451 Research® | Advisory

Best of both worlds:
Can enterprises achieve both scalability and control when it comes to cloud?

NOVEMBER 2016

SPONSORED BY

Hewlett Packard
Enterprise

ABOUT 451 RESEARCH

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK	SAN FRANCISCO	LONDON	BOSTON
20 West 37th Street	140 Geary Street	37-41 Gower Street	125 Broad Street
3rd Floor	9th Floor	London, UK WC1E 6HH	4th Floor
New York, NY 10018	San Francisco, CA 94108	**P** +44 (0)20 7299 7765	Boston, MA 02109
P 212 505 3030	**P** 415 989 1555	**F** +44 (0)20 7299 7799	**P** 617 261 0699
F 212 505 2630	**F** 415 989 1558		**F** 617 261 0688

Best of both worlds: Can enterprises achieve both scalability and control when it comes to cloud?

Executive Summary

Very little data is available on how effectively enterprises are managing private cloud deployments in the real world. Are they doing so efficiently, or are they facing challenges in areas such as performance, TCO and capacity? Hewlett Packard Enterprise commissioned 451 Research to explore these issues through a survey of IT decision-makers and data from the Cloud Price Index. The key findings of that analysis are:

Many enterprises demand the security and control benefits associated with on-premises deployments but struggle to efficiently manage the environment; public cloud is seen as a more flexible and scalable alternative.
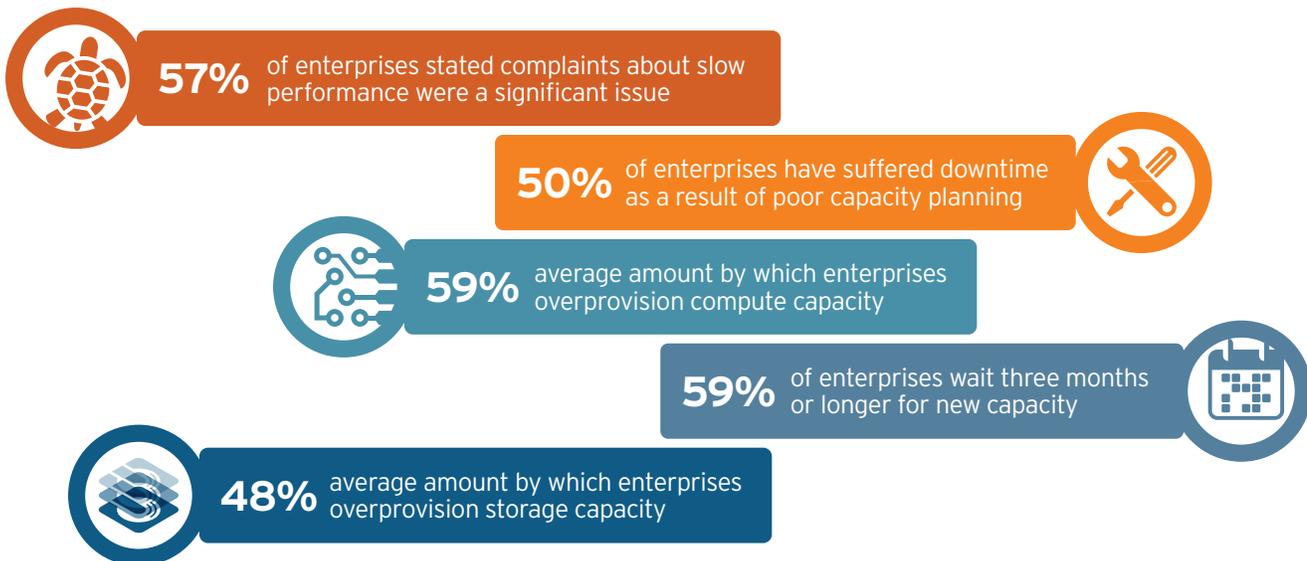
Enterprises struggle to balance private cloud capacity such that spending is kept to a minimum while providing end users access to all the resources they need.

This is a challenge that impacts the business performance and even the bottom line: Half of enterprises surveyed experienced downtime as a result of capacity-planning issues; 57% said complaints about slow performance were a significant issue.

As a result of this risk of underprovisioning, most enterprises play it safe and overprovision – by 59% for compute on average, and 48% for storage. This has a significant financial impact.

With 59% of enterprises waiting three months or longer for new capacity, it is clear to see why enterprises plan such a large overprovisioned capacity in advance.

HPE's Flexible Capacity helps resolve these challenges by providing on-premises infrastructure on a flexible and scalable basis. A TCO analysis of a typical enterprise scenario found it to be on par with public cloud and 29% less expensive than a self-managed private cloud.

