



Research®

CLOUD
PRICE INDEX
NOV 2019

Economics of Hybrid and Multi-Cloud

Owen Rogers, Research Vice President, Cloud Transformation

Jean Atelsek, Analyst, Cloud Transformation and Digital Economics Unit

As it becomes an explicit IT strategy (and one now recognized even by hyperscalers, with their cloud-to-ground initiatives), hybrid cloud presents a major opportunity for service providers and technology vendors. Reducing lock-in for customers through multiple venues might appear risky – what if they want to leave? Cloud-native capabilities, managed services, service guarantees, and ongoing support can make a preferred venue more attractive and – once embedded – decision-makers are less likely to want to move because they are more focused on value than cost. Optimizing utilization and labor efficiency across multiple venues on enterprises' behalf can squeeze costs and give customers the ability to choose the best venue for each workload. New pricing models on both public and private clouds are further enabling enterprise scalability while retaining control.

About the Authors



Owen Rogers

Research Director, Digital Economics Unit

As Research Director, Owen Rogers leads the firm's Digital Economics Unit, which serves to help customers understand the economics behind digital and cloud technologies so they can make informed choices when costing and pricing their own products and services, as well as those from their vendors, suppliers and competitors. Owen is the architect of the Cloud Price Index, 451 Research's benchmark indicator of the costs of public, private and managed clouds, and the Cloud Price Codex, our global survey of cloud pricing methods and mechanisms. Owen is also head of 451 Research's Center of Excellence for Quantum Technologies. Owen has previously held product management positions at Cable & Wireless and Claranet, and has developed a number of hosting and cloud services. He is a Chartered Engineer, a Member of the British Computer Society and a member of the Royal Economic Society. In 2013, he completed his PhD thesis on the economics of cloud computing at the University of Bristol. Owen was named 'Innovative Analyst of the Year' in the Institute of Industry Analyst Relations' global awards in 2018.



Jean Atelsek

Analyst, Cloud Transformation and Digital Economics Unit



Key Findings

- A hybrid cloud approach is being adopted by 57% of enterprises, primarily to enable developers and lines of business to select the optimal technology for each workload's specific requirements – the aim being to add value to the business. Cost is the second-most-cited reason for using hybrid cloud and, in principle, should be the most cost-effective means of providing scalable cloud consumption, provided the private cloud is being managed at high levels of labor efficiency and utilization, and the public cloud and private cloud are unified as a single cloud platform.
- For private cloud infrastructure to be cost-effective against public cloud, utilization should be above 70% over the lifetime of the platform. Negotiating a discount of 30% off list significantly improves this break-even to levels achievable by most enterprises. Enterprises can deploy tactics to increase their utilization, such as using flexible hardware procurement and consumption models to avoid provisioning unused capacity upfront. Aggregating applications from multiple departments and geographies can significantly improve utilization. In fact, the more workloads aggregated, the more likely a flat (and thus more predictable) demand can be achieved, if these workloads' demand requirements are unrelated to each other.
- In our benchmark scenario, unifying private and public cloud capacity saves 10% on enterprise direct spend compared to running each platform separately. Enterprises also consuming private cloud in an opex manner save 22% compared to those running separately managed public and private cloud platforms using a capex approach. Cloud-native frameworks are being developed to take advantage of the range of IT venues while moving attention away from infrastructure and onto the application. In this situation, applications can pull infrastructure dynamically to fulfill their needs, enabling true pay-per-use (rather than pay-per-provision) compute efficiency.
- Using multiple public cloud providers can make direct savings too. US customers can save half their cloud expenditure by mixing and matching services from multiple providers. If enterprises are willing to make upfront commitments, they can save 65% off list pricing. But total savings depend on how feasible it is to manage these cloud platforms, and there are challenges related to skills and management.

- Continuous workload migration is a key use case for hybrid cloud, as applications and workloads must be able to traverse on- and off-premises environments and be available anywhere at any time. Among businesses building their own software, 75% indicate that at least some is cloud-native and that the ability to move applications between environments without significant refactoring is important; 46% consider it very important. Such migration won't take place ad-hoc or in response to second-by-second changes, but will be there to provide options should migration be required.
- Far from a threat, hybrid and multi-cloud must be embraced as an opportunity by the tech industry. Creating lock-in or reducing options might just alienate customers. The tech industry players should enable options but build out their preferred venue so it is attractive through bundling capabilities, guarantees and a reasonable price. Once a provider is a trusted partner, most decision-makers, especially at senior levels, are happy to use a preferred venue even if it is more expensive. To make this happen, new levels of abstraction are required as are new levels of integration and of management. Services that are composable and software-defined (i.e., cloud-native) will help. The IT landscape of the next decade is all about diversity – any infrastructure, any cloud, any app. It's not chaos, but opportunity: Mastering multi-cloud is key to the next decade.

Executive Summary

Introduction

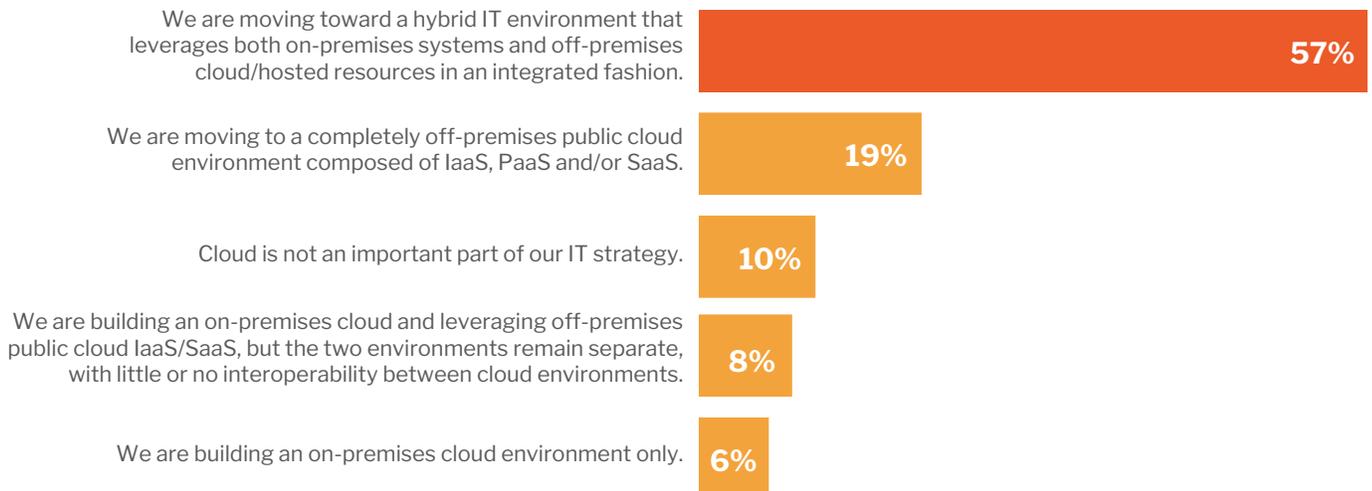
There was a time when hybrid cloud was just a buzzword. Many thought it was a simplified way of saying “I haven’t decided yet”; others thought it was a way for hardware manufacturers to show they were still relevant in an increasingly off-premises world. But hybrid cloud – defined as the use of public and private cloud venues – is now a reality: According to 451 Research’s *Voice of the Enterprise (VotE): Workloads and Key Projects 2019 survey*, 57% of enterprises today are pursuing a hybrid IT strategy in a unified way (an ‘integrated’ hybrid cloud). A further 8% are planning to use hybrid IT, but in a non-unified way (‘segregated’). Less than one in five are using public cloud only, and 1 in 12 are focusing only on private cloud. The use of multi-cloud – the use of multiple public cloud providers potentially alongside a private cloud – is also a fact: 71% of enterprises using public cloud today have more than one provider, with 45% spending more than 80% on their primary cloud vendor. Hybrid cloud now is the dominant strategy, and multi-cloud (even if it’s lopsided in favor of a single provider) is also common (see Figure).

Organizations’ Overall IT Approach

Source: 451 Research’s *Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019 Q. Which of the following best describes your organization’s overall IT approach and strategy?*

THE FUTURE IS HYBRID FOR MOST

% of respondents (n=571)



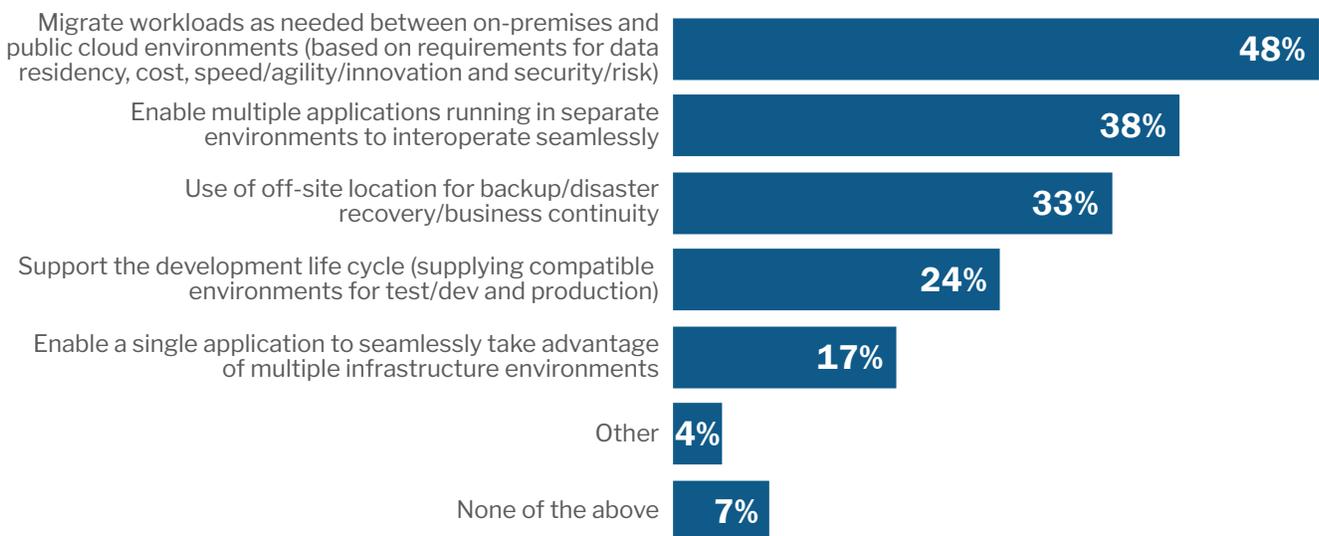
In one way, hybrid is as much a mindset or philosophy as it is a technological development – it’s a way for enterprises to say, “We’re opened-minded, we want to keep our options open.” But why keep options open? Because enterprises want to choose the best technologies available to them. They may want to use Amazon Web Services (AWS) for infrastructure, Microsoft for on-premises, Google for AI projects, Oracle for databases and IBM for bare-metal. Keeping options open so enterprises can choose best-in-class technology is a big selling point of multiple cloud environments and it allows departments and customers the flexibility to choose the best provider for their specific needs – 38% of our survey respondents say access to vendor-specific platform capabilities is their top reason for adopting multi-cloud. It also allows enterprises to develop resiliency and reduce lock-in by supporting multiple environments. Increasingly, we’re hearing that having multiple providers in their vendor portfolio aids contract negotiation as enterprises pit competitors against each other for lower prices.

Optimizing for cost is also high on the agenda. Yes, enterprises want value for money – they want capability that will give them a cutting edge in the business. But they also want to keep their costs in check. Some workloads simply can’t move off-premises, either due to regulation (requiring data to remain on-premises or in-country, for example) or because refactoring them for efficient cloud operation can’t be done cost-effectively. In the first phase of cloud adoption, some organizations lifted and shifted their applications into public cloud, only to discover that they cost much more to run at scale off-premises than in their own datacenters. Organizations are smarter about where to place their IT workloads, but sometimes leaving them on-premises is the best choice. Hybrid cloud, if executed correctly, can give the best of both worlds (see Figure).

