

Navigating Today's Computing Model Landscape

Nadeem Ahmad

Abstract

We are witnessing a shift in the computing model landscape. We are increasingly seeing enterprises move away from on-premise hardware deployments while seeing growth in cloud and hyperscale providers. Enterprises are now dealing with a hybrid computing environment encompassing four different computing architectures and multiple cloud-deployment services. This article recaps a recent review by the Innovation and Technology group at NTT Ltd. on the shift in the computing model landscape and the four main computing architectures being deployed within enterprise environments today. This article also covers where NTT Ltd. have helped clients navigate the decision-making process around deploying the various computing models and finally touches on our focus around providing high-value services to continue enabling and supporting clients with their hybrid computing challenges.

Keywords: cloud, application, data analytics

1. Changing dynamics in the market

The market is going through multiple transitions with a strong drive for enterprise clients to consider and adopt cloud computing models. In the datacenter and more traditional datacenter infrastructure space (servers, storage, and datacenter networking), there is no doubt that the new paradigm will consist of a hybrid computing model, incorporating a combination of on-premise computing resources (as per the classic bare metal servers and infrastructure), private cloud platforms (on-premise and/or hosted), and public cloud platforms.

2. Four computing models are challenging the enterprise application landscape

The market is also seeing new computing architectures emerging in quick succession, creating further challenges for enterprises as to which models and architectures to pursue. There are four distinct computing models evolving at the same time: physical server, virtual server, container-based architecture, and serverless computing. This results in enterprises having to support a hybrid computing environment

where the more diverse the mix, the more complex and expensive operations and service continuity become (**Fig. 1**).

In 2019, less than 50% of the physical servers globally have been virtualized, yet we are already seeing a rise in container-based architectures and serverless computing models. Each of these models has a relevant place and fit, and it is highly likely that enterprises will have to consider, plan, and invest in all these architectures. The answer as to which computing model or computing architecture to pursue is found in the application in question.

Many applications cannot be re-platformed to run on virtual machine (VM) architectures and will remain on a bare metal and dedicated infrastructure, most likely for years to come, as these are core business applications that will have to be rewritten to move to more modern architectures, and in many cases this is not worth the return. There are many applications that have been re-platformed to run on VM architectures, and many of the modern cloud platforms originated from these architectures and continue to grow. The next transition is defined by container-based architectures, and most applications will have to be rewritten to fully leverage container-based