57% of enterprises stated complaints about slow performance were a significant issue

50% of enterprises have suffered downtime as a result of poor capacity planning

59% average amount by which enterprises overprovision compute capacity

59% of enterprises wait three months or longer for new capacity

48% average amount by which enterprises overprovision storage capacity

The On-Premises Challenge

INTRODUCTION

Cloud has truly sparked a revolution in the IT industry and beyond, and from news headlines, analyst reports and vendor blogs, it's easy to assume that, one day, all workloads and applications will be situated in a handful of datacenters around the world. But anecdotally, that is not what we are hearing. Enterprises want a range of best execution venues to meet their differing needs. For some workloads, public cloud is the key. For others, particularly those with security or compliance concerns, on-premises dedicated equipment holds the answer. These concerns aren't unreasonable – constantly changing data-protection legislation, shifting borders and unions, and the real threat of data leakage and hacking explain why many aren't comfortable with going 'all in' on the public cloud. Enterprises want the best of both worlds – a hybrid strategy.

The big benefits of public cloud are flexibility and scalability – not just being able to grow to meet changing demands and capture opportunities, but also being able to shrink when demand drops to save money. The provider handles the capacity planning, and the provider holds the risk of underprovisioning and overprovisioning capacity. End users just consume what they need.

With on-premises deployments, the enterprise must handle these issues, buying or leasing equipment well in advance to ensure capacity is available without being wasted. On-premises infrastructure has the benefits of control and ownership, but also the responsibility of capacity planning. Public cloud has the benefit of no capacity planning, but control and compliance issues are still very much a concern.

But how well are enterprises doing this capacity planning? Very little is known about their capacity planning and management performance today. Are they managing on-premises infrastructure so well that it has economic advantages over public cloud, or are things so complicated that private cloud is a premium service?

Hewlett Packard Enterprise (HPE) commissioned 451 Research to perform a study of these issues through a survey of 500 IT decision-makers in the US.

RESULTS AND ANALYSIS

Aggregate Profile of Respondents Interviewed

TITLES	C level, IT vice president, manager/director, IT operations, systems administrator/engineer, IT administrator/storage administrator, system architect
INDUSTRIES	Transportation, government, insurance, wholesale trade, energy/utilities, communications, media and services, education, healthcare, retail/hospitality, professional services, financial services, manufacturing/construction, technology
REGION	United States
CONDUCTED	October 2016

ON-PREMISES EQUIPMENT IS DESIRED, BUT ITS MANAGEMENT PRESENTS BUSINESS RISKS

Enterprises see value in an on-premises approach: 68% and 52% of respondents rated control and ownership as key factors in their choice of a private cloud solution, reflecting the complexity associated with compliance and security in the public cloud domain. In fact, 90% of respondents stated a requirement for keeping some workloads on-premises, and that this was moderately or significantly impactful to the IT department.

Figure 1: Advantages for Enterprises in Operating Datacenters

Q. What does your organization perceive to be the top two advantages of operating in your own data center?

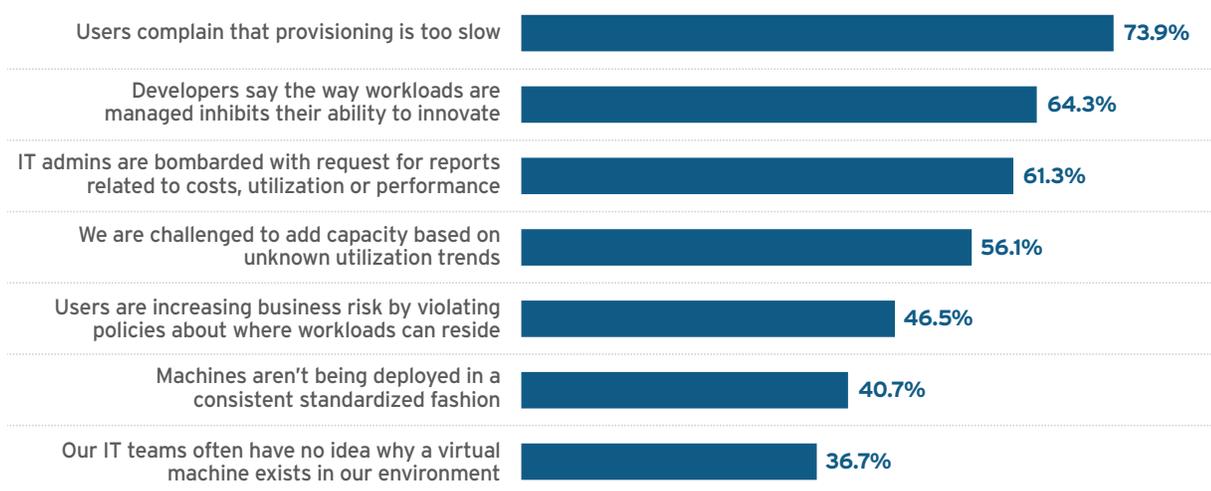| Advantage	Percentage
Control	68%
Ownership	52%
Business Processes	44%
Resell Value	15%
Accounting Practices	12%
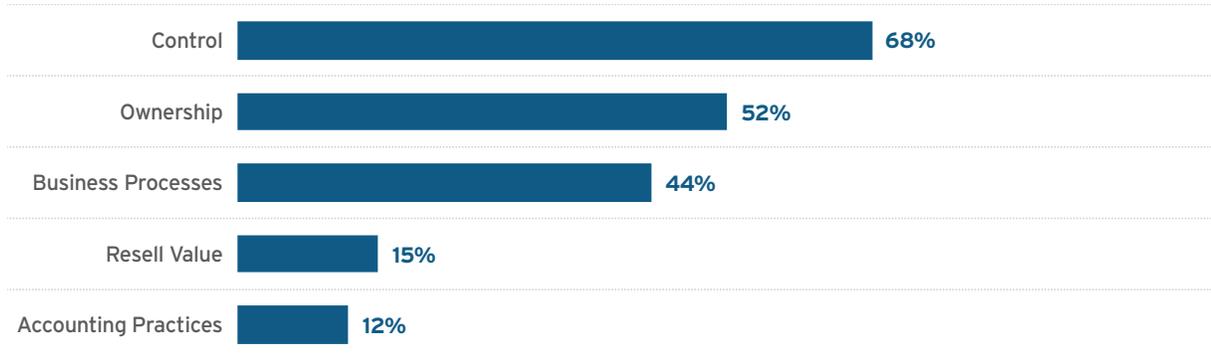
Source: 451 Research