Hybrid/Multi-Cloud Use Cases

Source: 451 Research’s *Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019*
 Q. Which of the following use cases are most important to your organization’s use of hybrid/multi-cloud environments?

% of respondents whose organization is moving toward a hybrid IT environment (n=336)



Methodology

This report utilizes data from our quarterly Cloud Price Index Public Cloud analyses, which captures pricing data from AWS, Google, Microsoft, Alibaba and IBM's pricing APIs and websites, plus our Q3 study, *The Economics of Private Cloud*.

WEEKLY	QUARTERLY	ANNUALLY			
Cloud Tracker	Public Cloud Country Benchmarks	Economics of Managed Services	Economics of Private Cloud	Economics of Hybrid Cloud	Public Cloud Review
		PDF Report	PDF Report	PDF Report	PDF Report
	Spreadsheet	Spreadsheet	Spreadsheet		
	Cloud Price Index Tool				
Analyst Note					

This report is the fourth long-form report published to Cloud Price Index subscribers in 2019.



Table of Contents

1. Utilization Is Everything	1
<i>Figure 1: Example Application Demand Profiles</i>	2
<i>Figure 2: Heatmap of Self-Managed Private Cloud vs. Public Cloud</i>	3
<i>Figure 3: Heatmap of Self-Managed Private Cloud vs. Public Cloud, With Example Workloads Included</i>	4
<i>Figure 4: Aggregated Customer Demand Profiles</i>	5
<i>Figure 5: Heatmap of Self-Managed Private Cloud vs. Public Cloud, With Example Workloads and Aggregated Workloads</i>	6
<i>Figure 6: Typical Time Taken To Procure Private Cloud Capacity</i>	7
<i>Figure 7: Aggregated Demand Profile With Stepped Capacity</i>	8
<i>Figure 8: Heatmap With Stepped Capacity</i>	8
Hybrid Cloud	9
<i>Figure 9: Aggregated Demand Profile With Stepped Capacity and Public Cloud ‘Bursting’</i>	9
<i>Figure 10: Heatmap With Hybrid Capacity</i>	10
Other Applications	10
Negotiation	11
<i>Figure 11: Heatmap With 30% Discount Showing Shift in Break-Even</i>	12
<hr/>	
2. In Practice	13
<i>Figure 12: Experience in Using Hybrid Cloud, by Integrated or Segregated</i>	13
<i>Figure 13: Specific Experiences in Using Hybrid Cloud, by Integrated or Segregated</i>	14
<i>Figure 14: Heatmap With Integrated and Segregated Average Performance Levels</i>	15

3. Multi-Cloud	16
<i>Figure 15: Multi-Cloud Savings Achievable Over Average Single-Cloud Price</i>	17

4. Cloud-Native and Containers	18
Portability Is Critical to Cloud-Native Apps	18
<i>Figure 16: Importance of Cloud-Native Portability</i>	19
Adoption of Cloud-Native Technologies	20
<i>Figure 17: Adoption of Cloud-Native Technology</i>	20

5. Let It Be	22
<i>Figure 18: Price-Sensitivity by Adoption Stage</i>	23
<i>Figure 19: Price Sensitivity by Buyer</i>	24

6. Conclusions and Recommendations	25
Recommendations for Service Providers and Technology Vendors	25
Recommendations for Enterprises	25

7. Further Reading	27
---------------------------	-----------

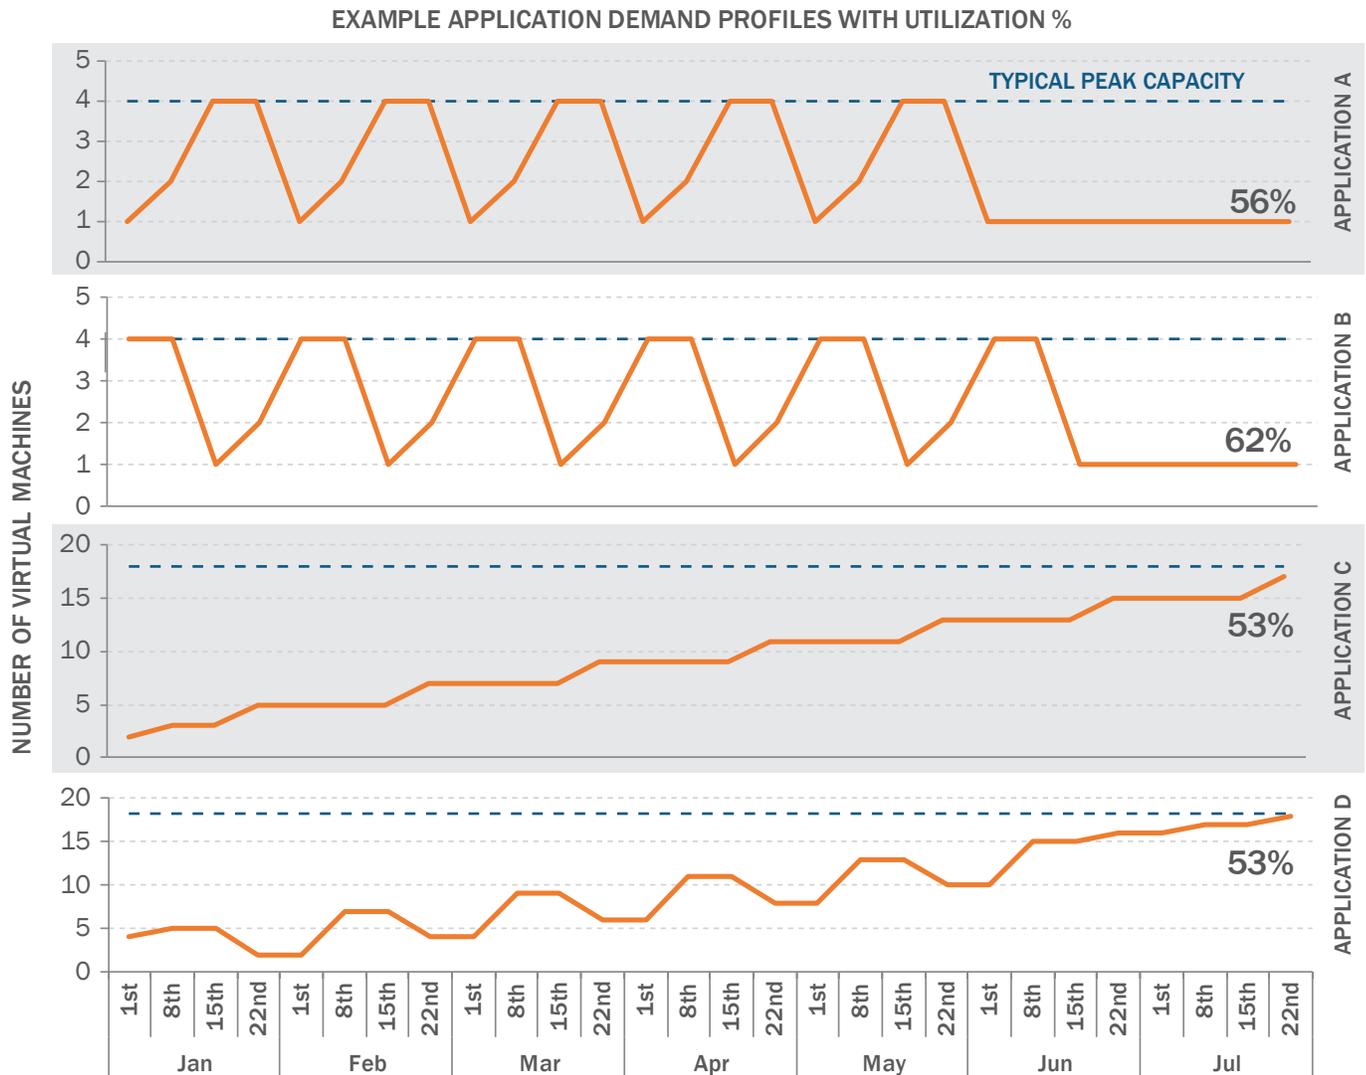


1. Utilization Is Everything

Figure 1 shows an example of changing VM demands for a portfolio of applications, with numbers intentionally chosen to be small for a simpler explanation. Over a period of a seven months, these cloud-native applications consume VMs in different cycles. The first might be an enterprise application that batch processes every fortnight. The second might be exactly the same app but deployed in a different time zone. The third app might handle backup or archiving growing as protected data grows, while the fourth might process and index increasingly larger datasets. If each of these applications were to have its own private cloud architected to support its typical peak capacity (indicated by the blue line), the utilization achieved over the course of the overall period would be the percentage shown. Capacity planning such that performance can be maintained is important for end-user experience as applications scale, but it must be fitted to demand as closely as possible to avoid waste – unused resources are sunk costs.

Figure 1: Example Application Demand Profiles

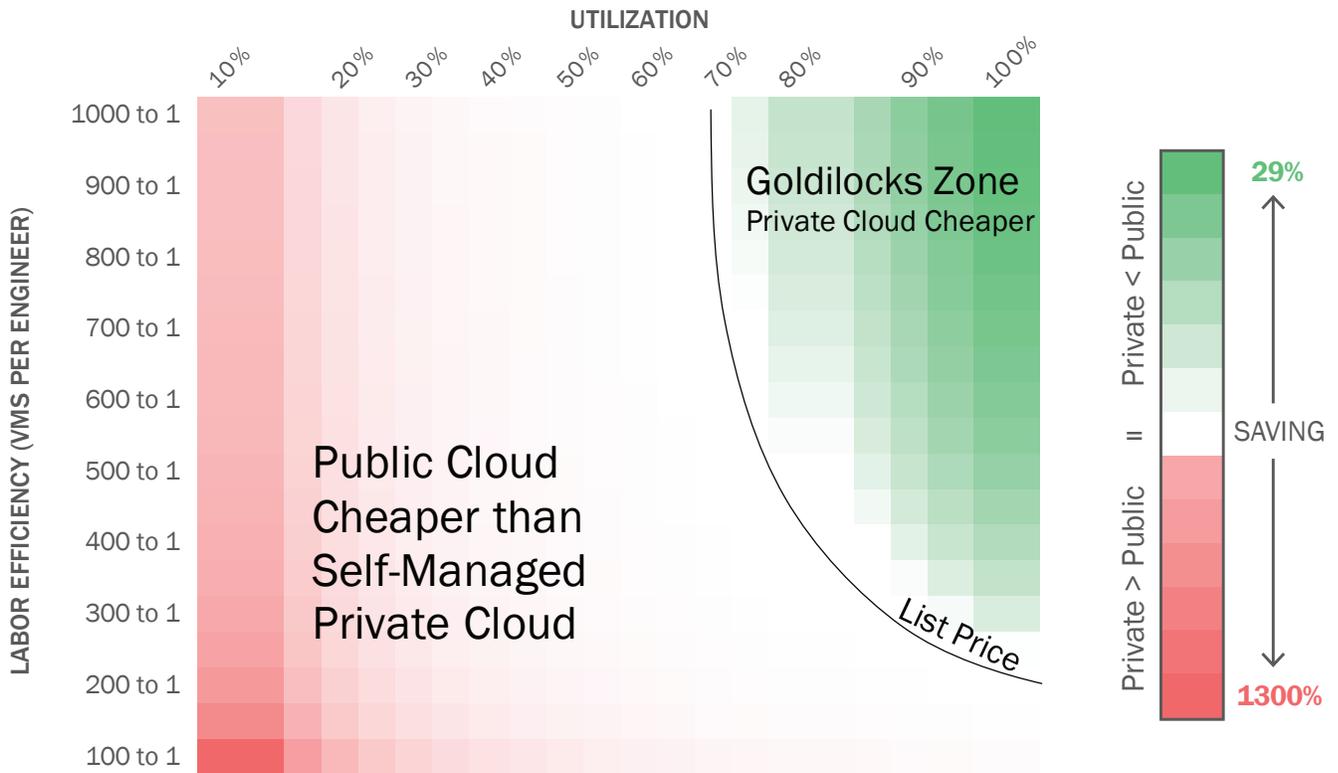
Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



For the past few years, 451 Research has published a heatmap that shows the break-even points for total cost of operations between public and private cloud in our Economics of Private Cloud report. Figure 2, the latest iteration, shows which option is cheapest depending on utilization and labor efficiency, and is based on the CPI public cloud benchmark and the average private cloud quote we received. As discussed, for fixed-capacity private clouds, utilization impacts unit costs. On-demand public cloud doesn't require such capacity planning, and thus isn't as susceptible to utilization as a factor.

