HYBRID IT HELPS BUSINESSES NAVIGATE THROUGH DIGITAL TRANSFORMATION

EXECUTIVE SUMMARY

In this era of digital disruption, businesses must be more agile to capture opportunities. Many viewed cloud computing technology as the way to do this, promising to address agility, scalability, and cost. But in moving to the cloud, many found that its security, compliance, and performance did not fully meet their needs. Additionally, previous common thought was public cloud is less expensive than private cloud. We now know that is not true in all cases. Savvy businesses realize **hybrid IT**, which includes both off-premises and on-premises services, enables better agility. After initial experience with public cloud offerings, businesses learned that many workloads are best hosted on-premises, primarily due to security, compliance, performance, control, and cost issues.

To capitalize on digital transformation opportunities, businesses require a hybrid IT environment that is software-defined, flexible, scalable, and composable, whether applications and data are in the datacenter, private and public clouds, or network edge, with the flexibility to operate seamlessly across these environments. As businesses build hybrid IT environments, they need a partner who can help them architect the right mix of private cloud, public cloud, and traditional IT to meet their agility, security, performance, control, and economic requirements, so they can take advantage of digital transformation opportunities.

TODAY'S CUSTOMER NEEDS

IT must change to drive digital transformation, accelerate innovation, and keep pace with an increasingly competitive market. The need for this digital transformation will be continuous, not simply a point in time. Key drivers for enabling a digital transformation include agility, scalability, and lower cost—benefits typically associated with public cloud. But IT has realized that other variables like performance, security, control and compliance can be disqualifiers for hosting an application outside of the datacenter. As customers work to find their right mix, a sharper focus on cost is forcing a deeper discussion about the true impact of off-premises versus on-premises.

In traditional enterprises, the core set of apps, data, and processes has been a static "processing engine". In the new era of digital transformation, enterprises must evolve this static environment to become a dynamic "digital core". The convergence of core transactions and analytics can help achieve the goal of becoming a real-time business

with the ability to make decisions based on immediate insights and take action in the business transaction streams.

To respond to the new way the business works, IT needs the power to compose and recompose resources efficiently and respond to new requests within minutes, providing fluid resource pools that users can access. IT must be able to predict future needs and allocate resources without the users even knowing that they need it. And this process must be automated, freeing up IT specialists for other tasks.

For many enterprises, their early cloud experience was in test/dev environments that needed fast deployment and had less criticality than production workloads. Public cloud providers like Amazon Web Services (AWS) and Microsoft Azure offer IT services that can be easily and quickly accessed. But production workloads often need better security, compliance, performance, control, and cost—all things that can be better met with on-premises services. A recent study highlights the expected growth of private cloud, consistent with our belief that the best approach for most enterprises is to use a mix of traditional IT, private cloud, and public cloud for a more balanced approach between cost, agility, and control.

FIGURE 1: THE DEPLOYMENT CONTINUUM

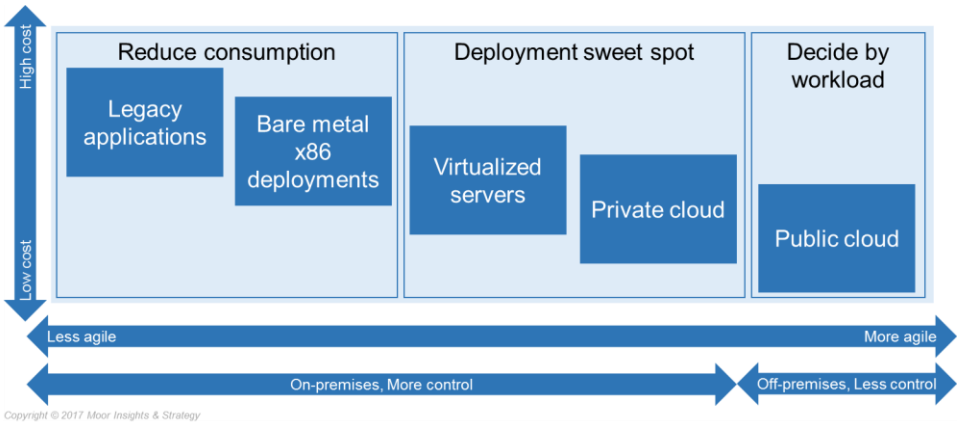

Source: Moor Insights & Strategy

With large established investments in existing architecture and people, it is unrealistic to consider a wholesale public cloud migration. Businesses focus on the easiest workloads to migrate, while the more difficult workloads will continue to run on existing platforms.

