



Technology Brief

Today's Computing Model Landscape

NTT Ltd.

27 April 2020 | Document Version 1.0

Nadeem Ahmad

+1215607971

nadeem.ahmad@global.ntt

Table of Contents

1.	Executive summary	3
2.	Changing dynamics in the market	4
2.1.	Four computing models are challenging the enterprise application landscape	4
2.2.	Where are all the workloads (and money) going?	5
2.3.	The struggle to maintain relevance is real: enter programmability	6
2.4.	With no end in sight – maturity and expansion of SaaS continues	7
2.5.	Rise of the ['white-box'] machines – the real threat to OEMs	7
3.	Game plan for ICT service providers	9
3.1.	Button-up that hybrid cloud strategy	9
3.2.	Get moving on developing in-house skills	9
3.3.	Summary and wrap-up	10
3.4.	Document references	10

List of Figures

Figure 1: Skills evolution.....	9
---------------------------------	---

1. Executive summary

Increasingly, within the industry, we are seeing an interest by enterprises towards moving away from on-premise hardware deployments and at the same time the growth of cloud and hyperscale providers. This will result in a significant impact on traditional OEM partners, and in correlation, an ICT provider's business. Even one of the top strategic partners, as the #1 server maker in the enterprise market, forecasts enterprise spend on cloud will rise to 40% of their infrastructure budget within the next 3 years.

This brief is meant to provide a review of the shift in computing models and architecture occurring in the market and provide some perspective as to how ICT providers' go-to-market and partner strategy will be impacted across our areas in terms of its relevance within hybrid or hyperscale deployments. In addition, this brief talks to the urgency with which ICT providers need to drive deep skills development for programmable infrastructure of all types in preparation for a hybrid cloud world. In multicloud, ICT providers can help their clients balance the risk of vendor lock-in while providing hard to get skills and scale them across vendors.

Hyperscale providers have such breadth of offering that staying up to date across multiple clouds will take a massive investment and has some complexity – in that space we can provide true value to our clients.

The market:

- The market is dealing with FOUR computing architectures currently: 1) the traditional physical server architecture, 2) virtualization, 3) containerization and 4) serverless computing. Enterprise clients are likely to have to support all four models over time. The longevity of these architectures depends on applications, and the financial and operational constraints around applications will require for these architectures to remain for quite some time.
- The market is also dealing with multicloud services in the form of multiple service types, such as infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS) and software-as-a-service (SaaS), as well as at the public and private level. It is highly likely that most enterprise clients will use a combination of these cloud services, resulting in a multicloud environment.
- The result is a hybrid computing environment, consisting of all four computing architectures and multicloud environments, and the more diverse, the more complex and expensive the operations and service continuity.
- Whether organically or through acquisition, we are also seeing vendor companies expand their product portfolio to cover the software-defined data center, software-defined workspace and connecting these with a software-defined WAN fabric. These organizations are also embracing the hyperscale cloud providers, including Amazon Web Services (AWS), Azure, Google Cloud Platform (GCP) and Alibaba. Some of them have even launched joint services with AWS continue to extend these offerings.
- Some vendors are going further in this software-defined world and are, or are in the process of, re-engineering their product offerings to be fully programmable (API-first in the true sense), over the last couple of years as they re-architect and re-code their offerings. There is strong competition in the race to have the first true software-defined suite that extends from the edge to the full computing hybrid. The result will be a product portfolio that enables clients to simplify operations, service delivery and cost. This will help to consolidate what has become a very cluttered market with regards to a full computing architecture and application transformation platform to enterprises.

2. Changing dynamics in the market

2.1. Four computing models are challenging the enterprise application landscape

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

The market is also seeing new computing architectures emerging in quick succession, creating further challenges for enterprises in which models and architectures to pursue. 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 architectures, and will remain on bare metal and dedicated infrastructure, most likely for years to come, as these are core business applications that will have to be rewritten in order 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 in virtual machine 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 architectures.

At the same time, we are seeing the strong emergence of yet another computing architecture in the form of serverless computing, where values are provided to cloud-based functions, the function executes, and values are returned for use. Serverless computing will require a 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, and so the expectation is that certain applications will be suitable or execute optimally on specific computing models and architectures, resulting in a situation where enterprises will have to invest in all these models and architectures in order to truly benefit from the advances in technology.

This brief distinguishes between computing models such as physical/dedicated infrastructure models, cloud-based models such as IaaS and computing architectures such as dedicated server, virtual machine, container and serverless architectures. It should be noted that hybrids happen within and across both computing models and computing architectures, resulting in a complex and challenging service challenge to most enterprises. This complexity is also contributing to the desire for more and more enterprises to leverage applications in a SaaS service type, to frankly avoid the complexity of the underlying infrastructure or rewriting applications. This is not always possible, and an increasing complex hybrid is a reality.

