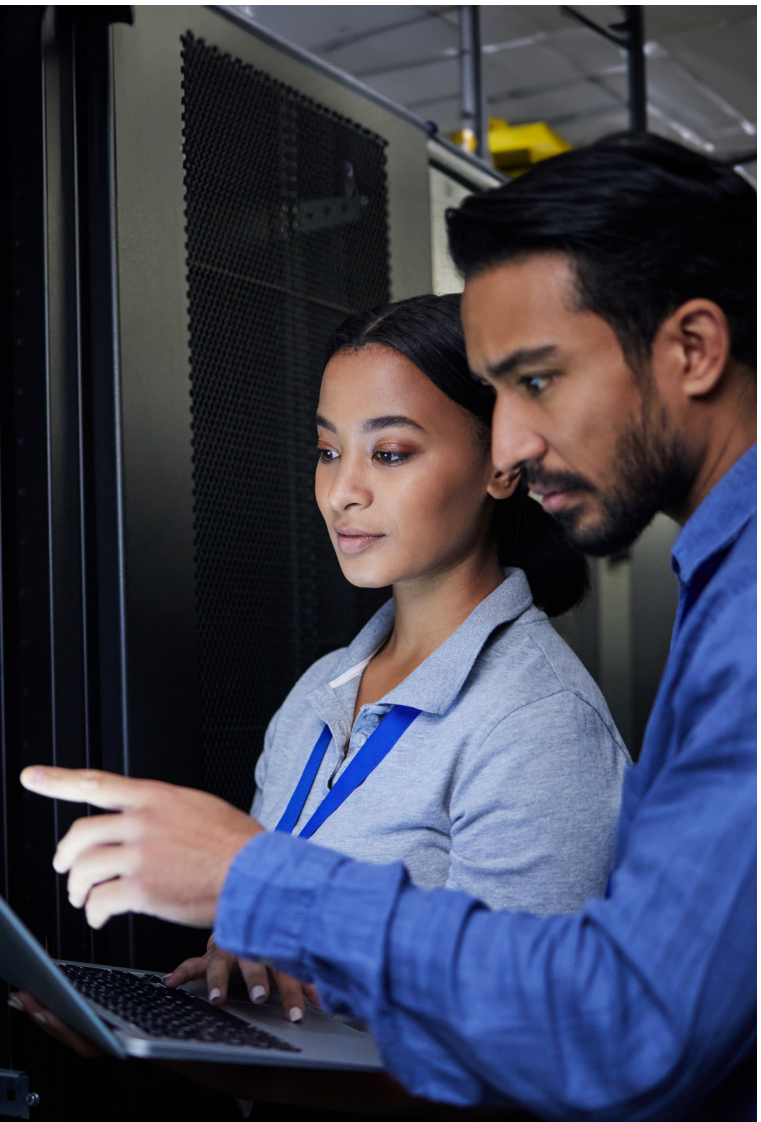

Fortunately, the market is responding to this phenomenon with Cloud Management Platforms and the associated practice of FinOps that include telemetry to detect charges for allocated, but un-used cloud assets and make recommendations to trim/turn off those assets along with support for detailed project/application level billing through the use of tagging.

As the entire enterprise is susceptible to the phenomenon of billing for un-used or un-expected cloud assets, the FinOps and any associated Cloud Management Platform should be an enterprise-wide practice and platform.

B. Risk and Compliance teams, in the early days, faced a major concern and deservedly so. Early versions of cloud platforms had little visibility to geo-placement of data. Nor did they provide services isolated from other enterprises. Furthermore, visibility to what other firms are sharing the infrastructure is not available. Similarly, while most enterprises are able to visit co-location data centres that they may lease for on-site inspection of physical security and access protocols, Hyperscalers generally only permit extremely limited access or referals to third party auditor reports. This often means that the CISO is unable to indemnify the security of the cloud environment, and this often means that the firm is out of compliance with their regulatory framework.
Again, just as with FinOps, cloud providers have matured as well as the IT practices of cloud consumers. Many cloud providers now enable geo-fencing of data, and some do provide limited isolation of asset use. And, to inoculate against exposure of sensitive data to third parties, the practice of data encryption in memory, in motion and at rest with the firm retaining possession of the encryption key generally suffice.