Figure 2: Heatmap of Self-Managed Private Cloud vs. Public Cloud

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019

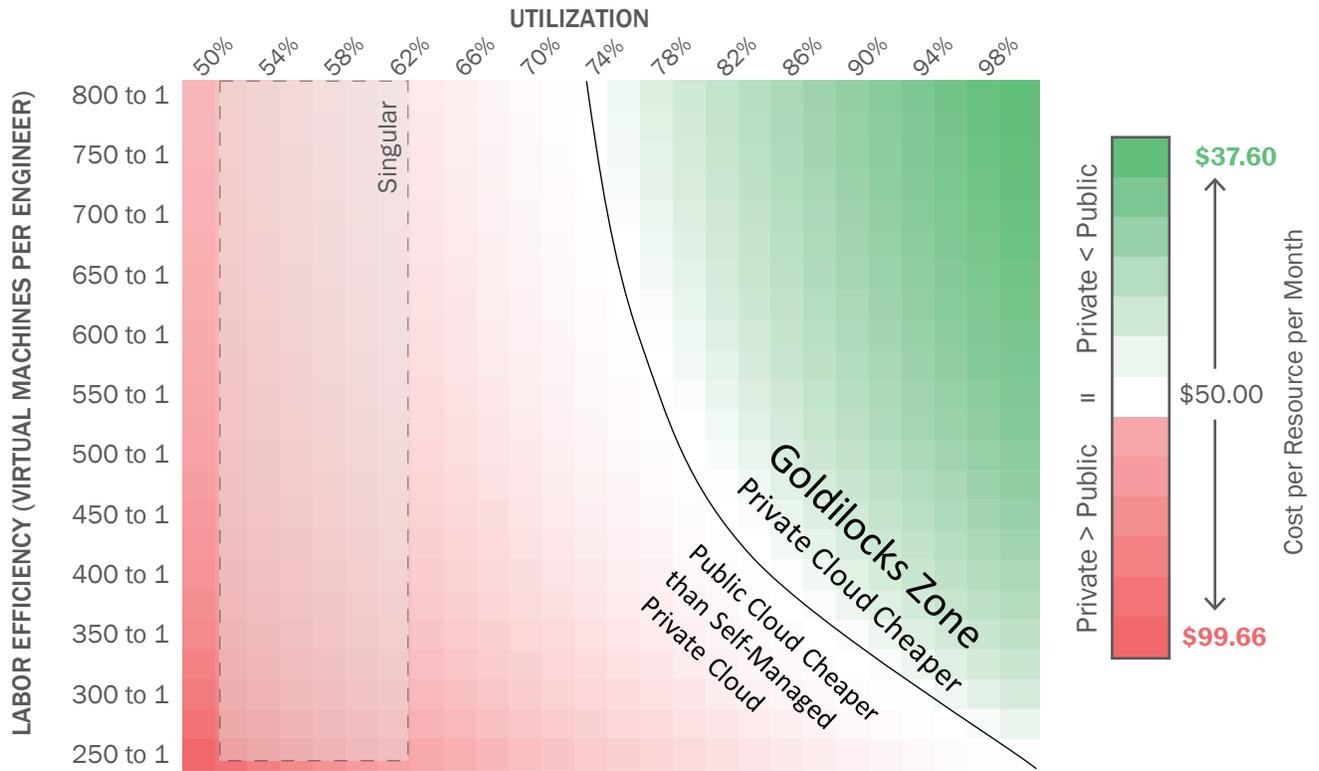


The objective of the heatmap is to see under what conditions buyers can save money using a private cloud compared to a public cloud. Of course, individual circumstances will differ. To assess the feasibility of this idea, we built a quotation from the list pricing of standard servers from HPE, software from VMware, and colocation from Rackspace in the northeast US. This was intentionally chosen to be conservative by taking a mix of nonrelated vendors and service providers.

The green area shows where private cloud provides more savings than public cloud does, the red area is the other way around, and the white area effectively represents the break-even point. This heatmap reflects common sense: The more the private cloud is used (utilization) and the better managed it is (VMs per engineer), the cheaper the unit cost of each private cloud resource will be, and the more it can save when compared to public cloud. Crucially, private cloud becomes cheaper at 70% utilization and labor efficiencies of 200:1.

Figure 3 shows where the individual workloads in our portfolio (which achieved 53% to 62% utilization) fall on the public vs. private cloud curve. None of these applications would have a lower TCO using private cloud compared to public cloud, as none are in the Goldilocks Zone.

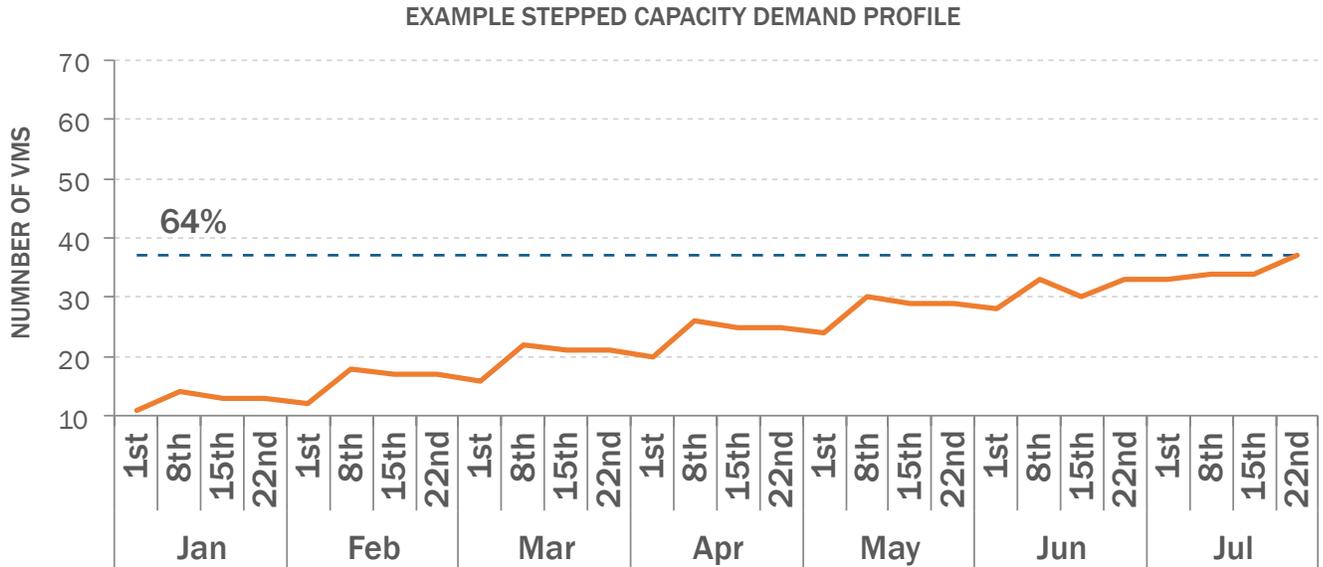
Figure 3: Heatmap of Self-Managed Private Cloud vs. Public Cloud, With Example Workloads Included
 Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



However, these applications can be aggregated onto a single private cloud. One private cloud, potentially hosted and managed by a service provider, can support all these applications, so why not operate them together as a single infrastructure? How does this affect utilization, and what does our total VM demand look like now? Figure 4 gives us a snapshot.

Figure 4: Aggregated Customer Demand Profiles

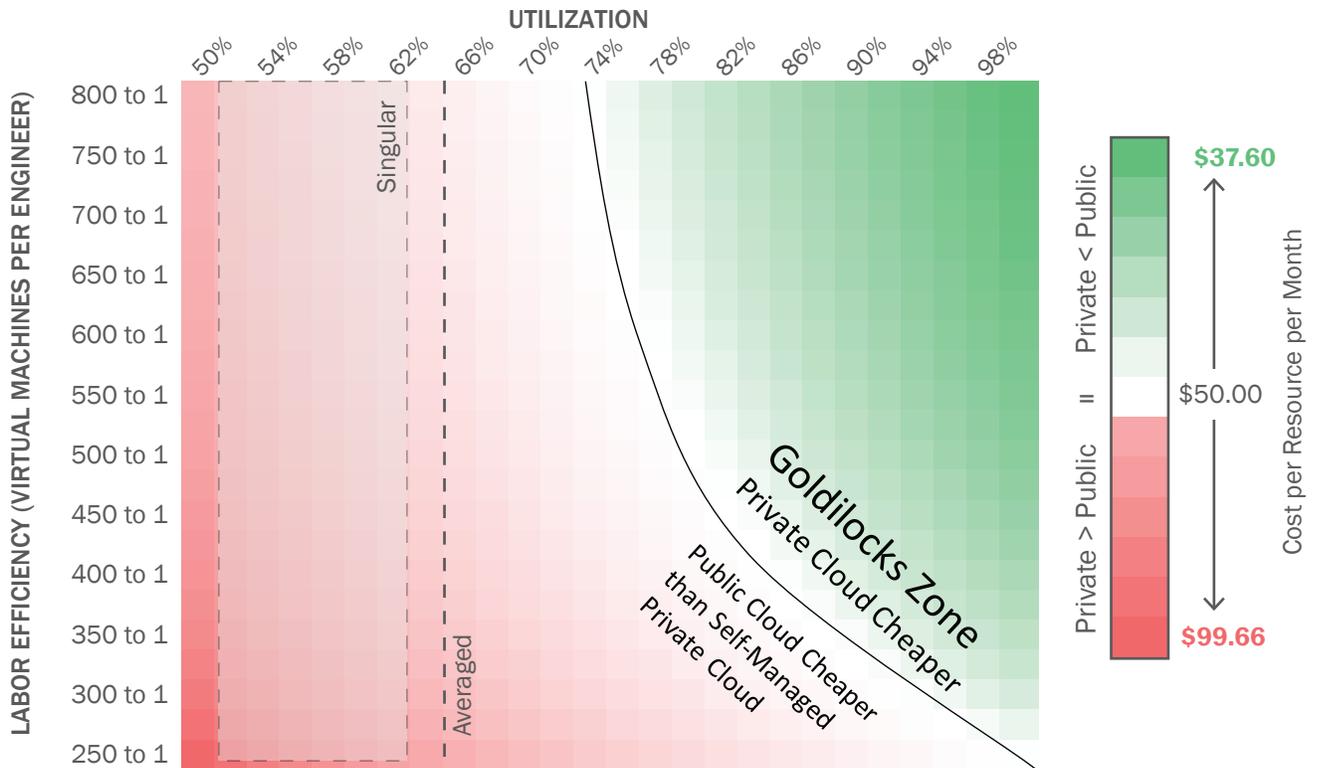
Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



We haven't cut any corners. All our applications are receiving exactly the same VMs they required previously, each containing the same resources. The difference is that by aggregating workloads together, we've increased our utilization. This is shown in Figure 5.