End users faced practical challenges in their use of cloud, suggesting that IT departments are struggling to effectively manage capacity. As Figure 2 below shows, more than 50% of respondents cited slow response times, compliance pressure, workload management, forecasting and provisioning, fluctuating capacity, and reporting and tracking as ongoing issues relevant to their organizations. It seems that enterprises have a desire and even a requirement to use private infrastructure, but they face pressures from end users and the overall business in terms of effective management.

Figure 2: IT-Related Capacity Management Challenges

Please indicate if the following IT-related issues are relevant to your organization.

Issue	Percentage
Users complain that provisioning is too slow	73.9%
Developers say the way workloads are managed inhibits their ability to innovate	64.3%
IT admins are bombarded with request for reports related to costs, utilization or performance	61.3%
We are challenged to add capacity based on unknown utilization trends	56.1%
Users are increasing business risk by violating policies about where workloads can reside	46.5%
Machines aren't being deployed in a consistent standardized fashion	40.7%
Our IT teams often have no idea why a virtual machine exists in our environment	36.7%

Source: 451 Research

These challenges are translating to meaningful threats to the business. A huge 93% of respondents said that product development was being moderately or significantly impacted by poor workload management; 91% said business risk was being similarly impacted due to poor controls on workload location; 82% said they had no idea why a particular virtual machine

451 Research® | NEW YORK ▪ BOSTON ▪ SAN FRANCISCO ▪ WASHINGTON DC ▪ LONDON

might be in existence. Crucially, rather than being just inconveniences, these issues have the potential to impact company revenue, reputation, stability and data-protection compliance.

When we asked about reasons for choosing public cloud, scalability through on-demand pricing was the major factor, cited by 54% as key. This is reflected in 45% who said they appreciate the ability to experiment with little risk. However, we were surprised to see 91% of respondents state that they were making commitments to public cloud providers – so although enterprises like being flexible, it appears they don't mind making some commitments if a benefit (such as lower costs or guaranteed capacity) is offered in return. Twenty-two percent like the capex-free approach associated with public cloud.
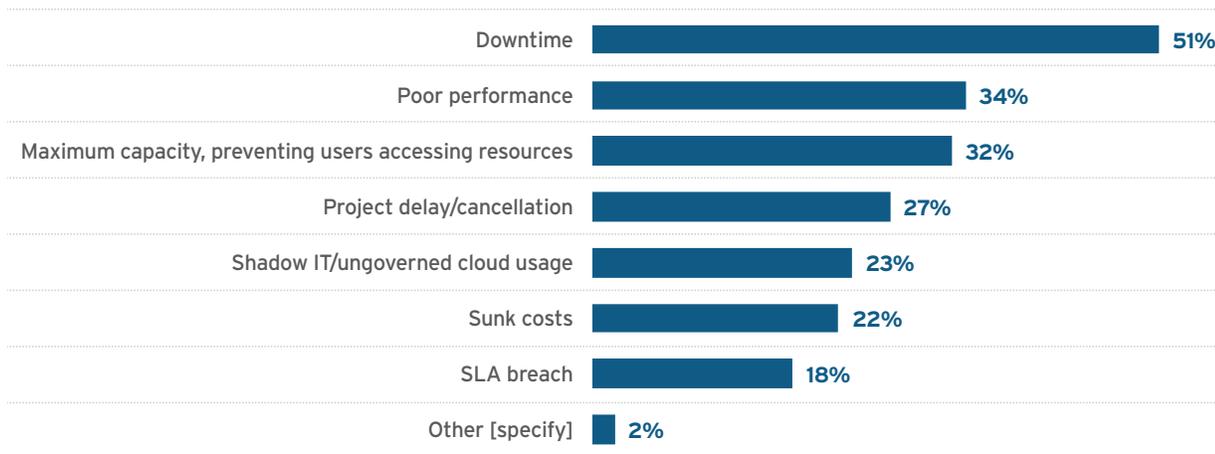
CAPACITY PLANNING IS A REAL BUSINESS RISK

In an ideal world, every CPU cycle, gigabyte of storage and even transistor bit in an on-premises infrastructure would be consumed. With everything deriving value, this brings unit costs down. If only four virtual machines are deriving business value out of a maximum of eight on a $1,000 server, for example, each virtual machine is effectively costing $250; if all eight virtual machines are being harnessed, the unit cost comes down to just $125. Ideally, this utilization needs to be as close to 100% as possible, but this is difficult in practice. Plan too little capacity, and end users might not be able to get access to resources when they are needed, meaning lost opportunities, poor performance and wasted time. Plan too much, and excess capacity might be unused and essentially represent sunk costs.