2.2. 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 2018¹. Meanwhile, every USD 1 of revenue growth for the largest cloud service providers (i.e. hyperscalers) has resulted in about USD 3 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 IT spending growth to be focused on digital transformation
- the change in mix of enterprise workloads by environment showing non-cloud shrinking
- the rapid growth and adoption of software-defined infrastructure

There are three clear leaders in the IaaS and PaaS hyperscale market, 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 both its SaaS offerings (predominantly Office 365) to drive utilization of Azure and its rich partner ecosystem to drive adoption and utilization. Of note Microsoft also have the largest incentive in this group to drive true richness in hybrid cloud linking cloud to their partner ecosystem and massive presence and breadth on-premise in 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.

Further, though hyperscale cloud providers have seen tremendous growth, they are not satisfied, and they continue to expand the number, variety and scale of the services that they offer to customers to meet their demand.

As an example, 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, but 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 customers demand. This is reflected in the growth in their revenue.

Today, enterprises are facing challenges within their organization 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 of mind 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 practices and toolchains and containers, are often public cloud first workloads, skipping the enterprise datacenter all together.

While there is 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 many PaaS components are proprietary and designed to work 'better-together' with other parts of their own ecosystem and their partners.

2.3. **The struggle to maintain relevance is real: enter programmability**

Considering the complexity and scale of the computing models and computing architectures, and the adjacent infrastructure such as the 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 just does not produce enough skilled resources to perform the operational services to maintain these environments. The problem is multi-faceted in that not only is there a shortage of skilled resources, these environments have become 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 on a journey to redevelop their products to be programmable, and this is a journey, as one must consider the life cycle of the installed base, the 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-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 approach, have to do so via complicated stages and typically one finds that not all the configurational items are accessible via their API, 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 models demand, and all of this becomes costly and lethargic.

Considering the current market, and the 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 ICT providers consider their objective to be a leading IT Technology and Services company for their clients, they will need, if they have not already, to establish a preference for vendor partners that offer API-first technology and that is fully programmable. At the same time, they will need to proactively remove technology products from their portfolio that do not provide programmability, as quite simply there is no longer a sustainable business model for non-programmable infrastructure, for the ICT providers or their clients.

2.4. With no end in sight – maturity and expansion of SaaS continues

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

- 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
- ERP: SAP Hana & cloud variants, Oracle, Dynamics 365, Intacct, Intuit, NetSuite
- People and payroll: Workday, ADP
- Service/Sales/CRM: ServiceNow, Salesforce.com, Siebel

Platform as a Service (PaaS) and serverless 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.

2.5. Rise of the ['white-box'] machines – the real threat to OEMs

The ill-fated reference to the old Terminator 3 movie aside, one final and important consideration, and maybe one of the greatest threats to traditional OEM business, is the rise of generic 'white-box' routers and switches. These unbranded devices generally have less powerful hardware than traditional OEM products, but the heavy lifting is accomplished via cloud-based software. This disruptive approach is known as software-defined networking (SDN).

Traditional OEM's core business model relies on the notion that their customers should buy all their networking or storage hardware from a single vendor to ensure that everything runs smoothly. They will sell large quantities of routers and switches, which are often bundled with additional software to 'lock-in' their big customers. However, many organizations now realize that white-box solutions usually work just as well as the traditional OEM hardware.

For example, AT&T has previously tested a high-speed network entirely comprised of white-box hardware. The test also successfully sent data between white-box hardware from two different vendors running on two different types of chips, which proved that it was not necessary for companies to buy all their networking solutions from a single OEM vendor. AT&T's white-box network ran on its own open-source network-management software, ECOMP. AT&T gave ECOMP to the Linux Foundation, which allows anyone, including rival Telco's, to use or modify the software.

AT&T subsequently announced that it would deploy 60,000 white-box routers across its network, which would cut its costs and boost the capacity of its networks as it gears up for the 5G transition. The revelation that AT&T is developing its own switches should not surprise anyone. AT&T's white-box transition enabled it to reduce its spending on traditional OEM networking solutions from USD 2 billion in 2013 to just USD 400 million in 2017. That figure could eventually fall to zero as

AT&T rolls out its white-box routers and first-party switches. AT&T is not the only major company using white-box hardware.

Facebook uses white-box hardware in its massive datacenters and has invented new switches that can replace traditional OEM hardware. Like AT&T, Facebook is giving away its designs for free through its Open Compute Project.