Figure 5: Heatmap of Self-Managed Private Cloud vs. Public Cloud, With Example Workloads and Aggregated Workloads

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



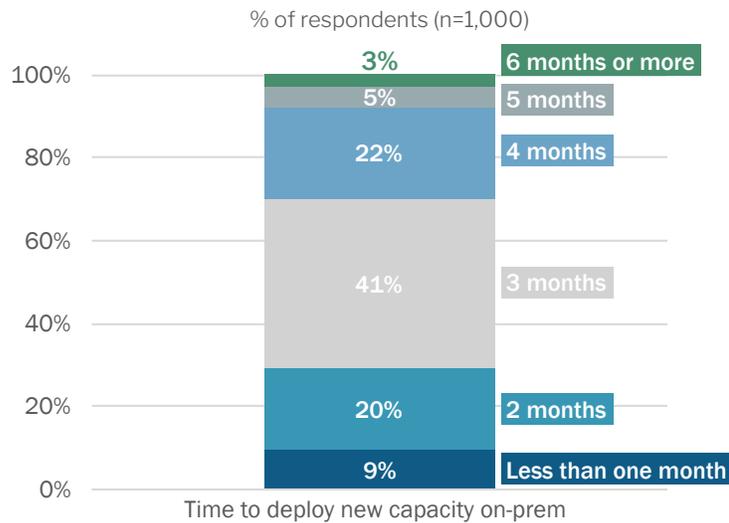
This workload superimposition means that our private cloud requires less capacity than the sum of its parts. In fact, the more randomly impacted, independent applications situated on a private cloud together, the closer the private cloud will be utilized to a fixed capacity (due to the mathematical Law of Large Numbers). And by having a fixed capacity, we can control our maximum utilization, so costs don't spiral out of control. We can also dictate contention and over-provisioning – anecdotally, we hear of public cloud resources not being turned down when they are serving no useful purpose, essentially generating a cost implication without generating value. In this private cloud model, capacity that isn't being utilized can be automatically turned over to other applications.

In the aggregated demand scenario, it's easy to plan maximum capacity as there is a regular peak demand. What happens if there is an upward trend? Then we must plan for the maximum capacity, which could be years or months in advance. The key here is to step capacity – grow it at intervals, as shown in Figures 6 and 7. In our scenario here, if we're operating at a labor efficiency of 500:1 VMs, we should beat public cloud on price.

This scenario requires a way to scale up (or down) private cloud capacity quickly, whereas procurement cycles for equipment have traditionally been weeks or months long. Figure 6 shows it typically takes three months or more to deploy new capacity, according to 1,000 enterprise decision-makers around the world surveyed for a custom study conducted for HPE GreenLake.

Figure 6: Typical Time Taken To Procure Private Cloud Capacity

Source: 451 Research, commissioned by Hewlett Packard Enterprise GreenLake, 2019



Vendors have bridged the gap in the form of pay-as-you-go private clouds, which install buffer capacity over a set baseline that represents the customer’s minimum monthly commitment; usage beyond the baseline level is metered and billed in arrears. This is a more ‘cloud-like’ billing model for hardware vendors willing to share the burden of capacity planning and the risk of excess resources remaining unused and unpaid for. The veteran provider for this model is HPE with GreenLake Flex Capacity; other vendors/products with similar flexible consumption offerings include Dell EMC Flex on Demand, Oracle Cloud@Customer, Nutanix Go, Lenovo Truscale Infrastructure Services, and managed service providers using these platforms as the basis for their private cloud offerings.

Figure 7: Aggregated Demand Profile With Stepped Capacity

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019

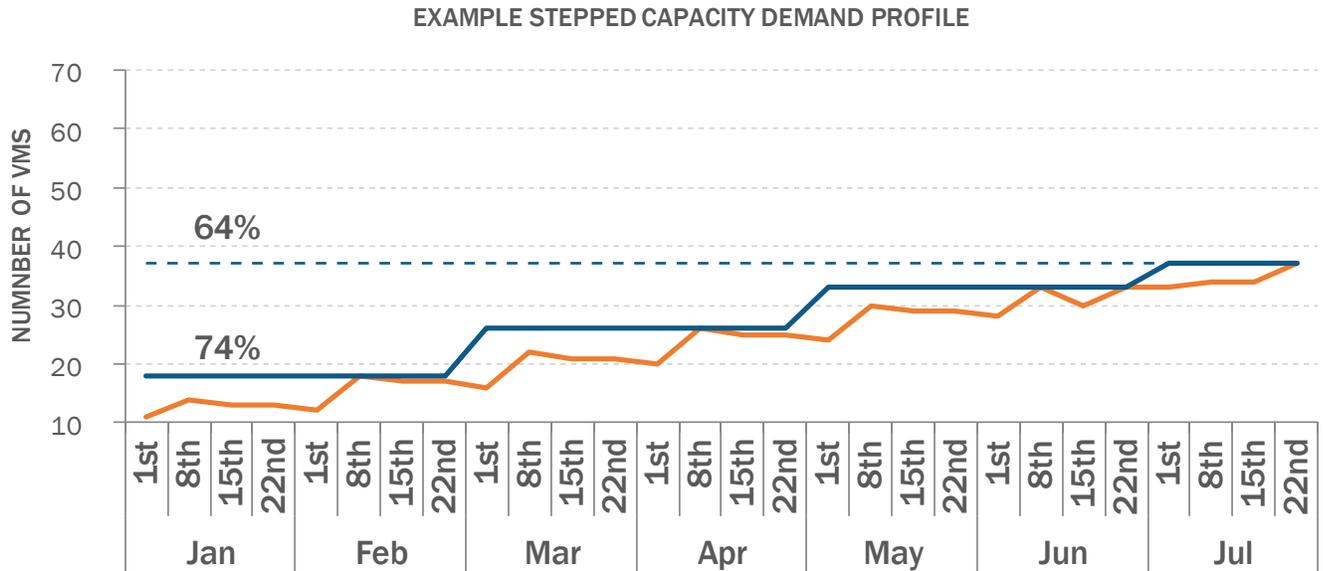
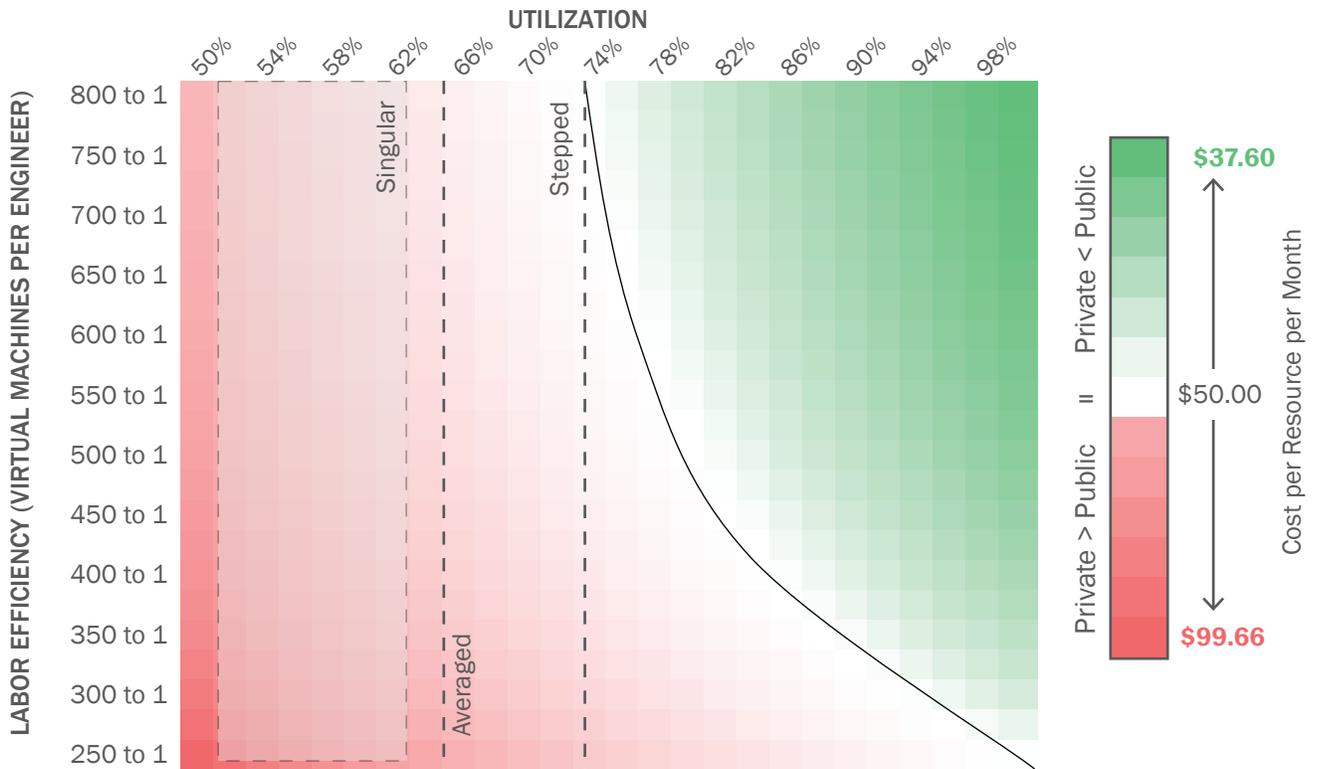


Figure 8: Heatmap With Stepped Capacity

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019

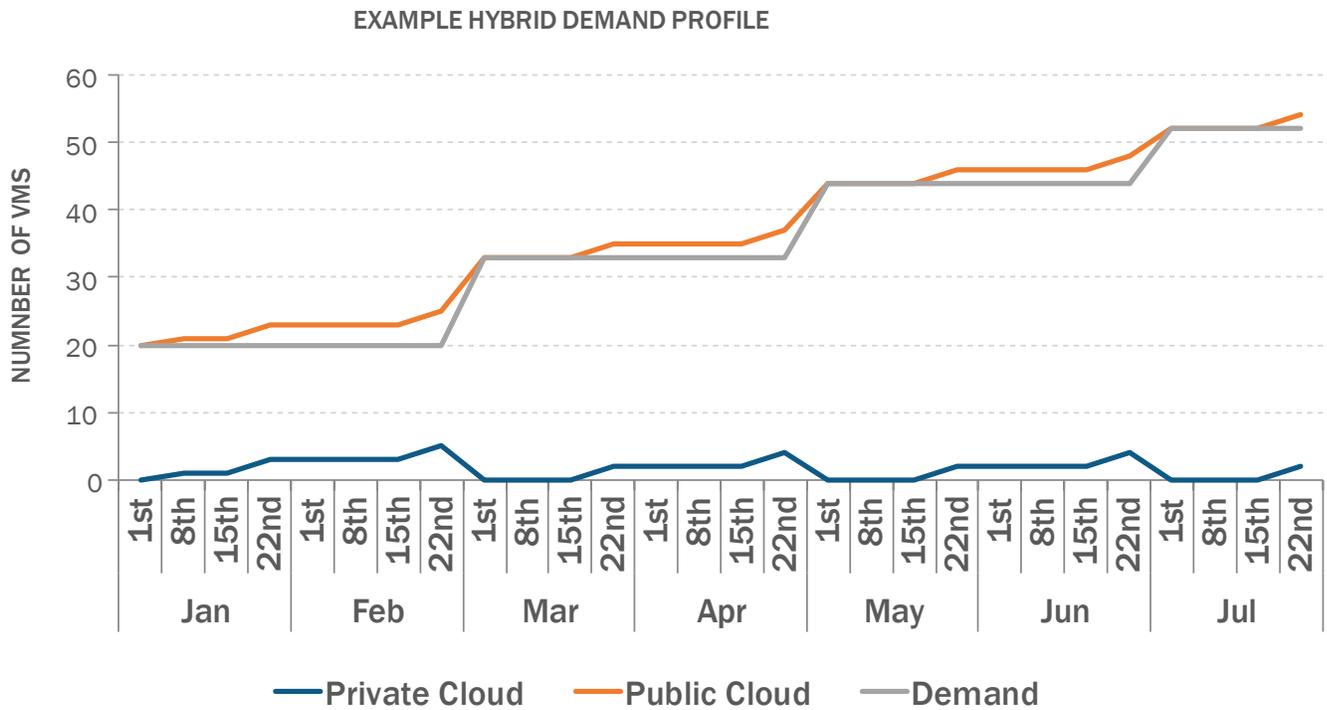


Hybrid Cloud

Is there a way we can increase our utilization even further? We could more closely step up our capacity, but the closer we get to 100% utilization, the bigger the risk we'll need to burst beyond the resources we have on-site. However, in a hybrid cloud model we can use a private cloud for delivering a stable and constant capacity, and a public cloud for bursting above our maximum capacity. Hybrid cloud, in the context of this report, means public and private cloud. In Figure 9, the private cloud is providing a minimum capacity of 20 to around 50 VMs, which are 100% utilized. When additional capacity is needed, the applications spin up additional VMs from a public cloud. With this hybrid approach, maximum efficiency is achieved: There is no sunk cost, no performance or resource constraints, and the cheapest resources are used when required.