C. IT Capacity and Performance discipline should have a new moniker: IT Consumption Costs and Performance. The elastic, and virtually unbounded capacity of cloud assets mean that this team no longer must estimate future capacity requirements and engage in long lead-time procurement processes. However, as is the case with the finance team, the capacity responsibility shifts to prudent optimisation of cloud asset mix and consumption.

While the issues with capacity have shifted more toward prudent selection of cloud services, performance issues are the same as on prem, but much more intense. Factors that need to be considered to optimise network performance are as follows:

i. Throughput –Network connectivity using shared internet or MPLS, with or without Software Define Wide Area Network (SDWAN), may have performance issues such as over-subscription, bandwidth, congestion, peak-hour considerations, etc. will have a material effect on data throughput. Service provider SLA's will govern the WAN performance especially with regards to quality of service.

ii. Cost vs Performance: The cost and performance parameters are to be considered for network connectivity. For example, a dedicated direct connection to the cloud provider provides better performance at higher cost, while network connectivity through a cloud gateway or a shared Internet / Multi-Protocol Label Switching (MPLS) will be lower cost and performance.

iii. Data Tromboning – The lack of optimised network design may lead to data tromboning where performance degradation results from data needlessly going back and forth within the network, from initial source to final destination. Material latency is added to the transmission and possibly packet loss.

Finally, with more complex network topology, traditional on-prem diagnosis are not robust enough to tune network performance. As such vendors and software organisations have specifically created visibility tools focusing on end-to-end network performance from an applications point of view.

D. Chief Digital Officer – This position is typically not present until the enterprise fully embraces the role that cloud plays in the client experience. Being a new position, the CDO will take on the challenge of defining client use cases and then designing the U/I and associated new, and modification of existing, applications. It is often the case that that modification means some level of re-factoring the existing estate. Furthermore, the role carries the additional responsibility of ensuring that the implementation of those use cases do not compromise the security, regulatory or cost objectives set out by the enterprise.

E. Regulatory compliance – The compliance team will need to work very closely with the security, application, and infrastructure teams to closely examine the existing regulatory framework to test the parts of that framework that would be no longer effective for a proposed new application or modification of an existing application to cloud. Understandably, the business will be extremely reluctant to modify the framework to accommodate a new, cloud-based approach. However, the current state of maturity of Hyperscalers has grown significantly through increased visibility, security and integration capabilities being delivered as part of the platform. Furthermore, there are several reference architectures for infrastructure and applications that have been proven to meet all but the most stringent regulatory requirements.