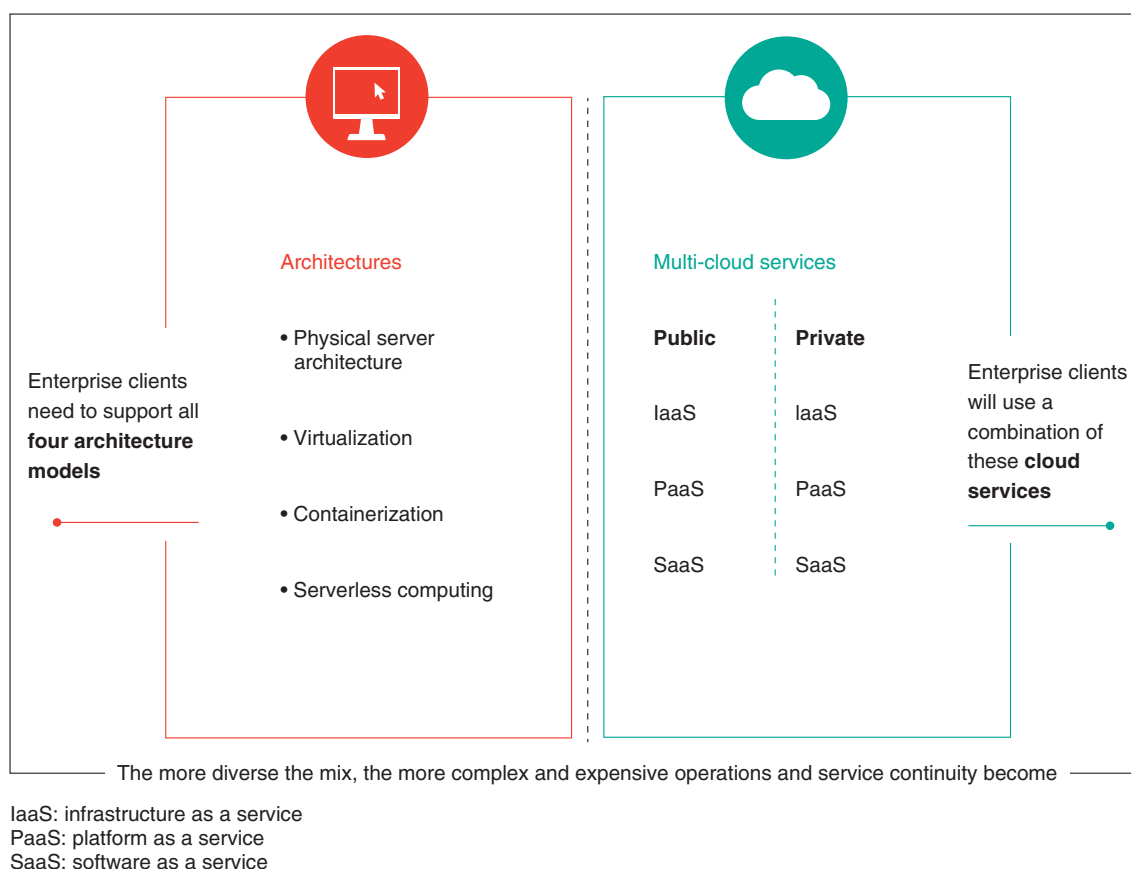


Fig. 1. Hybrid computing environment.

architectures.

At the same time, we are seeing the strong emergence of yet another computing architecture in the form of serverless computing, where value is provided to cloud-based functions, the functions are executed, and value is returned for use. Serverless computing will require further adjustment in computing architecture, especially in the context of data platforms and data abstraction, which implies an entirely different application architecture. Not all applications will be suitable for serverless computing, so the expectation is that certain applications will be suitable or execute optimally on specific computing models and architectures, resulting in a situation in which enterprises will have to invest in all these models and architectures to truly benefit from advances in technology.

This article distinguishes between computing models such as physical/dedicated infrastructure models, cloud-based models, such as infrastructure as a service (IaaS), and computing architectures such as

dedicated server, VM, container, and serverless. It should be noted that hybrids occur within and across both computing models and computing architectures, resulting in a complex service challenge for most enterprises. This complexity is also contributing to the desire for an increasing number of enterprises to leverage applications in a software as a service (SaaS) service type to avoid the complexity of the underlying infrastructure or rewriting applications. This is not always possible, and an increasing complex hybrid is a reality.

3. Where are all the workloads (and money) going?

Morgan Stanley estimates 44% of computing workloads will be orchestrated in the cloud by the end of 2021, up from 21% in 2018 [1]. Meanwhile, every 1 USD of revenue growth for the largest cloud service providers (i.e. hyperscalers) has resulted in about 3 USD of revenue decline for the major legacy non-cloud

infrastructure-technology providers. Data from other industry analysts and views from key market players support this trend. It is clear there is a market shift coming across three critical areas:

- The shift in information technology (IT) spending growth focused on digital transformation
- The change in the mix of enterprise workloads by an environment showing non-cloud shrinking
- The rapid growth and adoption of software-defined infrastructure

There are three clear leaders in the IaaS and platform as a service (PaaS) hyperscale market, Amazon Web Services (AWS), Microsoft, and Google.

- AWS is still the market leader. The rate of innovation and expansion of their services is significant, and part of their success with Amazon has been to compete where needed with razor thin margins, which could be dangerous as the market matures and competitive pressures rise.
- Microsoft has focused almost every part of its business on winning in IaaS and PaaS as well as making the most use of its SaaS offerings (predominantly Office 365) to drive utilization of Azure and its rich partner ecosystem to drive adoption and utilization. Microsoft also have the largest incentive in this group to drive true richness in hybridcloud linking to their partner ecosystem and massive presence and breadth on-premise at the enterprise.
- Google is a relatively late entrant; however, they are investing heavily in enterprise, have significant capacity worldwide, strength in high growth and consumption areas of data and artificial intelligence, and machine learning, and clear ambition to win.