"Our direction is to move to the cloud to a certain point, but not completely."
– Large Enterprise, Los Angeles[1]

[1] All customer quotations are from focus groups with both HPE and non-HPE customers.

HYBRID IT BRINGS AGILITY TO THE DATACENTER

A savvy investor will choose a diversified portfolio of stocks, bonds, and other financial instruments; they would never put all their eggs in one basket. Businesses should consider a hybrid IT approach that ties deployment strategies to usage, through a combination of public and on-premises. Hybrid IT enables a business to add agility while still maintaining control over workloads.

FIGURE 2: HYBRID CLOUD VS. HYBRID IT

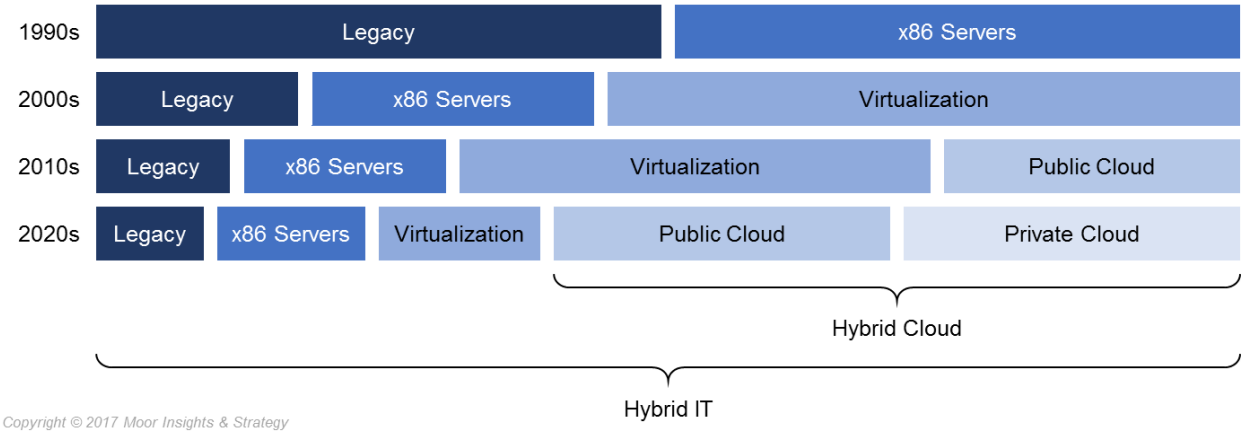

Source: Moor Insights & Strategy

Applications born in the cloud such as Office 365, Salesforce, and Workday make most sense to run as SaaS instead of on-premises. But that is not to say that every cloud-native workload must run in a public cloud. Newer server technologies like converged, hyperconverged, and composable platforms make it easier to run cloud-native applications in containers in a datacenter, something that may not have been considered just a few years ago.

Enterprises need to evaluate all workloads based on security, compliance, performance (high throughput and low latency), control, and cost to determine the best location. In this agile world, cloud deployments will be hybrid (public and private), and the IT environment will be hybrid as well (traditional and cloud).

> "*We're moving toward the cloud on-premises first, which will be stepping stone to public cloud. We'll use the transition step until risk numbers are more acceptable. Long term will be hybrid private and public cloud.*"
>
> – Large Enterprise, New York

THE NEW CLOUD REALITY

A subset of the broader hybrid IT discussion is specific to public versus private cloud. Initially, software as a service (SaaS) like Salesforce.com, Office 365, and Workday provided a significant savings over hosting those applications locally, primarily due to licensing savings. Conventional wisdom about off-premises changed, enabling businesses to take the next step—platform as a service (PaaS)—for new cloud-native applications. PaaS enabled developers to have access to more tools than IT was traditionally providing, leading many to posit that almost everything would move to the public cloud.

While the benefits of cloud were clear, limitations arose that made public cloud infeasible for many workloads. Some applications are better suited to run in public cloud, like web services, whereas others are better suited to run on-premises, such as databases with analytics that contain customer data. "Put everything in the cloud" was soon replaced by "Move the right workloads to the right environment", leading to more on-premises deployments than originally expected.