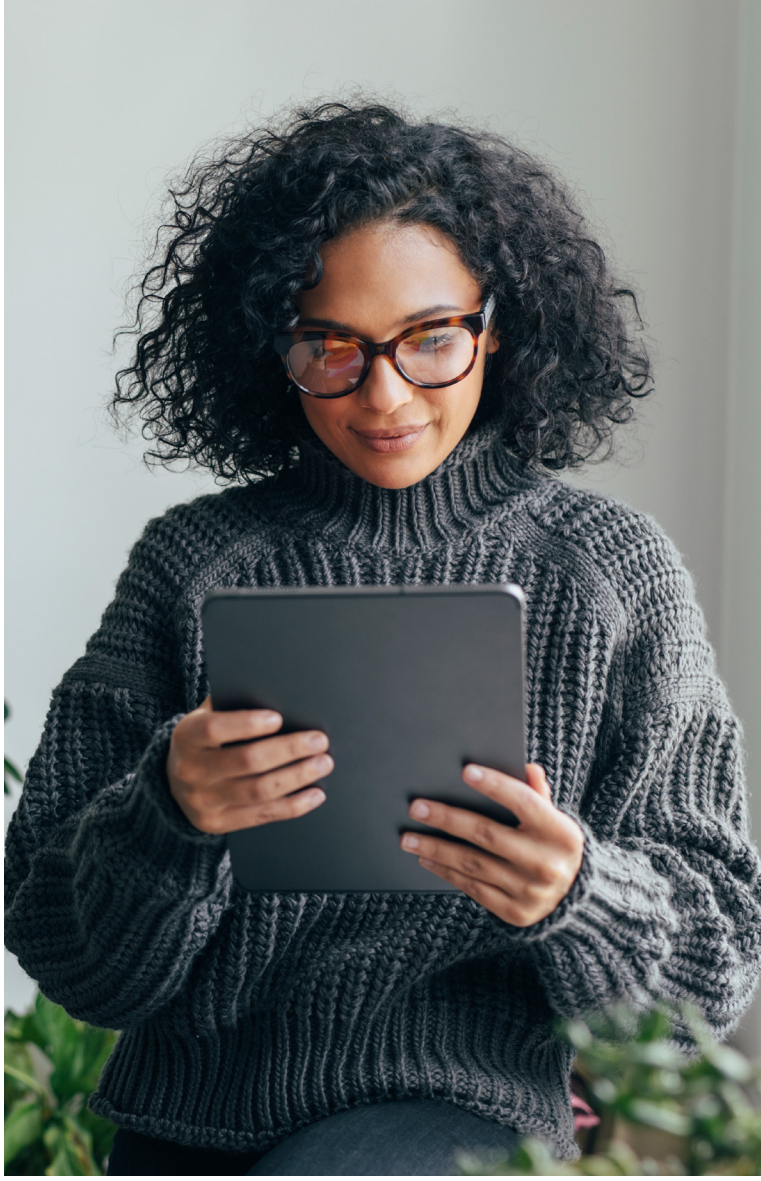

The best practice around interaction with the regulator is to, if possible, incrementally introduce modifications to the framework as the deployment of cloud systems progresses. For example, start with a simple application whose only modification is that data be encrypted in motion and at rest.Implementation of this single measure represents one step and can be approved by the regulator before the next, incremental step can be taken.

F. Sourcing – This community will likely see the biggest changes in their approaches to contracting. One of the biggest reasons that Hyperscalers can deliver self-service catalogues is the exclusive use of standards across their fleet. This means that there are very few (if any) items for negotiation. There may be several versions of hardware and software offered, but it is unlikely that the hyperscaler will accommodate any consumer requests to substitute any capability for another.

While this does represent a simplification in contracting, it becomes imperative that the sourcing team be clear with the enterprise that there is a material, residual risk that remains with the enterprise. This often means degraded non-functional requirements when compared with on-prem estates. Examples include the ability to extract service level penalties from an SI who may be running the on-prem estate, to sweat assets beyond their supported life to avoid the costs of testing upgrades, to use existing telemetry reporting schemes to ensure cost and regulatory compliance.

In short, the sourcing burden is reduced in part through simplified contracting, and in part by transferring risk from the supplier to elsewhere in the enterprise.

G. Business Continuity takes on a particularly large change in complexity in part because of broader dispersion of data and applications. Traditional mission critical systems achieve resiliency by having redundant copies of data stored farther away that the anticipated blast area of a disaster. For extremely sensitive data (electronic funds transfers, for example), traditional infrastructure ensures that any given modification of those data are safely, and permanently stored in both of the locations before the transaction is completed. This means that a credit card charge using an iPhone wallet will be delayed until the transaction is recorded in both the primary and the backup systems before the user receives acknowledgement. The speed of the network path between the phone, the primary system of record and its copy and back has the most influence on transaction completion time.

   As enterprises contemplate moving more sensitive data to the cloud, the requirement for tightly coordinated synchronisation across all the possible storage locations is often overlooked. Only with extremely sophisticated application and infrastructure architectures can an enterprise be assured that their mission critical data and applications are resilient in the face of failure of any of the components supporting the application – be those components on prem or in the cloud.