Other big service providers and enterprise customers likely will follow AT&T's and Facebook's lead soon. When they do, traditional OEM routing and switching businesses, which are major parts of the business for these companies, could be severely impacted.

In addition, last year Amazon announced its intent to offer its own flavor of networking solutions to market, the market reaction was immediate and saw a 3% drop in the largest OEM networking provider's stock shortly after the announcement. It also potentially signals risk that these solutions could be headed down market.

AT&T and massive telecommunications providers (including NTT, parent company of NTT Ltd.) have long had unique scale properties of their own, giving them the ability to focus R&D investment in areas like this. The rise in ODM and maturing of open-source software market has created a dynamic where they can bring their vision of independence from traditional OEM vendors to fruition.

However, while this has long been out of reach for other businesses, Amazon has the capability both from a sales and logistics as well as a technical perspective to truly disrupt in this space bringing 'white-box' style solutions to more of the market. Just consider its dual focus in penetrating the business market for product sales (they already have Amazon for Business market offering) as well as its own cloud scale, assets, ODMs and R&D driving innovation and independence with 'white-box' solutions.

Consider the scenario where someone in an enterprise simply says, 'Hey Alexa, I need another switch and wireless access for 30 people in our new Boston office.'

3. Game plan for ICT service providers

3.1. Button-up that hybrid cloud strategy

The market shift and resulting computing models will have a critical impact to ICT service providers' hybrid cloud strategy and how they execute on that strategy using leading edge tools and platforms to support their client outcomes. It will be important for these service providers to define its service offerings around these architectures and ensure the right skills exist, or are added, to drive this to market. At the same time, a strong requirement for support and managed services would be necessary and will need to be developed. These services will have to focus on software and platform delivered service provisioning, service delivery and service operations across the defined hybrid cloud and compute footprint. Technical services to deal with workload migration will be equally important, with a focus on automated provisioning and operations.

3.2. Get moving on developing in-house skills

Innovation and evolution of computing models and architectures also comes with changes to skills focus. ICT service providers need to move from legacy deployment and operational skills towards software-defined skills as Figure 1 below depicts. These services led businesses need to update their job frameworks, learning and career development programs ahead of the transition to newer technology delivery. And make no mistake, acquiring new talent is difficult in this market. Companies need to help their existing people in transitioning to these new skills.



Figure 1: Skills evolution

The question arises as to how an existing employee base can attain these skills. We are seeing the way companies engage with vendors for training and enablement changing.

Leading tools/platform provider in the software-defined space do not have formal partner training programs in place. Online learning platforms, such as PluralSight, are becoming a more attractive alternative. Industry experts develop content for subscribers to consume and external experts are developing content which we need to source ourselves. This requires investment dollars devoted to these training platforms.

HR organizations have already utilized PluralSight for different training initiatives within the departments and/or divisions. In addition, other internal company learning platforms enable internal (and external) subject matter experts to provide content and education through learning paths. Specific or initial emphasis should

be placed on software-defined networking courses covering topics such as Python or Ansible. The skills path to applying python to network automation is key as is PowerShell automation – writing scripts for deployment and automation.

A good first step is to identify personas. These are groups that staff members can be categorized within (as best as possible). For example, there is the DC Networking persona. Within it, is a network administrator named Jenny. Now we must identify what should Jenny learn to adapt to the evolution in the networking space. For this example, we would determine that Jenny should learn the fundamentals of software-defined networking, like OpenFlow, and how container architecture will change load balancers and firewall strategies. Ideally, Jenny's company would provide her the training, including perhaps access to a lab to test/engineer new elements in container-networking, on how to automate aspects of her work to ensure reliable outcomes for clients.

These are some examples citing certain relevant technologies, but a proper and specific methodology for developing the personas and then a curriculum for each persona must be developed and must be mandated as part of a skills transformation exercise across the department or division.

3.3. Summary and wrap-up

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

We have four distinct computing models evolving at the same time: physical servers, virtual servers, container-based, and serverless computing. This raises some big questions for organizations: What do we do with our applications? Which workloads should move the cloud? What must stay on-premises? Can I re-platform some applications to run on virtual machines? 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.

ICT service providers' go-to-market and partner strategies will be impacted across various service areas in terms of their relevance within the rapid growth of hybrid or hyperscale deployments. Hyperscale providers have such a breadth of service offerings that staying up to date across multiple clouds will take a massive investment on the part of enterprise clients and has some complexity. In that space, ICT service providers can provide true value to their clients.

What is also evident, is the urgency to drive deep skills development for programable infrastructure of all types in preparation for a hybrid cloud world. In multicloud, there is the opportunity to help clients balance the risk of vendor lock-in while providing hard to get skills and scale them across these vendors.

3.4. Document references

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