FIGURE 3: NEW CLOUD REALITY MOVES THE PUBLIC / PRIVATE LINE

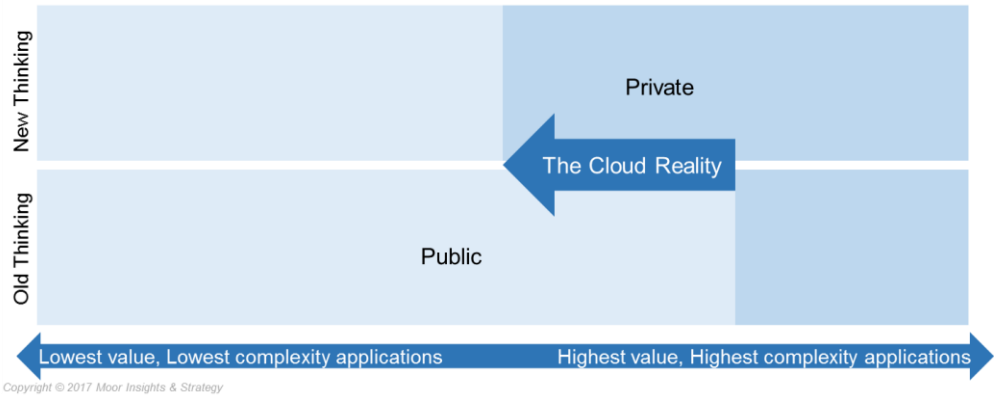

Source: Moor Insights & Strategy

As the understanding of public cloud limitations increases and on-premises compute gains new features and consumption models, old perceptions about where workloads will be deployed begin to shift. The "virtual line in the sand" that marked the difference between public cloud and private cloud is shifting as conventional wisdom indicates more private cloud deployments than in the past.

SECURING BUSINESS & DATA

Public cloud providers have service level agreements related to uptime, but security is less transparent. Security events are becoming more sophisticated, more complex, and

more expensive. Most importantly, security events are now moving past applications and operating systems into the hardware layer. Compromised firmware has been a favorite vector for recent IoT attacks, and servers are now being targeted. While the latest servers are beginning to address this, public cloud economics dictate that workloads are fungible and can be moved at will, making it very difficult to tie critical applications to stronger security platforms as one can do on-premises. The new exploits that target hardware make it more difficult to control the security level without control of the underlying hardware.

Off-premises environments have many variables outside of company control including security and data soverignty, making businesses hesitant to host applications off-premises.

> *"Financial reasons are why we look at the cloud…but for our production environment we're still hesitant about the cloud. The business is hesitant about change."*
>
> – Medium Business, Australia

Clearly, there is a "line in the sand" where IT is comfortable with some applications being off-premises, but not all applications will fit that model. Public cloud providers are not immune to security outages. AWS suffered from a DDoS attack in 2016 that left customers offline for hours.

> *"If there's a problem (in our datacenter), we can deal with it, we don't have to wait on a 4-hour SLA. (Our decision to) outsource is (based on) the criticality of the application and the resources needed to support the application."*
>
> – Medium Business, Chicago

COMPLIANCE

Increasing government regulation is intensifying compliance challenges. The European Union's General Data Protection Regulation brings stricter notification compliance requirements along with sanctions of up to €20M or up to 4% of revenue. In the US, regulations can be by industry (HIPAA for healthcare), by company type (Sarbanes-Oxley for public companies), or by transaction (PCI-DSS for the payment card industry). Data sovereignty laws, which dictate where data must physically exist, are common throughout Europe and Asia and are gaining momentum elsewhere. Regulatory environments drive many to on-premises solutions that help ensure agility along with compliance. The cost of non-compliance can weigh as a heavy tax on non-compliant workloads, rapidly shifting the economic cost picture.

Public cloud migration is often stymied by more generic approaches that may not work well for specific workloads. This dynamic can drive many to host their own private clouds, as their IT departments better understand the intricacies of compliance within their own vertical markets.

> *"Cloud providers will have to come up with a product that will meet the security requirements of government, healthcare, etc., like AWS for Government"*
>
> – Large Enterprise, Chicago

PERFORMANCE

Environments like test/dev and simple infrastructure may not demand the same level of performance and user experience, but as one moves up the criticality ladder to line of business applications, user experience becomes more important.