   This same phenomenon of tight coordination of version of applications requires sophisticated application promotion protocols. While this has always been true for traditional two site resiliency, the advent of blue/green deployment increases complexity as the enterprise is running different versions of the same application concurrently.

   Fortunately, the ubiquitous use of automation of infrastructure management and the nascent, but increasing, use of machine learning permit the enterprise to effectively manage this increased complexity. However, the business continuity cohort must understand these new issues and plan accordingly.

H. Product Managers should be prepared to move their product development teams to Agile product squads and include all disciplines.

   The job starts with developing specifications for the product and then become an active participant in the day-to-day activities of development – from product initiation throughout the entire lifecycle of the product. The means to chair all sprint planning sessions, review and each story before its sprint starts, be readily and regularly available to all developers to share the strategy as well as answer detailed implementation questions, be present for all sprint demos, ensure that the backlog has all functional and non-functional requirements recorded and identify the sequence requirement implementation.

Furthermore, in the case of an organization not completely operating in an agile manner, to serve as the interface between the agile squad and the balance of the organisation. This becomes crucial as a traditional development team will have KPIs. In an agile model, the most common KPIs include:

a. Return on investment vs. efficiency of delivery
b. The NPS of the product vs. the project success
c. Outcomes realized vs. fulfillment of requirements
d. Turnaround time of feedback vs. limited number of changes authorized
e. Focus on team output vs. the contribution of individuals
f. Assessment of improvement of team capabilities vs. growth of individual skills
g. Drive for team cooperation and reward accordingly vs. rewarding individual accomplishments

These measures can be very difficult for the senior leadership team to adjust to. The task of educating and re-enforcing correct behaviours for those not operating in the agile squads falls to the product manager.

I. Enterprise Engineering must establish a comprehensive catalogue of capabilities and a safe, secure set of CI/CD tooling to permit frictionless incorporation of those capabilities by the development squads. Furthermore, they must codify policy that controls the behaviour of the CI/CD pipeline to ensure that output is compliant with all the enterprise policies around data placement, security, licensing, performance, risk, regulatory compliance to name a few. This team must also set up, and lead, the tribes and guilds across the enterprise and the associated software engineering methodologies to be deployed and assess assets developed for harvesting and registration for hardening and promotion for enterprise-wide consumption. Note that assets can come from internal teams, Hyperscalers, traditional infrastructure/middleware vendors or SaaS offerings. To avoid introduction of the friction traditionally associated with use of common assets (internal or 3rd party), consumption of catalogue capabilities must be self-serve in terms of issuing credentials, provisioning services, monitoring, metering and maintenance of the assets.

This enterprise catalogue will be curated by an interdisciplinary team with representatives from all the disciplines the bank uses to keep is efficient and safe – i.e. security, compliance, risk, finance etc.

J. Two speed vs. New speed – the ever-advancing use of automation for the provisioning and management of infrastructure is now emerging in traditional infrastructure to the point that the need for two speed is diminishing in favour of new speed. In addition to the introduction of APIs for the configuration and management of traditional infrastructure, the penetration of New Speed is paced by the re-factoring of the application estate to capitalise on micro-services or at least the exposure of functional API capabilities to the CI/CD pipeline.