Our survey results show underprovisioning is a real risk (See Figure 3 below). Half of respondents experienced downtime as a result of not having enough capacity to meet demand; 34% suffered poor performance; 32% could not deliver resources to those whom require them; and 27% suffered canceled or delayed projects as a result. These are not minor inconveniences; they could result in customers shopping elsewhere, the loss of a first-to-market advantage, or missing the window to a potential new opportunity.

Figure 3. Underprovisioning Poses a True Risk to Enterprises

Q. Has poor capacity planning ever resulted in any of the following issues?

Issue	Percentage
Downtime	51%
Poor performance	34%
Maximum capacity, preventing users accessing resources	32%
Project delay/cancellation	27%
Shadow IT/ungoverned cloud usage	23%
Sunk costs	22%
SLA breach	18%
Other [specify]	2%

Source: 451 Research

So is it better to overprovision? On average, respondents are overprovisioning by 60%, effectively buying capacity that isn't even likely to be used. Six percent spend more than half a million dollars annually on excess capacity. Twenty-one percent buy so much capacity in advance they don't even need to do any capacity planning. A huge 89% of respondents state they are challenged to add capacity based on unknown utilization trends to a moderate or significant extent.

THE WAY AHEAD

With enterprises stuck between a rock and a hard place, it makes sense that we find average utilization of a compute cloud is just 58%. For storage, it is even less at just 56%. Of course, with better tools to forecast capacity, it might be easier to achieve high levels of utilization. In fact, our respondents believe they could improve utilization by 14-16% with such tools, and more than 70% rate such tools as attractive. However, there is still a big challenge in actually getting the hardware in place in time to be consumed. Fifty-nine percent of respondents wait longer than three months for new hardware – forecasting accurately with such a lead time is a challenge for anyone.

So, what does this all mean? In a nutshell:

- Enterprises want and need on-premises deployments, typically alongside a public cloud deployment. Both options have pros and cons depending on workload.
- There are significant business risks associated with the effective management of on-premises infrastructure.
- Many of these risks can be linked to capacity – too little, and revenue opportunities may be lost; too much, and money is being thrown away.
- Most decision-makers realize tools could improve capacity management, thereby optimizing costs and reducing business risks.
- But long lead times on hardware procurement can impact even the best forecaster of capacity.

HPE describes its Flexible Capacity program as giving enterprises the ability to consume on-premises resources in a pay-as-you-go model, while HPE handles the capacity and other facets of support and management. In Part 2, we compare the TCO of a private cloud built upon HPE Flexible Capacity against the Cloud Price Index public cloud and self-managed private cloud indicators.

Understanding the TCO Impact of Capacity

CLOUD PRICE INDEX

451 Research was commissioned to compare the TCO of a typical end-user self-managed private cloud and a public cloud against an HPE private cloud built on its Flexible Capacity program, using Cloud Price Index (CPI) and the aforementioned survey as sources. To ensure independence, none of those being surveyed were aware this project was commissioned by HPE, and HPE was not permitted to select a sample for analysis from the whole dataset.

The CPI is a service of 451 Research, and provides benchmark costs for a range of cloud services operating across public and private, commercial and open source, and self-managed and managed models, using an open, non-subjective, independent and publicly available model. It's very similar to the Consumer Price Index, which is essentially the price of a basket of typical groceries an average person buys every week. Our 'basket of goods' is a detailed specification of a cloud application, which allows us to collect data on market changes and the average prices of services within our cloud application. A more complete breakdown of the Cloud Price Index methodology can be found in the Appendix.

DERIVING A FLEXIBLE CAPACITY COMPARABLE TCO

451 Research was commissioned to compare the TCO of a typical end-user self-managed private cloud, but first we had to understand what constitutes a 'typical' deployment. Using CPI data and data from our survey group, we sought to compare TCO only for our medium virtual machine specification, which includes 1vCPU, 4GB memory and 150GB magnetic storage.

We asked our survey group for average levels of utilization and labor efficiency and found that current average utilization of a private cloud from valid responses was 58%. We then prepared a capacity profile (Figure 4 below), which matched this utilization over a three-year period, assuming linear growth from 0-2,000 virtual machines with fluctuations each month. It is impossible to simulate every demand profile (by definition, there is an infinite number), but a growing linear capacity with some fluctuations could certainly not be deemed 'unusual,' and we believe this is a fair enterprise experience. We suggest that enterprises considering Flexible Capacity perform their own analyses for their particular scenarios.