Though hyperscale cloud providers have seen tremendous growth, they are not satisfied and continue to expand the number, variety, and scale of the services that they offer to customers to meet their demand. The number of AWS services rose from 1 in 2006 to over 140 services by 2018. To us, this represents a shift away from traditional servers and buying patterns for some legacy applications and certainly for new initiatives. The applications and scale points to the fact that AWS' (and other hyperscalers') support for migration to cloud continues to expand, supporting larger instances, more dedicated configurations with the same programmability as their heritage, and customer demand. This is reflected in the growth in their revenue.

Today, enterprises are facing challenges within and competition from external and new entrants looking

to disrupt their business and even their entire industry. As a result, digital transformation and speed of innovation is top priority for both business and technology leaders. Increasingly, the line between IaaS and PaaS is blurring, and the speed and efficiency to consume one from the other is driving the shift to hyperscale providers. Cloud applications that are written to take advantage of instant scaling, latest DevOps (development & operations) practices, and toolchains and containers are often public cloud first workloads, skipping the enterprise datacenter all together. While there is a risk of provider lock-in for enterprises, the speed and richness of solutions is incredibly attractive. The stickiness with a single hyperscale provider is high for a solution given that many PaaS components are proprietary and designed to work 'better-together' with other parts of their own ecosystem and partners.

4. The struggle to maintain relevance is real: enter programmability

Considering the complexity and scale of the computing models and architectures and the adjacent infrastructures such as connectivity, the modern workplace, and diverse business models, it goes without saying that there is a major challenge to operate and service these complex environments and remain relevant in the industry, especially in the light of the speed of technology transitions. The industry does not produce enough skilled human resources to perform the operational services to maintain these environments. The problem is multi-faceted in that there is a shortage of skilled human resources, resulting in these environments becoming so complex and dynamic that humans can no longer perform the operations effectively or efficiently. All operational models are now underpinned by a large degree of automation, which in turn requires all these infrastructure and architectural platforms to be fully programmable.

Many infrastructure vendors have been redeveloping their products to be programmable, and one must consider the life cycle of the installed base, legacy of multiple platforms, operating systems, product families, and lack of modern software architectures. There is an increasing demand for all vendor products to be 'API (application programming interface)-first,' implying that products need to be programmable first and foremost and console or command line configurability secondary, if at all.

Companies that have a true API-first model will

most often have their console access via the very same APIs, whereas companies that have to transition from a console or command-line-first approach to an API have to do so via complicated stages. Typically, one finds that not all the configurational items are accessible via their APIs, which results in a complex hybrid model. The latter is not desirable as it requires ‘old’ and ‘new’ skills and does not support the automation that operational model demand, and all of this becomes costly and lethargic. Considering the current market and key technology vendors, it goes without saying that those who transition to API-first and truly programmable models will flourish, be preferred in the market, and accelerate their growth, and those who do not move swiftly enough will see market share losses and face severe market challenges.

As information and communication technology (ICT) providers consider their objective to be leading IT and services companies for their clients, they will need, if they have not already, to establish a preference for vendor partners that offer API-first technology that is fully programmable. At the same time, they will need to proactively remove technology products from their portfolio that do not provide programmability, since there is no longer a sustainable business model for non-programmable infrastructure for ICT providers or their clients.

5. With no end in sight—maturity and expansion of SaaS continues

The mass adoption of enterprise SaaS continues across a wide range of application categories, and each of these applications previously required servers and storage inside the enterprise and are now consumed as a service. There are SaaS providers in the following major categories:

- Productivity: Office 365, Google G Suite, Atlassian Confluence and JIRA
- File store: OneDrive for Business, Dropbox, Box, Google Drive
- Collaboration: WebEx, Zoom, High Five, Slack, MS Teams, Yammer
- Enterprise resource planning (ERP): SAP Hana & cloud variants, Oracle, Dynamics 365, Intacct, Intuit, NetSuite
- People and payroll: Workday, ADP
- Service/sales/customer relationship management (CRM): ServiceNow, Salesforce.com, Siebel

PaaS and serverless computing are poised to make further inroads into reducing the requirement for

enterprise clients to buy servers, storage, and networking infrastructure, not to mention the consulting, professional, and technical services that went into their design, build, and operation.