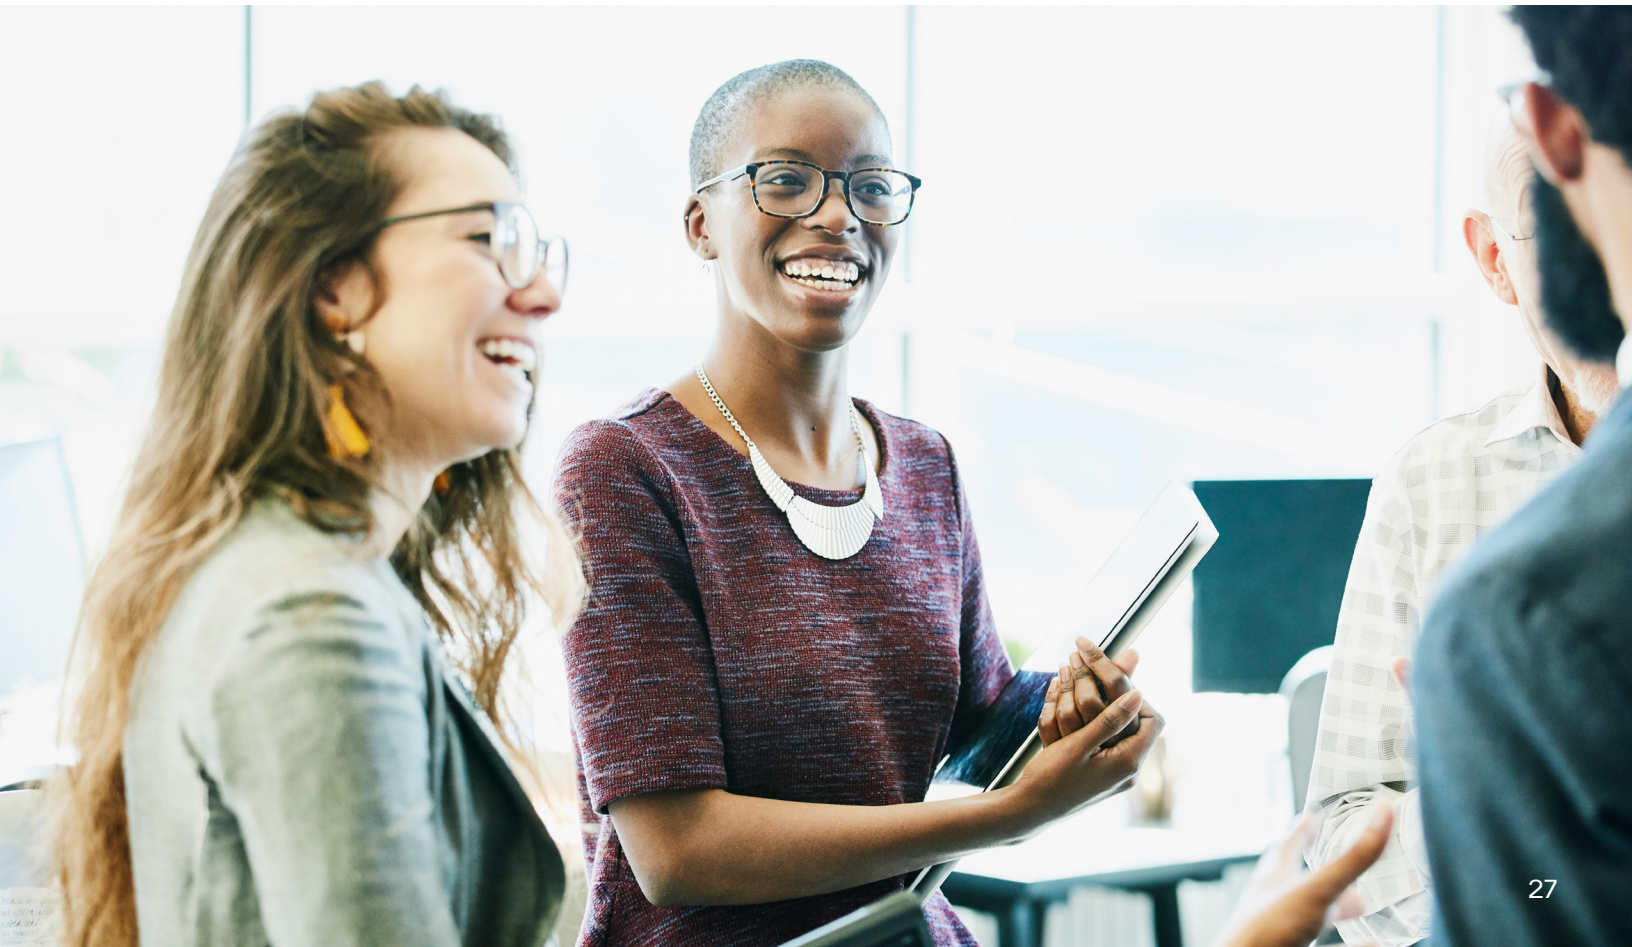

What does it mean for Financial Services Industry?