The need for greater performance, as well as control over the user experience, drives organizations to consider on-premises for both performance and scalability. High performance applications that demand lower latency must be highly tuned for the best performance, an aspect that is lost on public cloud's generic compute instances. A quick glance at typical US cloud latencies shows a best case scenario of ~63ms (and higher beyond the US), dramatically higher than typical datacenter latency in the low single digits. This is especially true for data locality and large workloads. As new CPU and compute performance technologies are introduced, IT will have a greater ability to expose those capabilities and tune on-premises applications because they have platform control. The "compute units" of public cloud are more generic mechanisms for establishing performance. Unless the business has committed to reserved instances and pre-paid for the ability to scale performance up, moving to a higher performing compute unit can be very expensive.

Each new platform generation will have more of an impact on-premises, as IT can optimize these systems to match the workload, increasing the on-premises advantage. Interestingly, public cloud economics depend on a very large percentage of workloads running on older technology or undifferentiated platforms. Matching compute to platform becomes more difficult with each technology change, as cloud providers will find it more difficult to match specific new workload needs against more generic compute capabilities. Scaling on-premises compute up and down is now easier with hyperconverged and composable systems that offer more flexible configurations and more granular scaling to better match applications.

ECONOMIC COST

"Swiping a credit card" to provision and turn on public cloud services is both agile and inexpensive, but costs can scale up quickly as data growth and services sprawl, making public cloud expensive at scale. Using public cloud services is so easy that costs can quickly get out of control; knowing how much your company is spending on public cloud can be difficult. Often enterprises are surprised a few months after using the public cloud at how much they are spending. A common complaint we hear is, "It is very inexpensive to put data into AWS, but extremely expensive to take it back out." Financial lock-in can become a large burden. The point where private cloud crosses over public cloud was once determined to be when the off-premises hosting bill exceeds $7,644 per month. Increased public cloud experience raises visibility to the true costs of public cloud computing, influencing the migration decision.

> *"We cannot get exactly what we want from cloud providers, so we mash up various services"*
>
> – Medium Business, Australia

As few have tried to move applications between different clouds or between public and private clouds, the public cloud's openness is often overstated. Cloud-specific tools often lock developers in and limit flexibility, increasing costs. While Microsoft comes closest, no vendor has a strong hybrid cloud strategy, putting the responsibility on IT to execute. A hybrid solution enables running small, lean, and fast on public cloud with a more reliable production operation on-premises for better scalability while still managing costs. Newer tools—like Docker that brings container functionality and Mesosphere that provides a DC/OS for orchestration and management—enable businesses to receive public cloud's granular control and scalability while on-premises. Much of the cost and focus of moving to the cloud is on creating new cloud-native apps, not making older legacy applications "cloud ready", because the cost of changing application deployment models often outweighs the savings.

HPE'S VISION: THE NEW COMPUTE EXPERIENCE

HPE is betting strongly on hybrid IT, as it believes hybrid IT has the agility, security, and economic control that businesses demand. Successful implementation creates a different platform strategy that speeds time to value with systems that are software-defined, scalable, and composable, so that applications may live in multiple locations, or even move over time. With public and private clouds, multi-cloud approaches, edge compute, and more, IT needs to be agile and deliver applications and services wherever it makes the most sense. The convergence of IoT, rapidly deployed apps, and data that

is moving away from a centralized core, closer to users, are all creating new value and new experiences, but only if IT can innovate. HPE has developed composable and hyperconverged infrastructures designed to help IT deliver this level of innovation on-premises.

To build a bridge to this future state and help businesses navigate through a digital transformation, HPE is building a new compute experience powered by Gen10 servers with Intel Xeon Scalable processors. The new compute experience will feature three key components: agility, security, and economic control.

AGILITY

Agility is about having the flexibility to deliver business results in a better manner, maximizing investments in applications and data through predictable performance and intelligent automation. These key attributes can only be delivered on-premises as IT has full infrastructure control, empowering IT to deliver new value instantly and continuously. Agility helps reduce complexity, enabling rapid deployment with better insight.

Today's complex infrastructure cannot support an agile business. To make the transformation effectively, businesses need agile and flexible solutions like HPE's composable, converged, and hyperconverged systems. The granularity of such solutions better matches the mantra that one size does not fit all. A large component of agility is fitting the tool to the task. This customization can help transform data into insight, enabling a business to more quickly make decisions and own their destiny. Rapid application development fuels this decision making and insight; applications can now change and deploy as quickly as the composable infrastructure, both matching the disruptive changes of the market.

SECURITY

Security is not just about securing systems, it is a broader discussion about securing business and data. Businesses choose hybrid IT, because having control over on-premises infrastructure is one of the best ways to ensure security as well as compliance requirements. To truly protect a hybrid the business and its data, IT needs security that was designed in from the start, from the edge to the core to the cloud.