6. Hybrid computing environment in practice

This concept of four computing architectures is not just theory or an academic exercise. At NTT Ltd., we have practical use cases of implementing a hybrid mix of architectures. However, decision-making for clients is not always centered around the best technical approach. Often, the factors that go into choosing to implement one architecture over another or a mix of architectures are a matter of operational challenges, financial discipline, and risk tolerance. Let us explore one of our practical examples of our multi-architectural approach to meet the needs of one of our more well-known clients.

7. Implementation of four computing architectures for Tour de France

NTT Ltd.’s Advanced Technology Group focus is on using sports as a technology incubator within NTT. They are focusing on the strategic innovation projects that NTT are executing with clients. The Amaury Sports Organization (A.S.O.) organizes the Tour de France, a bicycle race that takes place over 21 days around France involving riders traveling about 200 kilometers every day. Supporting the race is what is commonly referred to as a traveling circus as a large number of support and infrastructure vehicles are required. When we started and first partnered with the A.S.O., there was a real concern around the level of connectivity, e.g., just basic supply of items such as electricity, in many of the remote locations the tour visits.

Our first solution was using an enormous truck, i.e., a traveling datacenter. The truck traveled hundreds of kilometers each day and night to support the race. With this truck, we collected the data from the bicycles and riders, provided a large amount of analytics, and provided the summary and result of data analytics to the broadcasters. Our truck would pull up next to the broadcasters’ trucks, and a cable would be plugged in from one truck to the next. We had a physical infrastructure, i.e., physical servers in the truck. This worked, but it was quite challenging as one had to take all the technical staff on the road to ensure everything worked smoothly.

In the 2016 race, we had a stage of the race located

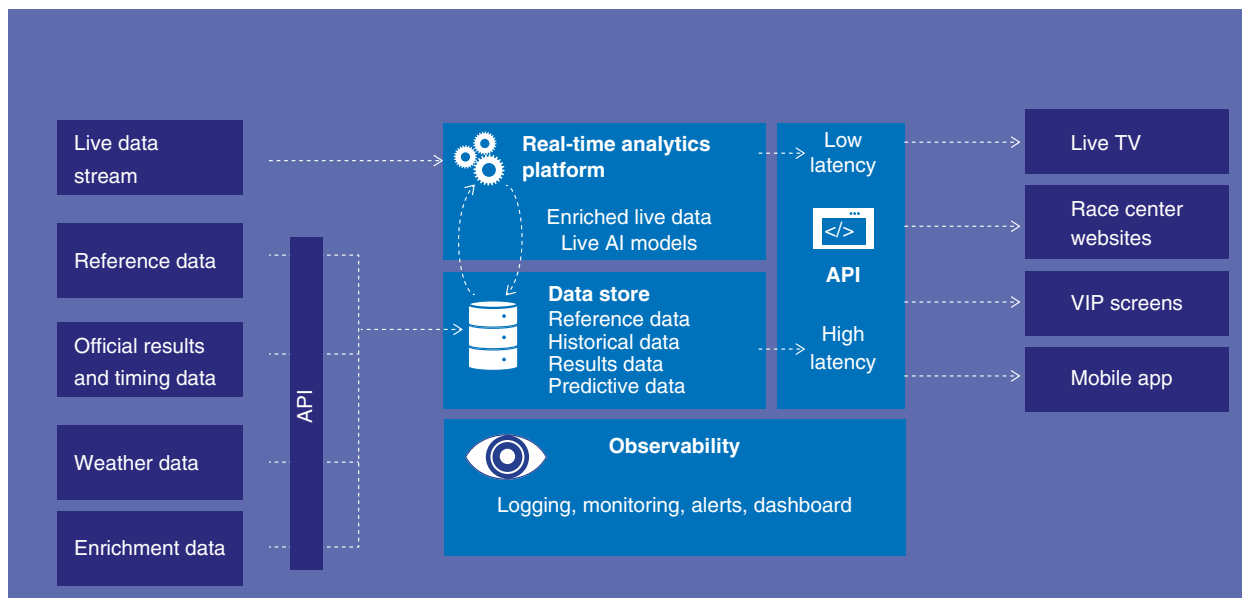


Fig. 2. Data platform.