The future-of-compute will vary for industries and for individual companies within an industry depending on the number of life years for company, sub-sectors, evolution of the industry and policies impacting that industry. The Banking and Financial Services industry is no exception with varying pace of adoption for public, private and / or hybrid cloud. The banking industry has been undergoing major evolution in recent years driven by regulatory changes, technology developments, customer experience, and global economic factors. To understand future of compute for Banking industry, it will be prudent to know some of these evolutions, with focus on Australia & New Zealand (ANZ) market:

- **Financial technology (Fintech):** Fintech ecosystem has been evolving rapidly in ANZ, with start-ups introducing disruptive innovation in traditional banking services through solutions such as lending and personal finance, robo-advisors, digital payments/ wallets, digital identity, cryptocurrency, and blockchain-based applications.
- **Digital Transformation:** With steady adoption of digital banking services including mobile payments, digital identity, and chatbots, traditional banks have been investing in the digital technologies to improve transparency, operational efficiencies and competitive advantage.
- **Regulatory governance:** Well-regulated framework governed by the Australian Prudential Regulation Authority (APRA), the Australian Securities and Investments Commission (ASIC), and the Reserve Bank of Australia (RBA). The data privacy and security regulations in Australia, such as the Privacy Act 1988 and the APRA Cross-Industry Prudential Standard (CPS 234) on Information Security, impose stringent requirements on banking sector for safeguarding customer data and ensuring data security. In addition to CPS 234, regulatory changes, including the Banking Executive Accountability Regime (BEAR) has improved accountability and transparency in the industry.
- **Cybersecurity:** With substantial number of data losses in cybersecurity attacks over the past 2 years in ANZ, banking industry has been investing in cybersecurity measures to address threats and ensure compliance with data privacy regulations such as Privacy act.
- **Branch optimisation:** Increasing digital transformation, self-service, reducing foot traffic and higher real-estate prices are some of the factors that have driven banks to reduce their branches footprint, while making the remaining branches digitally focussed.
- **Customer experience:** Improving customer experience is key focus for banks by leveraging data analytics to gain insights and provide personalised services including tailored financial products.
- **Open Banking:** Introduced in 2020 in Australia, Open Banking is first sector of Consumer Data Right (CDR) which gives customer the ability to share their banking data with third parties that have been accredited by the Australian Competition and Consumer Commission (ACCC).
- **Mainframe Modernisation:** With 92 of world's top 100 banks using mainframe for core banking systems, modernisation of mainframes has become key priority. Based on recent survey conducted by Kyndryl , to maximize value, majority of enterprises generally modernize "on" their mainframes, "integrate" with other platforms, or "move" certain workloads off the mainframe. According to the survey, mainframes increasingly become an integral part of hybrid cloud environments and continue to drive business value. However, the lack of a skilled workforce to support and secure these mission-critical environments is a significant cause of anxiety for many businesses.
- **Artificial Intelligence (AI):** AI and machine learning are being used for various purposes in BFSI, including customer service chatbots, fraud detection, risk assessment, and personalised financial advice.

The above trends are influencing the adoption of public, private and/or hybrid cloud across a bank's value chain, which will mean that future of compute will be different for a bank based on value chain element. For example, Credit and Lending part of bank's value chain, whether institutional, or retail, is adopting cloud differently to Treasury and Capital Management or Customer Acquisition and Relationship Management. Where the value chain element has stringent regulatory, compliance, security, and data privacy requirements, private cloud is

preferred option whereas the value chain elements focused on innovative solutions with Fintech or other shared supporting functions, public cloud is a commonly chosen pathway. Some of the value chain areas where a bank will most likely adopt private cloud are as follows:

1. **Core banking system:** The core banking system, processing financial transactions, payments, credit card processing, fund transfers, and loans, storing customer data and performing most critical processes, is most commonly hosted in a private cloud. As mentioned above, for 92 of the top 100 banks, core banking system is hosted by Mainframes.