Figure 9: Aggregated Demand Profile With Stepped Capacity and Public Cloud 'Bursting'

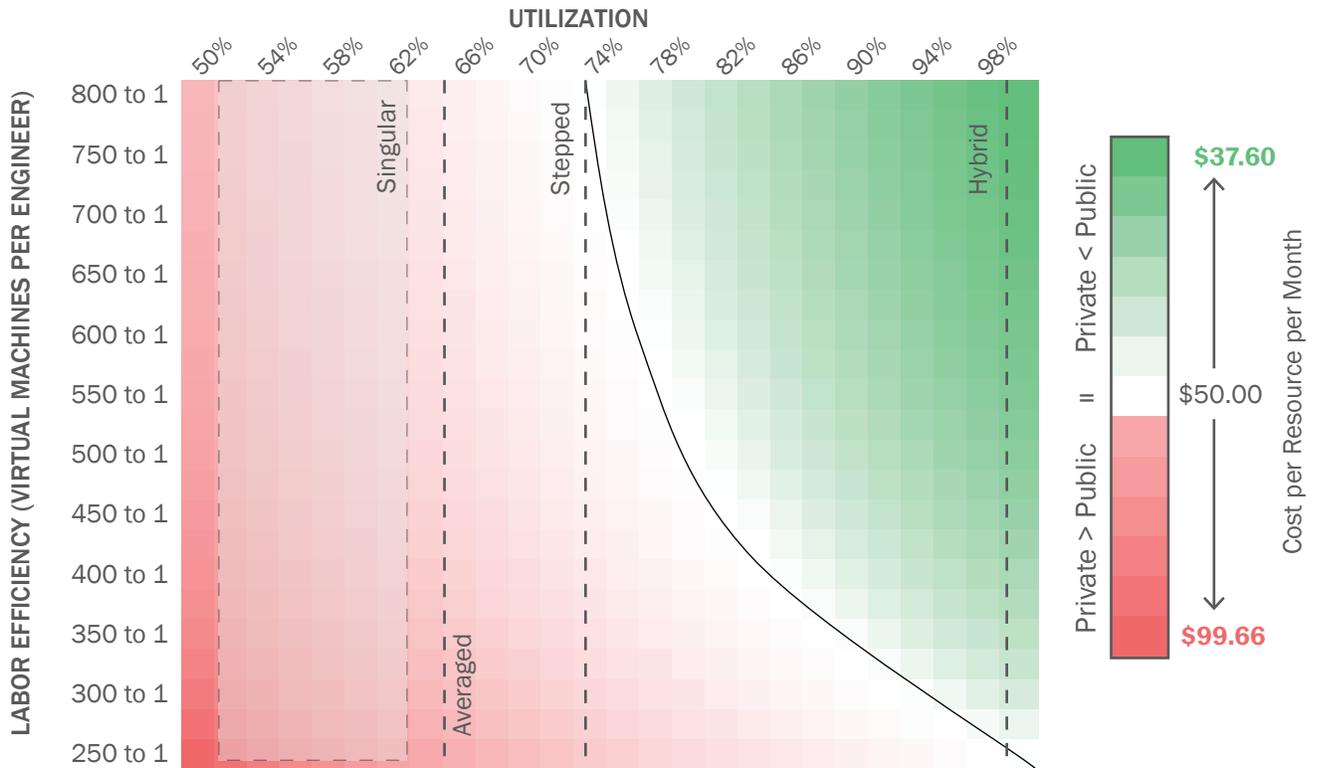
Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



At these high levels of utilization, it becomes easier to make private cloud TCO savings over public cloud. In Figure 10, an 800:1 labor efficiency, private cloud is as cost-efficient as public cloud, even without a private cloud discount.

Figure 10: Heatmap With Hybrid Capacity

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



Granted, managing on- and off-premises resources in an integrated fashion is not easy. Maxing out private cloud usage and launching public cloud resources only when elastic capacity is needed requires robust engineering and automation. Yet companies that need to operate in both public and private cloud domains are succeeding at doing so. A large share of enterprise workloads continues to run cost-effectively in-house. While dynamic allocation of resources based on workload demands is becoming a reality for net-new workloads in public cloud, legacy infrastructure and software can't simply be abandoned. Vendors and software developers are succeeding at bridging the gap between the two.

Other Applications

Workloads that aren't time-critical and can be scheduled ad-hoc are likely to help a private cloud reach higher utilization thresholds. In this manner, the private cloud can constantly be fully utilized – any spare capacity can be reallocated to applications already executing to speed up the result or can be reallocated to a new application. This approach is ideal for big data; critical time-to-results processing; batch processing; economic, financial and scientific analytics; or even overnight archiving and indexing. Service providers running their own private clouds could potentially offer spare capacity for a discounted rate, the proviso being that the

VMs can be terminated spontaneously to make way for full-price paying customers with real-time requirements. This means the infrastructure can be highly utilized, but customers who spontaneously need more capacity don't have to queue. (The largest public cloud providers do this by offering ephemeral 'spot instances' at deeply discounted rates (up to 90% less) versus on-demand compute capacity.)

Not all applications need to change their scalability every few hours or to a cycle every few days. Applications that have a fixed number of users or accounts, each requiring a fixed amount of resources, will have a fairly constant utilization. A company's workforce, for example, will usually stay constant over shorter timeframes, or at least vary within a certain range. Enterprise applications used extensively by all members of the workforce (or even teams) will typically have large utilizations, subject to good capacity planning. If there is a need to add a significant number of employees as a result of large-scale recruitment, the private cloud can scale by adding a bulk capacity increment. As there would typically be some notice period that such a recruitment drive would take place, it can be built into advance planning. A private cloud partner that can rapidly provide additional hardware in response to additional demand can help ensure high utilization rates while maintaining performance.

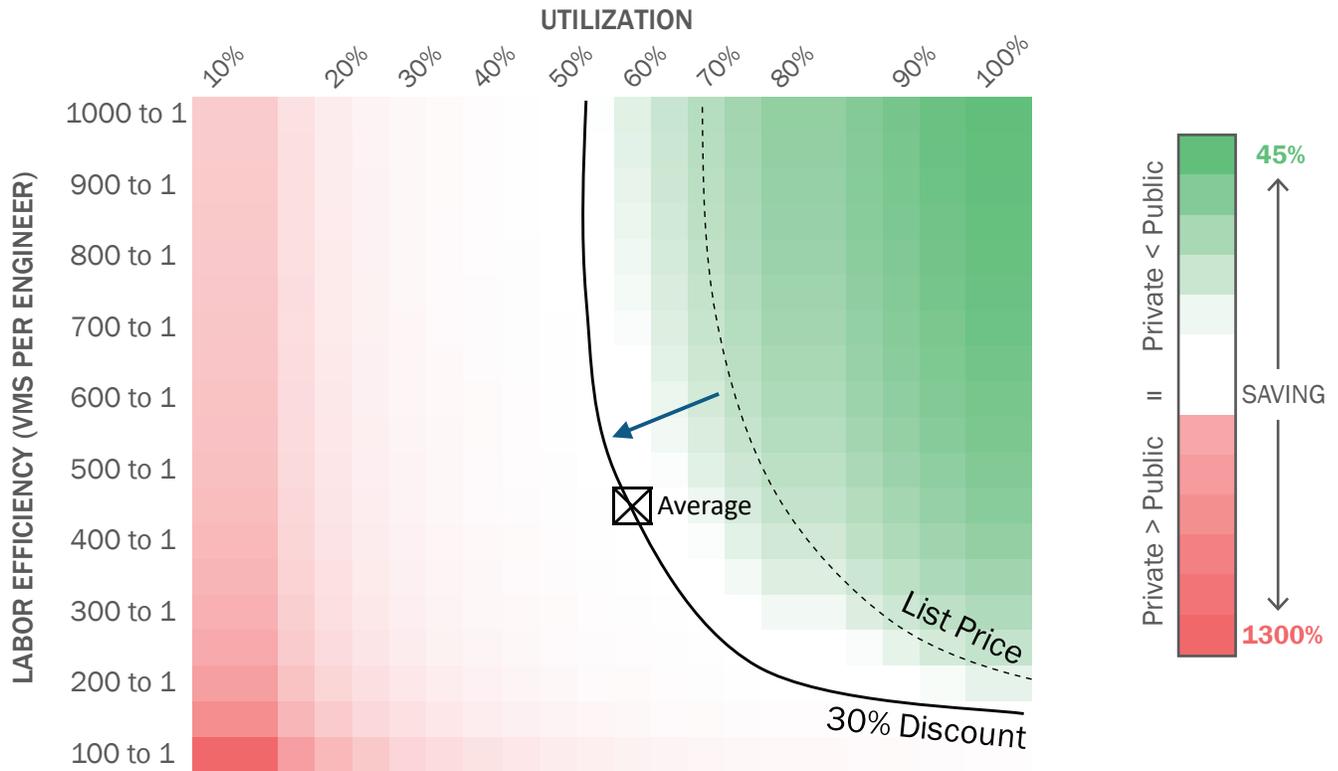
Negotiation

In 2017, 451 Research conducted a custom study for HPE and found enterprises were running at averages of 460 VMs per engineer at 58% utilization. A repeat of the study completed just as this report was going to print shows that these averages are still reasonable in 2019. This combination of utilization and labor efficiency is marked in Figure 11 with an 'X' and is outside the Goldilocks Zone, in which private cloud is cheaper.

But we need to examine the impact of discounts in our analysis. Most hyperscalers only give around a 7% discount for annual spends of more than \$1m. In our annual survey of *private cloud pricing*, our VM requirement only costs around \$186,000 a year based on CPI public cloud benchmarks and so wouldn't be subject to a discount. However, an upfront purchase of around \$250,000 for our private cloud hardware and software would usually attract a discount from hardware manufacturers and software vendors keen to show flexibility – anecdotally, we hear such discounts can be huge, up to 75%. With software, discounts can be generous because of the low incremental cost of producing an additional software unit – in other words, selling a piece of software has very low variable costs (all the cost is fixed during the software development). With hardware, discounts can be generous because of a highly price-competitive market. And of course, many enterprises have already invested in hardware and software that are already sunk costs and can be reused or have vendor relationships that extend far beyond hardware and software.