at a mountain called Mont Ventoux, which is known for having extremely windy conditions and severe storms. During that stage, a severe storm came through. It was so bad that the finish line, which was meant to be at the ski resort at the top of the mountain, had to be changed, and they finished the race only halfway up the mountain.

This does not seem like a significant operational challenge, but all support trucks that were going to be parked in the large parking lot up at the ski resort were distributed at schools, fields, and supermarket parking lots across several villages at the bottom of the mountain, and our truck was no longer parked next to the broadcasters' trucks. In fact, it was two and a half kilometers away.

Therefore, the decision was made at around two o'clock in the morning that we basically had to take everything that we have been running on physical servers and shift to running on VMs in the NTT cloud overnight, and we were actually up and running for the next day. We ran successfully on that model for the rest of the race, and that convinced everyone including our client that the applications and functions running on physical servers could run in a virtual world in a virtualized environment, and we stayed with that virtualized environment until 2018.

Over the last eight months or so, we have been making some very fundamental architectural changes, which are in line with what we have been talking

about thus far, i.e., the four computing models. If we look at **Fig. 2**, we see a representation of our data structure, which is the core of our solution. Very broadly, we obtain data from a range of sources on the left side, carry out some data processing and analytics in the middle, and distribute data to a variety of consumers on the right side. The real-time data analytics platform is grouped into four main areas because these areas were essentially built as different platforms and have different drivers in each area.

As an example, we need a robustly managed API. However, we had a hard look at our previously built API and through the course of last year, as much as everyone prefers to do it themselves, we realized that there was actually nothing about our API that was different from anyone else's. There was no additional benefit to building it ourselves, and what we could bring would be only additional cost, complexity, and the management burden across the platform and solution.

Therefore, we ended up going with a set of PaaS services and serverless functions for the API. We used a third party's API, which has the advantage in that we do not need to worry about any of the infrastructure or mechanics of the API infrastructure and simply plug in our functions and transformations we want and structure the data that we want to process.

This has significantly narrowed the scope, focus, and effort of that workstream. It also means that the

support requirement on our end significantly decreases. We do not need to worry about servers. We were previously running an API that can scale to support a million users and requires a very significant infrastructure and large amount of looking after, but using a third party's API basically bypassed that whole lot. The last race of the Tour de France in which we used such APIs, was uneventful for the people looking after the API infrastructure because it just worked. At the end of the day, this is the desired result for all IT operational support people both on the client and provider sides.

The real learning has been trying to understand the specific features of the hyperscalers' solution. One needs to get into such specifics, and we found the best way of doing this is running proofs of concept and proving out the technologies because one needs to not only make decisions in terms of technology but also in terms of costing models.

Previously, the approach was just about obtaining the right technologies, then we would run proofs of concept and end up with a choice of technology that supports the desired functionality, however, the cost of the technology solution was quite high due to our large amount of data. Today, with most hyperscalers, if there are multiple overlapping options in many areas, we can switch to another option that is functionally just as effective and have a far better pricing model. This is something to consider, particularly when one starts looking up at the higher end of the stack of solutions from the hyperscalers.

Thus far, we have seen operational challenges move our architecture from physical to virtual. We then saw that financial considerations cause a need to adopt serverless and "as a service" models from hyperscalers. Technologies are not always the deciding factor into which a computing model will work for our client environments.