Per the standard CPI specification of a self-managed private cloud built on commercial software, we determined the average cost per virtual machine based on this profile, but where capacity was fixed at a maximum from day one. Twenty-one percent of those surveyed stated they planned capacity in this manner, provisioning so much capacity initially that subsequent capacity planning wasn't a concern, which validates our approach.

Figure 4: Capacity Profile Used in TCO Assessment
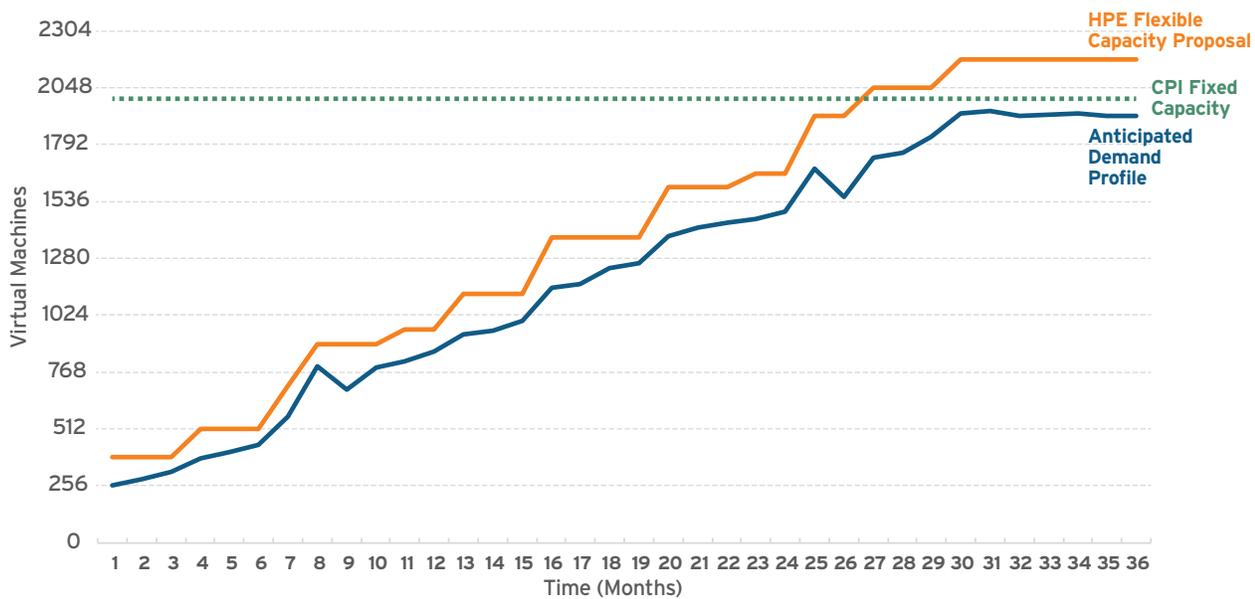
Source: 451 Research

We asked our survey group for their current labor efficiency needed to operate just the private cloud infrastructure, measured in terms of number of virtual machines under management per engineer. The average VM-per-engineer ratio was 464.

Based on CPI data from Q3 2016, the mean price per utilized virtual machine came out at $128 per month, or $0.13 per VM-hour, in scenarios when utilization is 58% and labor efficiency is 464 VMs/hour, and where full capacity is available from day one in a self-managed private cloud. See Appendix for an example.

HPE provided a quotation for delivering exactly the same profile using its Flexible Capacity approach. Most of HPE's Flexible Capacity customers choose a four-year contract, but we found in our survey that the majority of end users commit to three years when purchasing a private cloud. The CPI specification also calls for a three-year term. To provide a fair comparison, we analyzed the per-VM TCO of the first three years only of a HPE four-year contract.

HPE quoted a price of $0.034 per VM based on our specification, which includes hardware, software and services. The services element includes ongoing support to right-size the capacity for the customer, to procure incremental capacity when it is required, enhanced call handling, on-site 24/7 reactive care with a four-hour response time, and proactive services to make sure the customer's environment stays up and running. In a previous survey study conducted by an independent third party, HPE's Datacenter Care (which is included in Flexible Capacity) was found to provide an average 28% improvement on labor efficiency. We have used this figure as a conservative estimate of labor-efficiency improvements as a result of using Flexible Capacity (conservative because Flexible Capacity includes more management than was included in the sole Datacenter Care product). The full report is available from HPE.

A 28% improvement in labor efficiency means that where previously one engineer could manage 464 virtual machines, now 1-0.28 = 0.72 engineers can now manage the same number. So with this improvement, one engineer can now manage 1/0.72 x 464 virtual machines = 644 virtual machines.

The scaling profile provided by HPE indicated a 36% reduction in server requirements over the three-year period compared to the standard CPI self-managed private cloud architecture. To obtain a colocation fee for the HPE offering, we took total colocation costs to be (1-0.36) = 64% of the CPI specification colocation costs.

RESULTS

Based on the same specification and scaling profile, the table below shows the respective cost per virtual machine based on HPE's Flexible Capacity, a self-managed cloud with fixed capacity from day one (from standard CPI dataset), and the public cloud (from standard CPI dataset).