2. **Institutional Banking:** The institutional banking system performing large scale lending, complex financing, capital market transactions, tax administration and new evolving markets such as ESG and carbon.

3. **Risk Management and Regulatory compliance:** The platforms and applications performing regulatory compliance and risk management functions will be hosted on private cloud platforms.

4. **Treasury Operations:** The trade finance and treasury operations functions including foreign exchange, liquidity management, and cash management.

5. **Blockchain:** Banks exploring private blockchain technology for next generation of lending mechanisms, transferring funds, tracking disbursements and payments and monitoring compliance with covenants, especially when made interoperable with cryptoassets and stablecoins, will be leveraging private cloud.

With regulations such as CPS 234 from APRA making Board responsible for ensuring the enterprise maintains its information security, risk management has become key when the enterprise us deciding on the future of compute. The adoption of private or public cloud differs across the value chain of a bank based on combination of factors such as data sensitivity, regulatory requirements, TCO analysis, performance needs and risk management. As a result, banks have adopted hybrid approach to “Compute” to balance these requirements effectively.

Conclusion

The future of compute for enterprises involves the adoption of cloud services, which offer numerous benefits such as flexibility, standardisation, innovation, security, efficiency, and sustainability. However, there are also challenges and unintended consequences that must be addressed, including culture and operating model issues, security and compliance concerns, economics of cloud, environment and governance impact, and risk management strategies. Kyndryl is well-suited to help enterprises navigate these complexities and realise the full potential of cloud services, with its comprehensive approach to cloud-native services, agile squad delivery models, prebuilt integrations, and a proven Cloud Center of Excellence model. By adopting cloud services and continuously modernising their technology, enterprises can stay up-to-date and competitive in the rapidly evolving digital landscape.

Kyndryl's Relevant Value

Kyndryl is well-suited to help enterprises adopt cloud services due to its comprehensive approach to cloud-native services and its alignment with well-architected frameworks of major cloud providers that build on 30 years of on-prem architectures. Relevant capabilities include:

1. **Sovereign Cloud:** Kyndryl assists organisations leverage benefits of Cloud technologies while mitigating associated risks using our sovereign cloud services. We help with adoption of sovereign cloud, public or private, that operates in a particular country or region and meets a governing body's data privacy and jurisdictional standards. With a sovereign cloud, all data and metadata stay on sovereign soil, preventing other nations from accessing it. Sovereign clouds offer enhanced security measures, including encryption, access controls, zero trust, and network segmentation, which may be tailored to specific countries or regions.
2. **End-to-End Cloud-Native Operating Model:** Kyndryl provides a complete cloud-native operating model that is aligned with AWS and Azure's Well-Architected Frameworks. This ensures that the cloud services are designed and implemented following best practices for security, reliability, performance efficiency, and cost optimization.
3. **Agile Squad Delivery Models:** The use of Agile Squad delivery models allows for flexible and efficient project management, ensuring that cloud services are delivered quickly and can adapt to changing requirements.
4. **Prebuilt Integrations:** Kyndryl offers prebuilt integrations, which can speed up the deployment process and reduce the complexity of integrating cloud services with existing systems.
5. **Kyndryl Service Technology Platform (KSTP):** KSTP delivers consistent, secure, managed hosting of solutions on the cloud, focusing on the delivery of solutions that manage clients effectively.
6. **Cloud Center of Excellence (CCoE):** Kyndryl has a proven Cloud Center of Excellence model that has been deployed in multiple client projects worldwide. This model packages Kyndryl's deep expertise into a collection of services to design, build, and run a CCoE, which is crucial for guiding enterprises through their cloud adoption journey.
7. **Empowering Continuous Modernisation:** Kyndryl's cloud services are designed to empower continuous modernization, enabling enterprises to stay up-to-date with the latest technologies and practices in the cloud domain.
8. **Collaboration with ServiceNow:** Kyndryl works directly with customers to unite people, process, and tools, modernizing and managing mission-critical systems that organizations depend on. This includes integrating ServiceNow's capabilities to enhance IT support and service desk operations.