Looking at Fig. 2 again, the data store is primarily an SQL (structured query language) server; it is just a data store. It has a large amount of data, and we have made the choice, in this case, to stay with the VM-based solution not because there are no other alternatives or because it is the best option. It is just that in our case for 2020, we have several other parts of the solution changing, so we did not want to also change our data store architecture. This represents another item for organizations to consider and is their tolerance of risk. We have deliberately chosen not to do anything with that data store and chose to leave it on a VM because we have enough technology risk elsewhere in the stack and do not need to change it. We

will probably change it next year because there are some advantages to shifting to PaaS solutions, but it is not an imperative this year. We now have two computing models incorporated into the environment, i.e., we have serverless functions on the API and VMs on the data store.

The last area is the real-time analytics platform, which is where NTT Ltd. brings a large amount of its intellectual property into the services that we provide to the A.S.O. and broadcasters. It is where we take the sensor data from the bicycles and enrich the speed and location data we obtain from the bikes and apply machine learning to turn a couple of data points into over fifty data points that then support a range of consumers and consumption models. Therefore, as opposed to the above third party's API, it made sense for us to build our own data analytics platform.

However, building our own real-time analytics platform is complex, and we want to be able to continue to implement changes and rapidly evolve that solution. In this case, therefore, we have gone with a containerized solution that is built across a range of Docker containers, which makes deployments much easier. One might ask, "How does running on containers make things easier?" It is because one does not need to deploy an entire VM, operating system, plus everything else. One just deploys the container so what one needs to run on is much smaller and simpler in terms of footprint and resources. We can describe this in code, which means we can say we want a container that is this big and looks like this and has these items in it much more easily in code. This means we can make a deployment directly from code. The code for the infrastructure can be managed, just like application code, and run just like the rest of the application. In fact, it becomes a single package—the container and everything that is in it. We can then make deployments multiple times a day, which means that we can make changes and test the changes rapidly, giving us the velocity that we really want around the development in that space. The above gave a bit of context and insight in the three computing models being implemented. Remember, we moved away from an entirely physical infrastructure to VMs, serverless computing, and container-based architectures.

There is one last point around the complexity of having all three architectures in play. Referring to Fig. 2 once again, there is a small box at the bottom called "Observability." Running hybrid architecture solutions like this introduces much complexity. The challenge is that if one changes one piece of the

puzzle, another piece somewhere else can entirely break or malfunction, and it can become too complicated to understand all the ways the system can fail. While one can put in place alerts, one can be guaranteed that what will break is not something one has placed an alert on.

What we have built across the whole solution is an observability platform for visualizing the outputs of all the layers then all the functions of the layers. This allows us to monitor the outputs and understand what functions are working or not working. Complexity is definitely a challenge and is the reason we have implemented that fourth observability component in our solution to address and manage that challenge.

The Tour de France platform we built provides an excellent example and provides interesting insights into the factors considered when implementing a hybrid mix of computing architectures. We started out with a traditional physical infrastructure then, due to operational challenges, we had to pivot to VM implementation. Due to financial considerations, we incorporated serverless computing, and further operational, portability, and efficiency considerations drove the move to a container-based architecture.

The lesson here is not always a matter of the best technology but the best fit of technology driven by business considerations, aside from the technology itself. In the Tour de France example, the team had to address considerations such as operational challenges, financial considerations, security, and of course risk tolerance.

8. Summary and high-value services

Unlike physical computing, virtual and container-based computing make it easy to move applications, as they are decoupled from the underlying computing infrastructure. The use cases for these models are becoming clearer and adoption is on the rise, and there is now the promise of serverless computing, which offers even greater agility and cost savings because applications do not have to be deployed on a server. Instead, functions run from a cloud provider's platform, return outputs, and immediately release the associated resources.

We have four distinct computing models evolving simultaneously: physical server, virtual server, container-based architecture, and serverless computing. This raises interesting questions for organizations: What do we do with our applications? Which workloads should move to the cloud? What must stay on-premise? Can I re-platform some applications to run

on VMs? Can I rewrite others for the container or serverless world? Most importantly, how long will it take to get this done so I can unlock the benefits of these new models? The computing-architecture decision has a direct relationship with the future of each application, in fact, it may even dictate the future direction of each application.