Figure 5: Per-VM-Hour Breakdown and Totals by Model

	HPE Flexible Capacity	CPI: Self Managed with Fixed Capacity from Day 1	CPI: Public Cloud
Colocation	$0.005	$0.009	
Hardware/Software	$0.034	$0.054	
Labor	$0.053	$0.068	
TOTAL	$0.092	$0.131	$0.092

Source: 451 Research

Figure 6: Per-VM-Hour Breakdown and Totals by Model
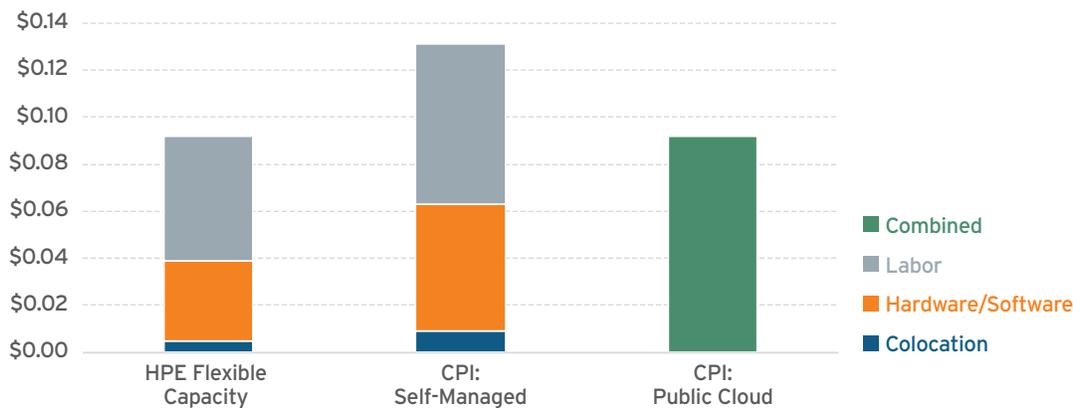
Source: 451 Research

The current CPI average price for our defined virtual machine size on a public cloud is $0.092, based on quotations and estimates from 50 cloud providers covering 90% of the global IaaS market. Rounded to three decimal places, the cost of HPE Flexible Capacity on a per-VM basis is the same as public cloud. This suggests that by using Flexible Capacity, enterprises could gain the benefit of retaining control that the survey has shown they value, without needing to pay more than public cloud.

The self-managed price is based on CPI data on public cloud from quotations provided by VMware and Microsoft as part of our commercial benchmark indicator. The indicator has been normalized to reflect the 58% utilization and 464 VMs/engineer parameters found in our survey. Here, we find that Flexible Capacity is 29% less expensive than the equivalent private cloud built using full capacity from day one and fully self-managed.

As a result of the ability to scale, Flexible Capacity has made a 36% cost saving on colocation compared to a fixed approach. As a result of bundled services, Flexible Capacity has made a 28% cost saving on management compared to a fully self-managed option. Since customers only pay for capacity when they are using it, the result is that compute capacity is 38% less expensive than the CPI self-managed private cloud average.

451 Research® | NEW YORK ▪ BOSTON ▪ SAN FRANCISCO ▪ WASHINGTON DC ▪ LONDON

Conclusions

In our derivation and analysis of a typical enterprise use case, a virtual machine delivered on a private cloud managed through HPE's Flexible Capacity program is roughly the same price as one delivered from the public cloud. Crucially, the Flexible Capacity program delivers the benefits of private cloud as viewed by our survey group, primarily control, while also enabling the benefits of public cloud as viewed by our survey group, namely scalability and low capex. From our analysis, it appears this 'best of both worlds' capability is not at a premium compared to public cloud.

Compared to a self-managed private cloud where peak capacity is provisioned from day one, a virtual machine delivered via Flexible Capacity is 29% less expensive. Because Flexible Capacity manages capacity, we believe utilization can be improved using the program by ensuring there is capacity available for new requirements without wasting money on unused servers. This could potentially reduce those risks and impacts identified by our survey group, including downtime, poor performance, delayed projects and wasted budgets.

HPE builds a buffer into capacity forecasts to account for unexpected demand, and regularly meets with customers to re-evaluate this capacity. We encourage customers and HPE to continue this dialog on an ongoing basis so that requirements can be constantly reevaluated based on changing demands. Our comparison is based on a typical enterprise use case, based on data collected from IT decision-makers, but we do make the assumption that the forecast is fairly correct. With a trusted advisor providing ongoing support for on-premises deployments, enterprises are in a better position to achieve lower TCO with greater benefits.

Appendix

The CPI is a service of 451 Research, and provides benchmark costs for a range of cloud services operating across public and private, commercial and open source, and self-managed and managed models, using an open, non-subjective, independent and publicly available model. It's very similar to the Consumer Price Index, which is essentially the price of a basket of typical groceries an average person buys every week. Our 'basket of goods' is a detailed specification of a cloud application, which allows us to collect data on market changes and the average prices of services within our cloud application.