HPE claims that through its HPE Secure Compute Lifecycle, it will have the world's most secure industry-standard servers. This security begins with the innovative silicon root of trust. Millions of lines of privileged firmware code run before a server operating system boots; if compromised today, it goes undetected. HPE offers industry-standard servers with firmware anchored directly into the silicon, where this silicon root of trust

enables firmware runtime validation. It can detect compromised code or malware residing in the firmware and issue alerts through audit logs. It also has the ability to recover firmware to known good state after detection of compromised code; often the only option today is a service event and a possible hardware replacement. Additionally HPE is adding the Commercial National Security Algorithm Suite (CNSA Suite), the highest level of cryptographic algorithms in the industry typically used for handling the most confidential and secret information. The HPE 3PAR and StoreOnce storage products help ensure end-to-end data integrity so that the right data is always being written, stored, and even recovered in the event of an incident. The network that ties these secure servers and secure storage components together is protected by HPE Aruba ClearPass Manager, which enables the authentication and network access. The newly acquired Niara products use AI and machine learning to profile user and system behaviors, rapidly identifying anomalies that can point to trouble. Security is proven by HPE through its certifications, from CNSA for handling top secret information to NIST 800-53 Certification that certifies the security controls for federal use with mapping to other controls such as HIPP, NERC, DFARS, and ISO27001.

Economic Control

To gain better economic control, HPE is enabling businesses to exploit a better way to consume IT services, only paying for what they use—and doing so within on-premises infrastructure, giving them the best of both worlds. Flexible payment models and usage-based billing enable IT to align costs with business outcomes, even in the world of unpredictable customer demand. The ability to scale on demand eliminates the need to overprovision or incur exponential costs as services rapidly scale up.

Cloud technology drives the conversation from CapEx to OpEx, and HPE can do the same on-premises. While the benefit of public cloud is a pure OpEx model that often supersedes mixed CapEx / OpEx in the typical on-premises infrastructure; HPE's Flexible Capacity model enables a public cloud-like consumption-based model from within the datacenter. With the ability to tie cost to actual IT consumption, HPE Flexible Capacity is built on an OpEx-based model that enables IT to begin the shift toward a consumption model with the ability to scale infrastructure on demand. Buffer capacity can be used on-premises that is only paid for when used, as on-premises applications scale up or down. The only thing billed is the actual metered usage, which eliminates overprovisioning and enables better control over CapEx and OpEx costs. Through HPE's Advisory Migration services, businesses can more effectively transition from legacy IT to hybrid IT as they migrate their services, assets, and data.

As the transition to a hybrid IT environment begins, costs for education, training, and support can rise as businesses must bring new skills to the IT staff. The HPE Datacenter Care Operational Support Services can help augment existing IT skills, again converting costs toward a consumption-based model and enabling the IT teams to more rapidly scale up capability. This program enables IT to free up resources that can be spent with the business units to be more responsive, boosting agility in application deployment while keeping an eye on the cost models.

The combination of products and services from Hewlett Packard Enterprise are designed to help businesses move from a traditional IT world toward a hybrid IT environment by delivering agility, scalability, and cost savings—all without having to relinquish control.

CALL TO ACTION

As businesses move toward a digital transformation, they must carefully consider locality of their compute or their unique right mix.

Moor Insights & Strategy believes hybrid IT will be the prominent mode of deployment, as there are conditions that prevent many workloads from moving to a public cloud. The need for security, compliance, performance, and economic control will clearly tip the scales toward a hybrid IT world. As businesses map out their hybrid IT plans, they should consider HPE solutions for the flexibility, security, cost savings, and scalability they are designed to deliver to IT environments.

To begin the journey of digital transformation, there is a clear roadmap businesses can follow, beginning with learning more about hybrid IT and how vendors like HPE can help impact a company's future direction. Once grounded in the concepts of hybrid IT, businesses can begin to chart their path by looking at their application requirements to understand the value of locality. Finally, stepping out beyond their current infrastructure to see where new disruptive technologies like IoT and mobility are heading will help them move toward that digital future. These steps can play a part in enabling better innovation and closer relationships with customers as the world is transforming. IT's current infrastructure will not be able to keep pace without a serious change.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR
John Fruehe, Senior Analyst at Moor Insights & Strategy

PUBLISHER
Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

EDITOR / DESIGN
Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES
Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES
This paper was commissioned by Hewlett Packard Enterprise. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.