Figure 11: Heatmap With 30% Discount Showing Shift in Break-Even

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019



A discount on self-managed private cloud makes the break-even point easier to achieve. If a buyer managed to save 30% on hardware and software through reuse of hardware, porting of licenses and negotiation, they only need to achieve 50% utilization and 150:1 labor efficiency to break even. And the savings achieved will be greater beyond this point. Crucially, the discount now puts the average enterprise on the break-even threshold for the Goldilocks Zone.

In other words, cloud buyers are probably more likely to get a discount from a hardware manufacturer or software vendor than a public cloud provider. Buyers should reuse licenses and hardware and maintain lines of communication with partners to achieve these savings, because they heavily impact the case for private cloud cost-efficiency, and its role in hybrid cloud.

2. In Practice

In the previous section's scenarios, we assume that:

- The application is architected such that it can scale (perhaps using cloud-native frameworks).
- The cost of managing the venues does not significantly impact the savings made through higher utilization.

However, the data does suggest that companies that have managed to integrate their on- and off-premises clouds together are experiencing lower costs and a better experience than those that are operating them in a segregated fashion. In a custom study for HPE GreenLake conducted in 2019, we asked 1,000 enterprises around the world for their experiences in using public and private cloud. Figure 12 shows that businesses cite a better experience when using an integrated hybrid approach. Figure 13 shows they are also managing to control their environment – users of integrated hybrid clouds have fewer complaints about performance, orphaned resources or visibility. This is the opportunity that has led the major public cloud providers to move into on-premises datacenters with offerings such as AWS Outposts, Azure Stack, Oracle@ Customer and Google Anthos.

Figure 12: Experience in Using Hybrid Cloud, by Integrated or Segregated

Source: 451 Research, commissioned by Hewlett Packard Enterprise GreenLake, 2019

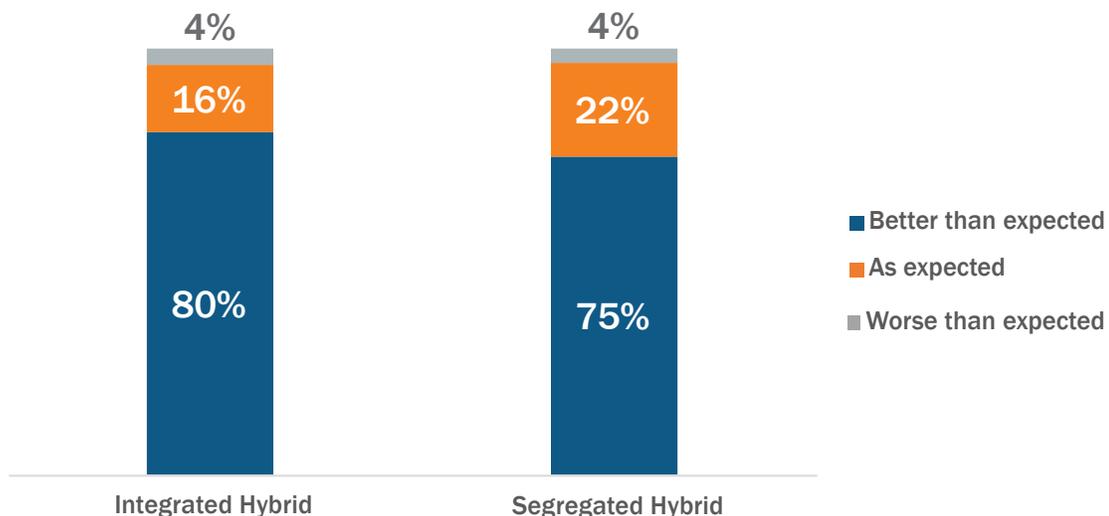
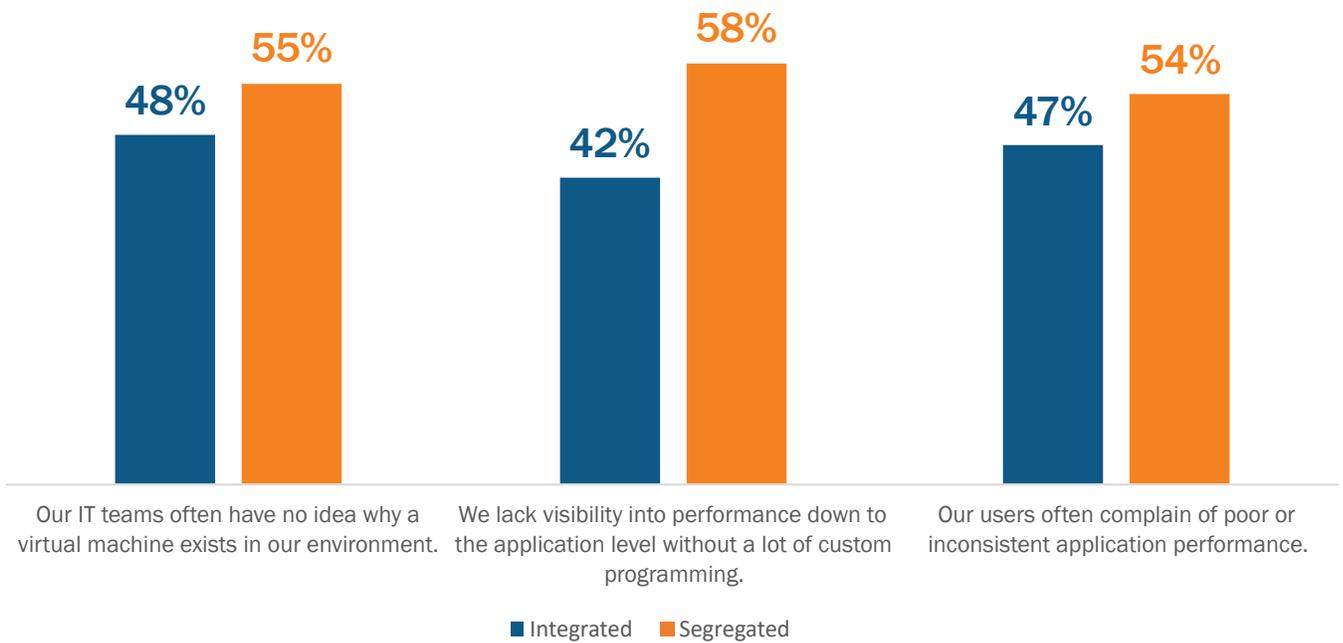


Figure 13: Specific Experiences in Using Hybrid Cloud, by Integrated or Segregated
 Source: 451 Research, commissioned by Hewlett Packard Enterprise GreenLake, 2019

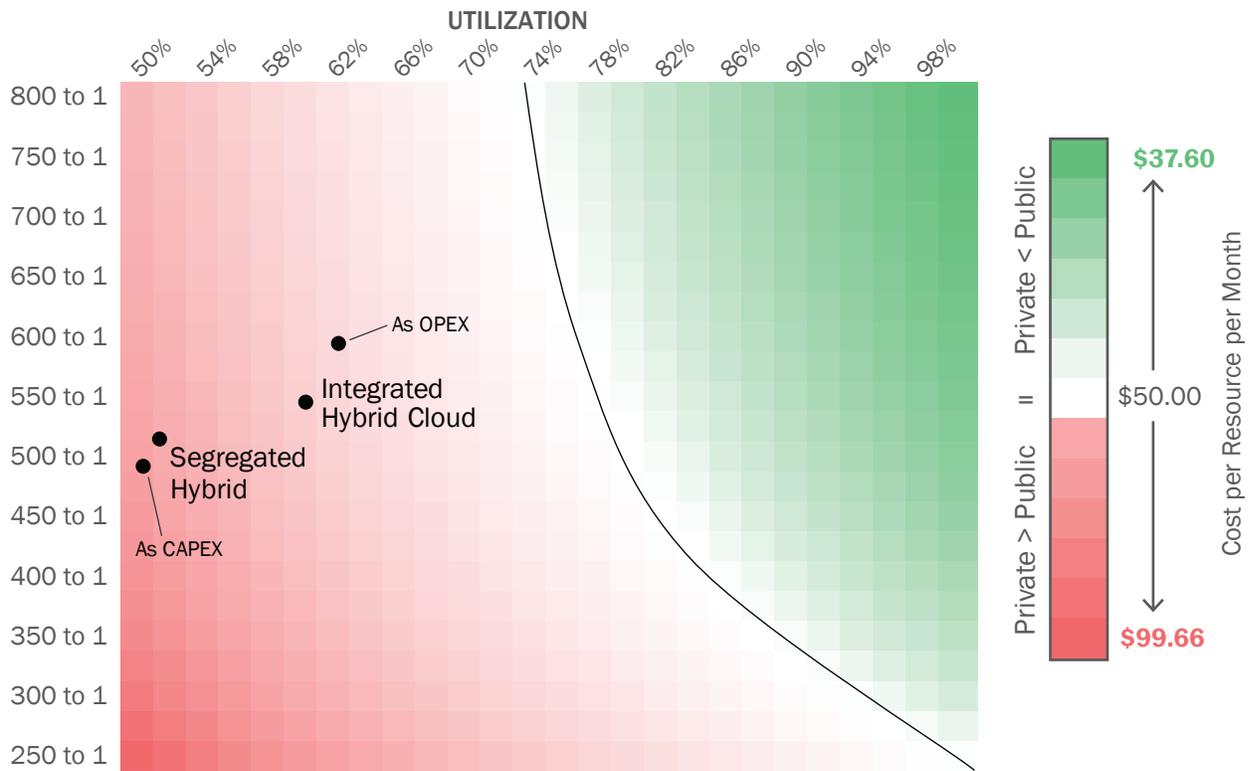


Enterprises following an integrated approach are also experiencing higher levels of utilization and labor efficiency. Integrated hybrid has an average 60% utilization and 546:1 VMs per engineer labor efficiency. Contrast this with segregated, which has averages of 59% utilization and 514:1 labor efficiency. This translates to a 10% difference in the cost of each resource: \$138 for the integrated, \$152 for the segregated.

When we layer in how enterprises are paying for server capacity, the gap becomes more obvious (see Figure 14). Enterprises that are using an integrated hybrid approach and paying for capacity in an operating expenditure model are operating at 62% utilization and 592:1 labor efficiency. Those paying a fixed capital expenditure are operating at 50% utilization and 491:1 labor efficiency. This translates to a hefty 22% reduction in resource cost.