By defining standard specifications of typical applications, it's possible to measure changes in the typical price of real-life scenarios. The CPI should be used to provide a benchmark to aid planning and decision-making. Enterprises, service providers, vendors, investors and partners should use the CPI, alongside other data and experience, to assess the cost and value of the IaaS, PaaS and SaaS they offer and consume. Value is the key term here – cheaper is not necessarily better. More expensive offerings are likely to be a better value if features, support and geography justify the additional expense.

COLLECTION AND ANALYSIS OF PUBLIC CLOUD DATA

To calculate the CPI, we use a weighted average based on market share. To capture the total cost of hosting our detailed specifications from a range of providers, we seek quotations directly from cloud providers and also derive estimates using online pricing and API interrogation. Providers are asked to submit quotations that deliver as close a match as possible to our required specification, using their cheapest datacenter in that region. Providers are rewarded with a sample of the research in return for their quotation, and we agree not to release any individual provider's quote – the data is only released as averages and anonymized distributions. The current summary specification for our 'large basket' application is:

- 2 x 'web server' medium virtual machine – 1vCPU, 4GB memory, 150GB magnetic storage
- 2 x 'app server' large virtual machine – 2vCPU, 8GB memory, 300GB magnetic storage
- Standard (non-premium) Linux operating system, e.g., Ubuntu, CentOS
- Total incoming internet bandwidth to application: 30GB per month
- Total outgoing internet bandwidth to application: 120GB per month
- 100GB object storage per month
- 8GB sized relational database instance, 20GB database index, 100 IOPS performance
- 20GB indexed NoSQL storage, 500 reads/sec, 250 writes/sec, 16GB memory
- 200GB load-balanced total inbound traffic

Each total cost is multiplied by that provider's market share percentage to produce a provider-weighted cost. Market share data is obtained from 451 Research's Market Monitor service. The weighted costs are then added together and divided by the total market share represented by all providers. In other words:

$$CPI = \frac{\sum_{i=1}^{n} s_i p_i}{\sum_{i=1}^{n} s_i}$$

where for provider i of n providers, s_i is market share ratio and p_i is the market price.

The equation gives us a value that is an approximate representation of the average price an end user would pay per hour for running that application (or for that underlying cloud service) across that percentage of the market.

We currently collect data from 50 public cloud providers, which together cover 90% of the global cloud market in terms of revenue, across four regions. Crucially, our approach doesn't just focus on hyperscalers or the new players – it is a mathematically valid estimate of the average price an end user would pay for that service based on our specification. Our dataset is refreshed every quarter, and the data used in this report is from Q3 2016.

COLLECTION AND ANALYSIS OF PRIVATE CLOUD DATA
Infrastructure Costs

In a similar vein, every summer, 451 Research takes the role of a hypothetical medium enterprise looking for a dedicated (single-tenant) hosted private cloud solution, that as closely as possible meets the following specification:

- ~500VMs, each with:
    - 1 vCPU (assuming 1 vCPU to 1 hyperthread)
    - 4GB memory (uncontended)
    - 150GB magnetic storage (uncontended)

- o 0.5Gb/s peak bandwidth, with likely 120GB transfer out per month and 30GB transfer in. It is unlikely this bandwidth will peak simultaneously across all virtual machines at the same time.
- 40,000GB capacity object storage (x3 replication)
- 24,000GB capacity block storage (RAID 10 or 3,000 IOPS)
- Orchestration tools to allow self-service and provisioning of compute, object and block storage plus:
    - o Identity management, e.g., to control roles and polices regarding self-service
    - o VM image management, e.g., to allow catalog of images to be maintained and provisioned to compute nodes
    - o API management, e.g., to allow provisioning of resources through CLI
    - o GUI, e.g., to allow provisioning of resources through web interface
    - o Network management, e.g., IP assignment
- 24/7 telephone support to the software vendor
- Access to hypervisor and orchestration updates and patches
- Single-site deployment. Controllers, network and power will be in dual configuration to increase availability.

The full specification above was sent to a range of service providers, which were asked to quote for delivering the requirements as a fully managed service (including hosting, hardware and cloud) over three years, as closely as possible using their standard services only.

We also asked providers to provide detail on how they'd deliver a 5,000-VM private cloud, and from this data, it appears that expansion is linear. In other words, a 5,000-VM offering is approximately the cost of 10x our 500-VM offering. This means that our findings should relate to private clouds architectures beyond our specification. HPE does not typically support Flexible Capacity for a use case of just 500 virtual machines, so we have scaled our data to match a 2,000-VM cloud based on these principles.

451 Research also designed its own private cloud architecture based on HPE standard servers and equipment, and colocated in the northeast of the US. In our architecture, we used servers with 16 2.2GHz cores running a total of 32 hyperthreads, 128GB of memory and a total 4.8TB storage capacity. Based on this specification, 32 of the aforementioned size of virtual machine should fit per node; at peak capacity, this brings total compute nodes to 16.

In our analysis, self-managed means the hardware and software are managed by the end user, but the datacenter is managed by a third party (i.e., colocation). We sought quotations from commercial software vendors, OpenStack distributions and others. This dataset gives us the ability to calculate the average total cost of the infrastructure needed to support a self-managed private cloud over three years.