This article reviewed some of these decisions and considerations during a practical client implementation of all four computing architectures where factors such as operational challenges, financial considerations, portability, efficiency, and risk tolerance all played a role in making the decision to deploy each of the four computing architecture. ICT service providers, including NTT Ltd., are adapting their go-to-market and partner strategies given the rapid growth of hybrid and hyperscale deployments. Hyperscale providers have such a breadth of service offerings that staying up to date across multiple clouds will require a massive investment on the part of enterprise clients as well as some complexity. In this regard, ICT service providers can provide true value and guidance to their clients.

What is also evident is the need for ICT providers to drive high-value service offers to enterprise clients to guide and support them on the journey ahead. This is why NTT Ltd. is now focused on selling and architecting over forty high-value services. Some relevant services include:

- **Infrastructure consulting** services encompassing several services focused on optimization and transformation wherever our clients are in their transformation journey. Whether it is resolving ageing infrastructure, managing complexity, securing infrastructure, moving to a multi-cloud environment, or improving user experience, we can help our clients evaluate their current operating state and transform and improve the performance of their infrastructure while optimizing their spending.
- **Cloud consulting services** including a hybrid migration assessment to obtain clear insights into a client's current infrastructure and applications to create a roadmap to managed services on a hybrid cloud. Cloud consulting services also include cloud migration services, which are end-to-end services beginning with auditing workloads for suitability for migration, planning the migration to the best-fit venue, then migrating these workloads to the cloud without disrupting daily operations and business.
- **Managed Hybrid Infrastructure Services**

offering comprehensive management and monitoring of on-premise, cloud, and hybrid IT infrastructures to improve operational agility, reduce risk, and optimize cloud and technology infrastructure investments. With our highly trained vendor-accredited people, processes, and tools necessary to manage the day-to-day operations of our clients' multi-platform multi-site IT infrastructures, we ensure the resilience and performance of such infrastructures so one can focus on achieving business outcomes.

- **Managed application services** around the implementation and management of ERP and CRM applications on premise, in the cloud, or in hybrid IT operating environments with consistent governance, security, and management practices.
- Finally, our **Hybrid and Multicloud service** allowing our clients to use the best platform for each workload and taking away the difficulty of managing a hybrid IT environment. With systems spread out across multiple environments, our client's need a way to manage all these environments from a single interface. With our 'Solution Insight' cloud management platform, we enable our clients to gain insight into how all their systems work together. This helps them bet-

ter understand which applications can be moved to the cloud and which need to stay in their current environments.

This short list is not an all-inclusive one, as we have several other high-value services we are emphasizing in our engagements with clients across digital transformation, workplace, datacenter, networking, and security.

The shift in the computing landscape is upon us as enterprise clients need to support four different architecture models at the same time. Enterprises will use a combination of cloud services across public and private clouds. The more diverse the mix, the more complex and expensive operations and service continuity become. Through our high-value service offers, we can help enterprise clients navigate decisions and considerations when embarking on a hybrid computing-environment-implementation journey.

Reference

- [1] Morgan Stanley, <https://www.morganstanley.com/ideas/it-hardware-2018>

Trademark notes

All brand, product, and company/organization names that appear in this article are trademarks or registered trademarks of their respective owners.



Nadeem Ahmad

Group Senior Vice President, Technology Strategy, NTT Ltd.

Nadeem Ahmad is a technology professional with 25 years of ICT industry experience and is a social constructivist at heart. He believes technology does not determine human action but that human action shapes technology. He joined NTT Ltd. (previously Dimension Data) over twenty years ago initially leading technical delivery and operations teams within the United Kingdom, Europe, and Americas but has spent the bulk of his time in the Group CTO Office (now the Innovation and Technology group) focused on business and technology innovation initiatives across the Group.

His role at the Innovation and Technology group involves defining NTT Ltd.'s technology strategy and positioning, in collaboration with go-to-market practice leaders, across relevant technology domains. This includes identifying and monitoring key technology trends and emerging domains, assessing their impact on enterprise clients and their potential to affect new solutions or services, as well as providing input into new offer creation. This extends to providing group guidance on emerging technology domains and translating key trends and market developments into business impact via thought leadership, internally and externally.

LinkedIn profile: <https://www.linkedin.com/in/nadeemahmad1/>