Confidentiality

This white paper is confidential to you and Kyndryl. It is provided subject to the confidentiality provisions of the existing relevant agreement between us, if any, and on the condition that it may not be disclosed, in whole or in part, to any party other than those of your employees and professional advisors who need to see it to evaluate the proposal (provided that such employees and advisers agree to treat the proposal as confidential) and that it may not be used for any other purpose. Otherwise, disclosures may only take place with Kyndryl's prior written consent.

1. https://www.imf.org/external/pubs/ft/fandd/2021/07/india-stack-financial-access-and-digital-inclusion.htm
2. https://www.federalreserve.gov/publications/files/money-and-payments-20220120.pdf
3. *Top 12 Largest Data Centers Around the Globe (theworldranking.com) lists Azure/Microsoft at 7m sq. ft. Amazon Web Services (AWS) Data Center Locations: Regions and Availability Zones - Dgtl Infra lists AWS at 38m sq. ft.*
4. *Stripe, Inc. - Wikipedia*
5. https://www.pwc.com/us/en/tech-effect/cloud/cloud-business-survey.html - Source: PwC US Cloud Business Survey. June 15, 2021: CIO base of 109, CFO base of 84, Tax leader base of 53, COO base of 67, CRO base of 70, CHRO base of 84 and Board member base of 57
6. https://www.boq.com.au/About-us/media-centre/media-releases/2019/2019-03-071 https://www.webberinsurance.com.au/data-breaches-list
7. https://www.forbes.com/sites/oracle/2018/10/29/4-hidden-costs-of-cloud-infrastructure/?sh=651a83dd2107 https://www.cloud4c.com/insights/blogs/beware-of-5-hidden-costs-of-cloud
8. *On-prem license novation to cloud for IBM Cloud, Azure, and Oracle.*
9. *The Shift Project*
10. *JPM investment in cloud for innovation*
11. *Overview of Vertical Scaling*
12. *GPU platforms*
13. *The role of coupling when moving applications to the cloud.*
14. *Open source overview.*
15. *Consistency/latency trade-offs.*
16. *Federated Search.*
17. *SRE – Site reliability engineers who spend a large portion of their time building and maintain automation programs for the infrastructure they are responsible for.* https://en.wikipedia.org/wiki/Site_reliability_engineering
18. *Moral hazard*
19. https://en.wikipedia.org/wiki/Cloud_management
20. *FinOps overview.*
21. *Links to policies of: Azure, Google, & AWS*
22. *Need a citation about geo-fencing and isolation/bare metal offerings*
23. https://en.wikipedia.org/wiki/Data_at_rest https://en.wikipedia.org/wiki/Data_in_use
24. *Military and National Security requirements often mandate complete isolation from any infrastructure used by another entity. This is the principal motivation behind Secret Cloud and other extremely secure, private cloud implementations.*
25. *Note – while data stored on-prem may be encrypted at rest, it has not been a strong requirement that it be encrypted in motion throughout the private data-centre LAN.*
26. *See the Consistency property of ACID transactions:* https://en.wikipedia.org/wiki/ACID
27. *See* https://en.wikipedia.org/wiki/Blue-green_deployment
28. *See Inner Source* https://en.wikipedia.org/wiki/Inner_source
29. https://www.kyndryl.com/au/en/about-us/news/2023/09/mainframe-modernization-survey-results

This document is current as of the initial date of publication and may be changed by Kyndryl at any time without notice. Not all offerings are available in every country in which Kyndryl operates. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.