Labor Costs

To bring labor into the mix, we have used industry-standard salary data extracted from salary-comparison website Indeed. com in August 2016. According to this source, OpenStack engineers cost $140,000, and VMware and Microsoft engineers cost about $100,000. Note that both of these average salaries have risen about 10% since our last Private Cloud Index 18 months ago, showing increased demand for cloud skills. A multiplier of 1.7 was used to capture the associated costs of benefits and workplace. A sensitivity analysis found variations in this multiplier had little or no effect on overall findings.

In our analysis, engineers are ultimately responsible for ensuring the private cloud is available to internal and external developers and end users, which may include detecting, diagnosing and resolving faults, installing capacity, configuring users and resource allocations, updating and patching hypervisors and cloud software, optimizing performance, and supporting developers with cloud-specific queries and configuration requests.

Example

For a 2,000-VM OpenStack source, self-managed private cloud, the total cost for hardware, software and hosting from the CPI dataset is about $2m over three years. If our 2,000-VM cloud is being managed at a ratio of one engineer for every 500 virtual machines, we need 2,000 ÷ 500 = four engineers. The average annual cost of a qualified engineer is about $100,000, so our direct labor cost over three years is $1.2m. These engineers need benefits, office space and administration, and an industry-standard multiplier of 1.7 is used to cover these costs, bringing total labor costs to $2m. This brings the total cost to run our private cloud to $4m over three years.

If the private cloud is used at full capacity over the three years, the price per VM is $4m ÷ 36 months ÷ 730 hours ÷ 2,000 VMs = $0.08. However, very few clouds run at full capacity – it's fairer to divide by the actual VMs that are deriving value. At 60% utilization, only 1,200 VMs on average are being consumed each month. This brings our final value to $4m ÷ 36 months ÷ 730 hours ÷ 1,200 VMs = $0.13.

Hewlett Packard
Enterprise

Why choose between on-premises and on-demand IT when you can have both?

Flexible Capacity combines the simplicity, agility, and economics of public cloud with the security and performance benefits of on-premise IT. You determine your own "right mix" of hybrid IT and workload placement—an approach to IT consumption and scalability that can help you address your desired business direction in the digital economy.

With its agile pay-per-use service, HPE Flexible Capacity can help your IT organization:

Avoid IT expenses stemming from overprovisioning

- Improve time to market by maintaining a safe buffer of capacity, ready for use when you need it
- Keep capacity ahead of demand with regular monitoring—and a simple change order to replenish
- Pay for only the capacity used, not the capacity deployed
- Reduce IT risk with HPE Datacenter Care

How HPE can help

With HPE Flexible Capacity, your organization gains the agility, pay-per-use billing, and rapid scalability of the cloud service delivery model. Surges in demand are easy to manage, so you get to market faster without wasting capital on overprovisioning IT capacity for peak loads. But the technology resides in your own data center, so you maintain control of security, data privacy, compliance, and performance.

With HPE Flexible Capacity, your actual metered usage is billed monthly as a service. This helps you track costs, and includes the infrastructure, software, and support needed to run on-premise workloads with confidence.

Include certain Microsoft® Azure services in your HPE Flexible Capacity service, and you gain:

- Pay-per-use for hybrid IT; the cost for your IT usage is based on actual metered consumption of resources on premise or in Microsoft Azure
- A single contract for your on-premise servers, storage, networking, operating software, and certain Microsoft Azure cloud services
- A single monthly invoice that includes usage of resources in your data center and in Microsoft Azure
- A unified usage portal that reports on both services used and capacity consumed, so you can understand your usage and plan for the future; it also provides the data you need to be an IT services broker for your organization
- An enterprise-grade support experience with HPE Datacenter Care, a part of Flexible Capacity that provides support and guidance for your IT environment—and can be extended to include Microsoft Azure, giving you one accountable partner for hybrid IT

Provide stable, reliable, efficient services

With HPE Flexible Capacity, you can stop worrying about your capacity: you have a compute resources buffer in place, ready to use when needed without lengthy procurement processes. You benefit from:

- Maximized business agility—increase capacity in minutes, not months, from the pre-provisioned buffer.

- Lower risk of underprovisioning and missed business opportunities. Our account support team works with you to monitor usage and suggests actions to replenish capacity before it is needed.

- Reduced business risk, because key workloads can be kept on premise to meet security, privacy, latency, compliance, and control requirements, complemented by public cloud services as appropriate.

- Optimized cash flows—no upfront capital needed—and monthly payments. Align your costs better to the business. Consumption-based cash flows provide a more accurate ROI measure, whether NPV, IRR, or payback.

- Enterprise-grade support with HPE Datacenter Care, making it simpler to operate your hybrid IT. With one contact for your IT environment, backed by extensive global resources and 24x7 coverage.

- Simpler transitions to new technology. HPE Synergy, for example, delivers a dynamic IT environment; and Flexible Capacity can provide the needed room to grow that environment, because you still only pay for actual metered use.