Figure 14: Heatmap With Integrated and Segregated Average Performance Levels

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019





3. Multi-Cloud

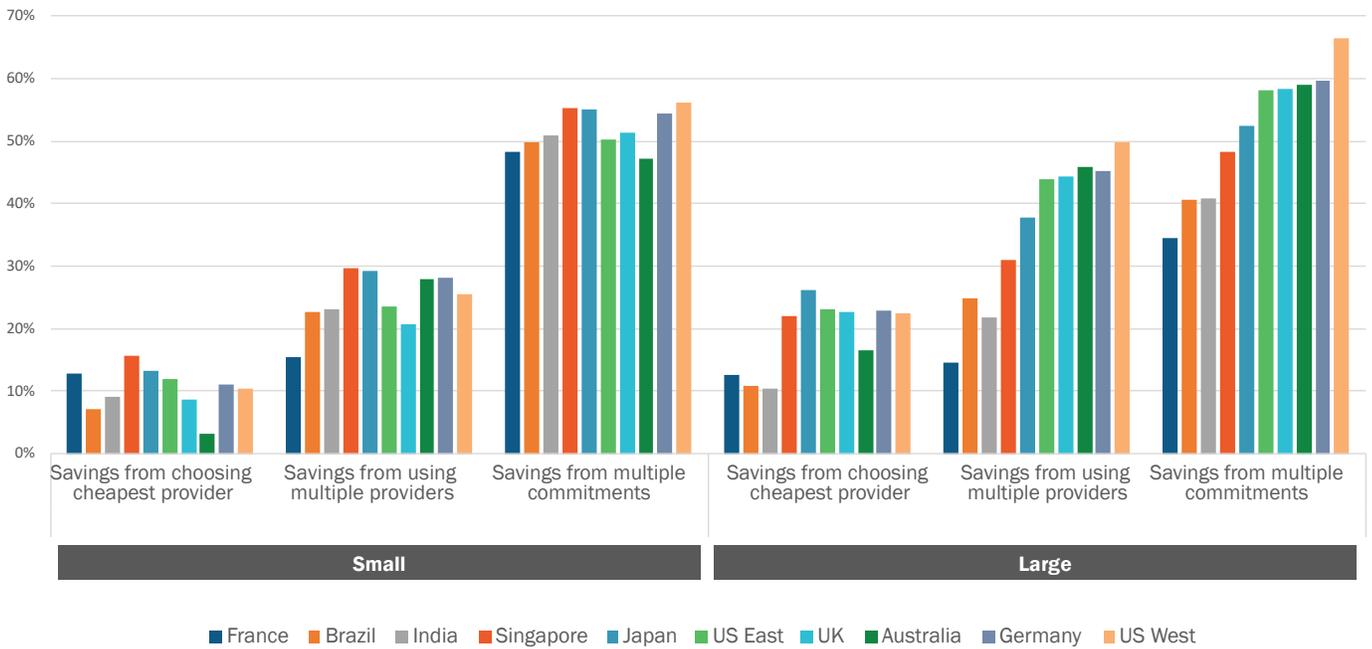
There can be huge value in using multiple public clouds, and the more complex the application, the greater the savings that can be gained on direct expenditure (see Figure 15). Shopping around for a cheaper single provider to deliver all requirements will save just 3% in Australia on the average cost of the Public Cloud Price Index's small application specification, with a more enticing 16% savings in Singapore. But mixing and matching different services can save a massive 30% in Singapore, and a still worthwhile 16% in France. Shopping around for our large basket specification can make huge savings too, from 10% in India to 26% in Japan. Mixing and matching can halve costs in the US West region.

When multiple commitments come into play, the overall savings can be vast. Anyone can consume reserved instances, sustained-use pricing models and pre-provisioned database capacity – there is no minimum spend required to make an enterprise eligible. As such, an enterprise can make commitments to multiple providers without needing exclusivity. And when this happens, the savings compared to on-demand pricing is substantial – for our small basket, the average savings is around 50%; for our large basket, savings of up to 65% can be made versus on-demand expenditure.

Of course, this misses out on a lot of complexity tied to managing these multiple clouds: skills to manage and operate, platforms that work between providers, and, of course, the cloud-native frameworks that can enable such decoupled applications – but for savings of up to 65%, surely there is an opportunity to be capitalized on. And this cost-savings ability is in addition to the value enabled by being able to use multiple providers.

Figure 15: Multi-Cloud Savings Achievable Over Average Single-Cloud Price

Source: 451 Research's Cloud Price Index: Economics of Hybrid and Multi-Cloud, Q4 2019





4. Cloud-Native and Containers

Containers and other cloud-native technologies are evolving to take advantage of the on- and off-premises resources embodied by hybrid cloud. The function of containers is to decouple applications from the underlying infrastructure, creating lightweight runtimes that can be deployed where they make the most sense for reasons of cost, resiliency or proximity to data stores or other difficult-to-move resources. Containers alone enable higher utilization of infrastructure by packing application components into a smaller form factor and minimizing the ‘hypervisor tax’ required when spinning up underlying VMs.

Portability Is Critical to Cloud-Native Apps

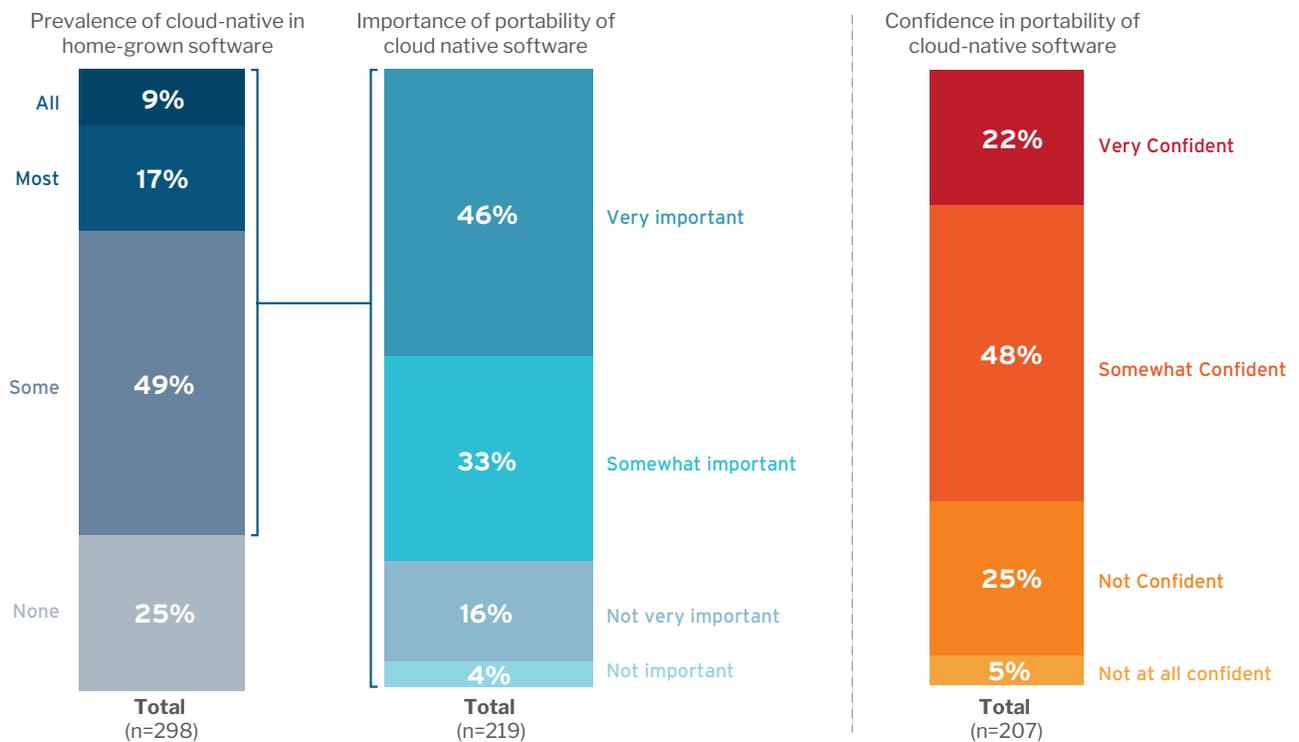
As detailed in 451 Research’s 2019 VotE: Cloud, Hosting & Managed Services, Workloads and Key Projects survey, cloud is the primary platform for new software development, and the drive for application portability is a likely contributor to the appetite for hybrid cloud, for which the top use case is ongoing workload migration. Among businesses building their own software, 75% indicate that at least some is cloud-native and that the ability to move applications between environments without significant refactoring is important; 46% consider it very important (see Figure 16).

Figure 16: Importance of Cloud-Native Portability

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019

- Q. How much of the software developed internally at your organization is cloud-native or cloud-enabled?
- Q. How important is it that the cloud-native/cloud-enabled software your organization develops is portable and able to run in other environments without significant refactoring?
- Q. How confident are you that the cloud-native/cloud-enabled applications your organization has developed will be able to run in other environments without significant refactoring?

PORTABILITY CRITICAL TO CLOUD-NATIVE APPS



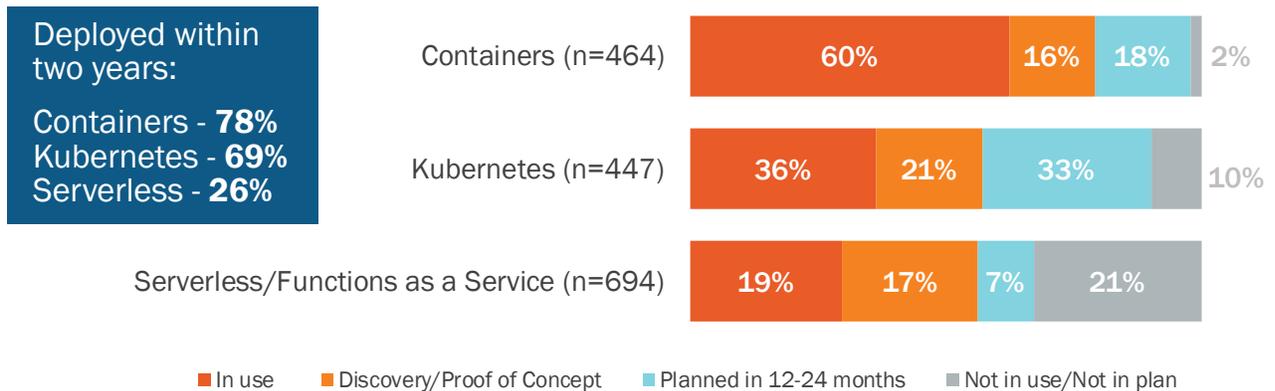
Kubernetes lit a fire under container adoption by creating a distributed framework for building, scaling and managing containerized applications. Implementing Kubernetes is not for the faint of heart, but all the major cloud providers have created their own managed Kubernetes platforms to ease adoption, and a large ecosystem of open source software has grown around it. As an orchestration standard, Kubernetes creates a hybrid/multi-cloud development and management layer that abstracts away the underlying servers, effectively moving the demarcation between infrastructure and other layers up a level, from the operating system to the application. Kubernetes serves not only as connective tissue between infrastructure and applications, but also as a conceptual tool to help companies figure out how to approach digital transformation: Rather than focusing on infrastructure (either physical or virtualized) and adapting software to run on top, it shifts IT attention to application-driven infrastructure.

Adoption of Cloud-Native Technologies

Nowhere is this service-oriented view of IT clearer than in serverless computing. The promise of serverless applications – in which reusable functions and triggers are assembled into software that works independently of the infrastructure that executes it – has attracted a wave of startups, open source projects and cloud providers because the potential benefits are too compelling to ignore: faster development, hands-off provisioning and dramatically lower costs. Rather than having VMs sitting idly by – burning cash, compute resources and energy – these applications invoke compute resources only when needed, operating on a pay-per-use rather than pay-per-provision basis (as virtualized hardware does). Although the technology for building and operating serverless applications at scale still has rough edges, partly due to the lack of an open standard, some born-in-the-cloud companies are adopting a ‘serverless-first’ strategy given the favorable economics and speed of development that it makes possible.

Figure 17: Adoption of Cloud-Native Technology

Source: 451 Research’s Voice of the Enterprise: DevOps Q1 2019 and Digital Pulse, Budgets and Outlook 2019



The cloud-native ‘run anywhere’ mentality is not without challenges – storage, networking and security are big issues that companies are wrestling with as they attempt to optimize applications that are distributed by nature. Yet the tools for doing so safely and at scale are being refined, and software giants with embedded customer bases such as Microsoft, IBM, VMware, Red Hat and Oracle are reconfiguring their products (and building partnerships) to help usher in this next wave of IT transformation.

One technology that aims to address issues of authentication and networking within cloud-native applications is service mesh. As large enterprise IT environments become more heterogeneous and smaller chunks of code (i.e., microservices) are deployed, service-to-service communications and authentication become more difficult. A service mesh detaches the control plane (which applies policy and configuration instructions to traffic entering and traversing the application) from the data plane (which is a programmable network of proxies that observes, authenticates and routes service traffic based on the control plane instructions). Envoy, the data plane developed at Lyft, has achieved impressive take-up since it was open-sourced at the end of 2016 – it is the default data plane for many service meshes including Istio, a Kubernetes-native open source project launched in 2017 by Google, IBM and Lyft.



5. Let It Be

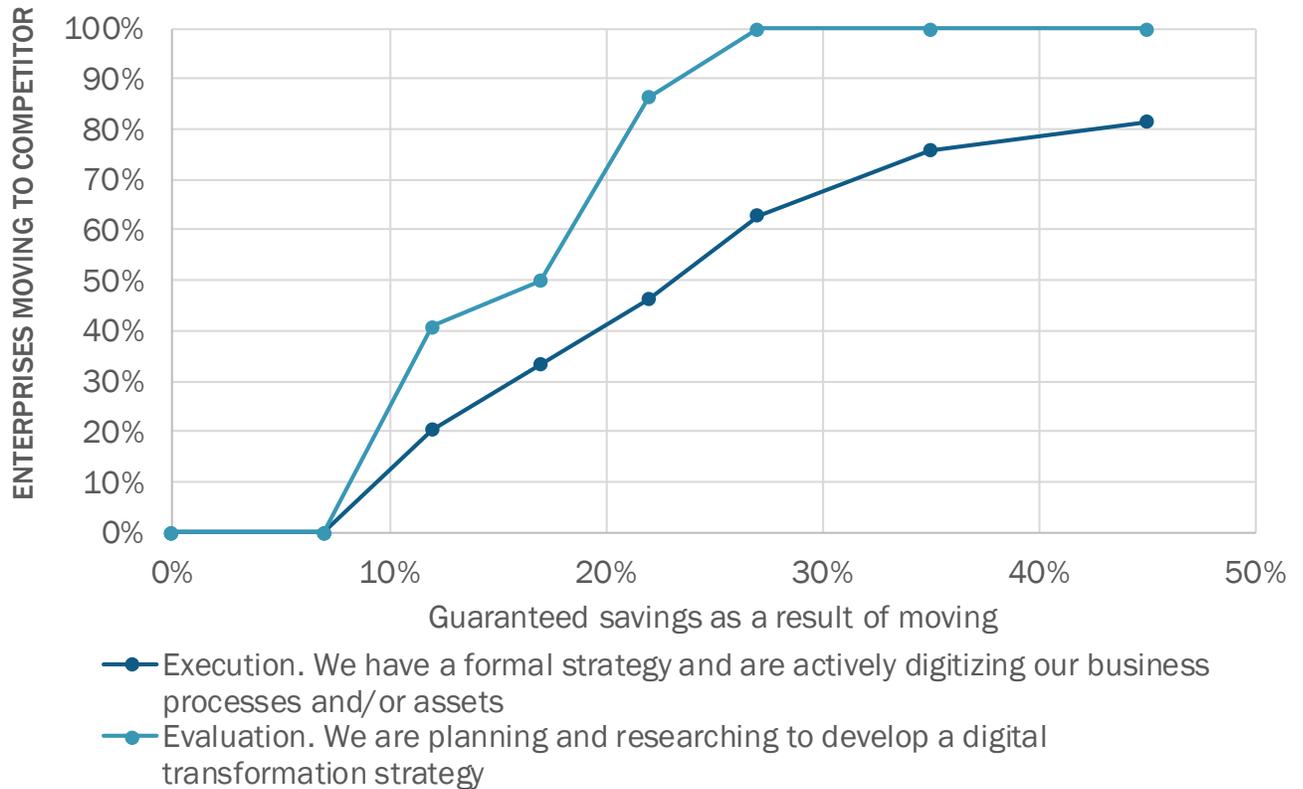
For many service providers and technology vendors, hybrid and multi-cloud are risky strategies – the risk is that enterprises move workloads increasingly out of providers’ or vendors’ remit. A lack of lock-in means fluidity, and workloads move to public cloud where reselling margins are thin instead of the more value-adding managed private or public cloud.

Our advice is to embrace it. Locking down customers or restricting options won’t win any friends, it might just alienate customers. Give options but make the preferred option is compelling enough that it dominates. And remember, data has gravity. Once the preferred option is in use, it’s more likely to be consumed on an ongoing basis. Lock customers in not through technology, but through a positive experience and ongoing usage. Give credits and rewards to encourage take-up, throw in capabilities and services, guarantee what public cloud providers won’t, and enable experimentation.

Figure 18 shows the percentage of enterprises (vertical axis) that would move venue for a guaranteed discount (horizontal axis). Incumbent providers are far stickier: At a 27% discount, for example, all surveyed enterprises that are evaluating clouds would move venue; just 60% of those that have currently deployed cloud and are executing on a strategy would move.

Figure 18: Price-Sensitivity by Adoption Stage

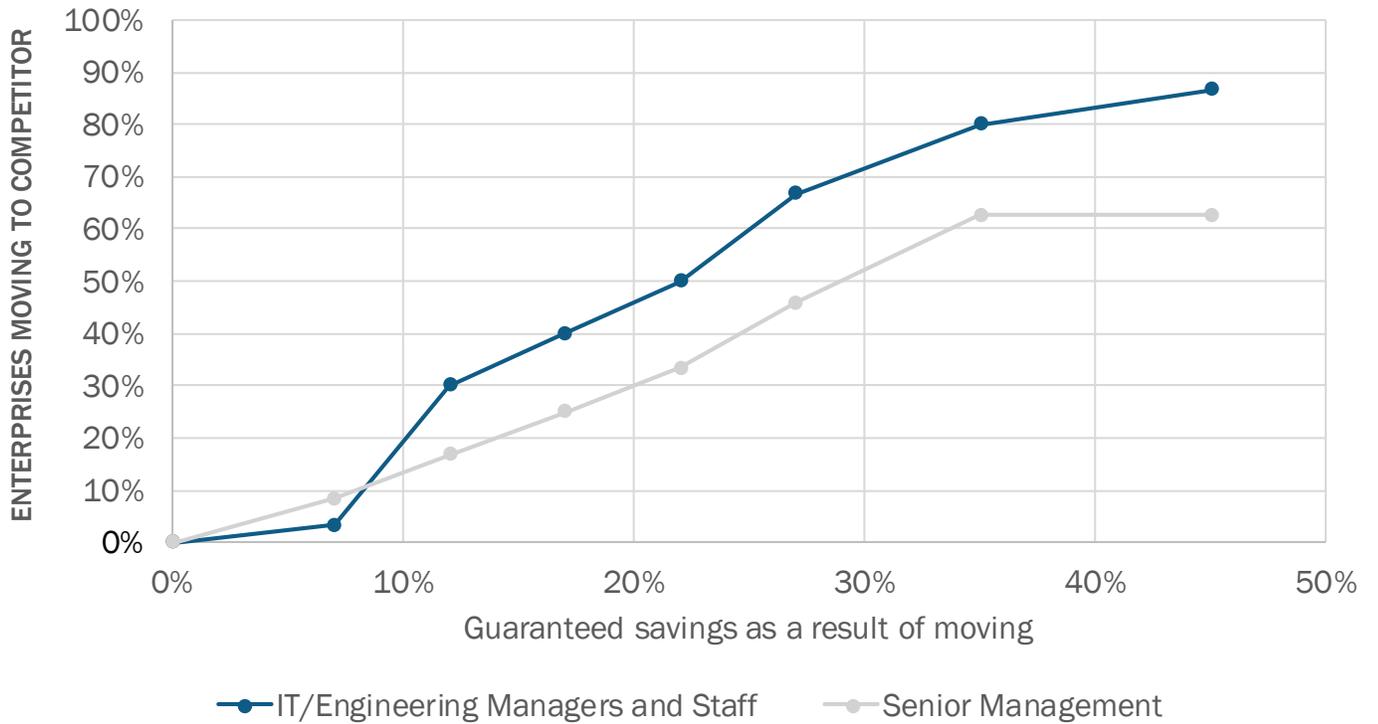
Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets & Outlook 2019



Senior management are keener to stay put than IT management. The C-level and senior management are more interested in cloud used to deliver value, whereas perhaps IT management are under increasing pressure to cut cost. At a guaranteed savings of 45%, three-fifths of decision-makers in senior management would move, while nearly 90% of IT managers would jump ship (see Figure 19).

Figure 19: Price Sensitivity by Buyer

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets & Outlook 2019





6. Conclusions and Recommendations

Recommendations for Service Providers and Technology Vendors

- **Recognize that hybrid on- and off-premises cloud will be the default for most enterprises for the foreseeable future.** Make mutual visibility (if not interoperability) a priority to avoid stranding workloads. Even if you have a preference for the model you most want to sell, enable all the options but build out the proposition around your preferred venue such that enterprises are drawn to it more than the others, through cost, bundled features or guarantees.
- **Honest assessment of customer environments is critical to understanding the best execution venue for individual applications.** This is not a one-and-done proposition: performance, cost and portability all factor into the equation, and those metrics (and the workloads themselves) change over time. Add ongoing value by constantly improving the technology and its value to the business.
- **Understand factors that can weigh in favor of keeping some operations on-premises,** including existing hardware and facility investments, data sovereignty requirements and limited resources for refactoring applications to run effectively off-premises. Project when customers are likely to reach an inflection point for updating their infrastructure.
- **Remember that companies that build their own software prize application portability** to ensure they're not beholden to any one vendor or venue for hosting their intellectual property. This 'run anywhere' philosophy has its challenges, but providers and vendors able to meet customer needs without locking them in to a given platform can earn loyalty and trust. Your ideal objective is that customers should want to stay because they are happy, not because they are trapped.

Recommendations for Enterprises

- **Focus on two practices that will go a long way in minimizing the total cost of operations for hybrid IT environments:** maximizing utilization of private cloud infrastructure and setting up and scheduling public cloud resources to scale with fluctuating demand. Press your cloud or managed service provider for help in improving these measures; they will also need to support you in this on a regular, formalized basis.

- **As with any portfolio, IT and cloud investments must be periodically rebalanced** – for purposes of technology refresh, access to new capabilities or vendor management – to maximize business benefit. Do this with eyes open for opportunities to improve performance, reduce cost and move the business forward with newly available capabilities.
- **Embrace open source solutions.** Especially for cloud-native development, open source options offer a low-risk way to try user-vetted tools, influence and contribute to roadmaps and benefit from community engagement at the leading edge of software development.
- **Explore proofs of concept for cloud-native development where appropriate.** Good candidates include net-new applications to leverage company data to assist in business decision-making, optimizing IT resources to minimize cost and increase performance, and front-end tools and adjustments to improve the customer experience.



7. Further Reading

Voice of the Enterprise (VotE): Workloads and Key Projects 2019

Cloud Price Index: The Economics of Private Cloud, September 2019

Voice of the Enterprise: DevOps, 1H 2019

Voice of the Enterprise: Digital Pulse, Budgets & Outlook 2019

The Impact and Evolution of Cloud Native, September 2019



About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

© 2019 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such. 451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.



NEW YORK

Chrysler Building
405 Lexington Ave,
9th Floor
New York, NY 10018
P 212-505-3030
F 212-505-2630



SAN FRANCISCO

505 Montgomery,
Suite 1052
San Francisco, CA 94111
P 212-505-3030



LONDON

Paxton House
Ground floor
30 Artillery Lane
London, E1 7LS, UK
P +44 (0) 203.929.5700
F +44 (0) 207.657.4510



BOSTON

75-101 Federal Street
5th Floor
Boston, MA 02110
P 617-598-7200
F 617-428-7537