Flexible Capacity: A Scalable On-Premise Datacenter Platform with Public Cloud Advantages

March 2015

Managing and operating a datacenter entail working with a myriad of complex technologies that can be difficult to run efficiently and effectively. Decisions regarding equipment, systems design, support, and management of the overarching cost structure have far-reaching implications. To address these issues, enterprises should look for solutions that carry a predictable monthly spend based on what datacenters and workloads consume. Enterprises should measure such solutions with service-level agreements (SLAs) and performance matrices — these are measurement capabilities enterprises should evaluate when working with a service provider. Managing datacenter costs continues to be a challenge for many enterprises. At the same time, they want to simplify deployment and utilization of application workloads and improve agility. Datacenters need to provision IT seamlessly, accessing new capacity in minutes, not weeks as was the case with the typical datacenter model. IDC has seen enterprises shift from making capital-intensive hardware purchases to purchasing capacity as needed on a monthly basis. As a result, organizations no longer have large capex and unused capacity on their balance sheets.

This Technology Spotlight explores the trends affecting enterprise datacenters and discusses the role that HP plays in providing services that meet today's dynamic computing needs.

IT Complexity: The Impact on Managing and Supporting Datacenters

For today's enterprises, IT has become increasingly complex. Within the datacenter, gone are the days of information silos in which servers, storage, networks, and applications operated as separate entities, with information largely locked away. Datacenters increasingly are adopting virtualized and converged infrastructures to deliver solutions on what IDC calls the 3rd Platform — IT services centered on cloud, social, mobile, and analytics technologies. As information fuels the business, line-of-business (LOB) managers are placing unprecedented demands on IT for capacity. Meanwhile, the CFO continues to put pressure on IT to regulate costs, even as the CEO looks to IT to innovate and deliver solutions to market faster. To some degree, cloud models are setting the agenda and expectations among LOB managers who are demanding quick and inexpensive access to infrastructure and platforms. As a result, IT departments are under pressure to act quickly or risk losing control of the datacenter.

Managing these new environments requires a considerable amount of retooling among IT employees. Since most enterprises are asking IT departments to manage ongoing IT operations with fewer resources, IT managers are starting to look to support providers for assistance in improving overall performance while reducing and improving resource allocation in the IT environment. In today's competitive environment, IT is asked to innovate; spending time on mundane routine tasks no longer adds value to the organization. In recent IDC research, 55% of IT managers polled stated they would need to train employees on how to rapidly detect and remediate problems in a converged infrastructure.

This sentiment is reflected by attitudes regarding support. IDC research indicates that only 18% of respondents stated that they wanted to support their virtualized environment themselves; hence support providers have responded by altering packages and features to include new tools and utilities to help address the most pressing management and support issues. Capacity on demand is a great fit to help resolve these issues.

The Benefits of Flexible Capacity

IT managers have critical SLAs with internal and external customers that must be met. Identifying support paths and key contact points is critical. This involves knowing which support provider to contact when problems occur and the best way to engage with the provider.

Knowing which support provider to call is particularly important for datacenters with multiple workloads spread across many systems. In the event of an incident, the issue may reside virtually anywhere across the storage, server, network, or software infrastructure; uncertainty over the origins of an issue can drastically increase time to resolution. Flexible capacity can reduce this uncertainty by providing a knowledgeable single point of contact for all system issues. In addition to fewer downtime issues, flexible capacity helps datacenters achieve the following benefits:

■ Improve data integrity and query response times. With information residing close to the end user, the speed at which the information can be accessed increases dramatically without the costly overhead of data lines.

■ Accelerate the ability to scale environments so as not miss a business opportunity or a compute need. Excess capacity resides in the datacenter with no capital outlay, which enables quick deployment of resources.

■ Provide data points around SLA and key performance indicator (KPI) measurements for workloads that can be shared with the CFO and the CEO. Datacenter managers will be able to quickly and effectively let senior management know how a workload is performing and at what cost.

■ Eliminate larger capital investments (since flexible capacity is paid for with opex) as well as costs associated with sizing for peak loads or future growth (i.e., eliminating the need to have unused capacity on the balance sheet).

■ Regain control over workloads and eliminate rogue IT. With LOB owners placing all kinds of information out in cyberspace, this capability allows IT to act like a cloud provider without sacrificing security and governance.

In addition, flexible capacity:

■ Eases creation of an effective virtual environment if the datacenter is still in a siloed infrastructure architecture (single server per workload).

■ Allows for better patch and firmware management. Flexible capacity solutions help reduce the amount of time spent coordinating, deploying, and verifying patches and upgrades and provide a well-thought-out methodology for keeping systems up to date.
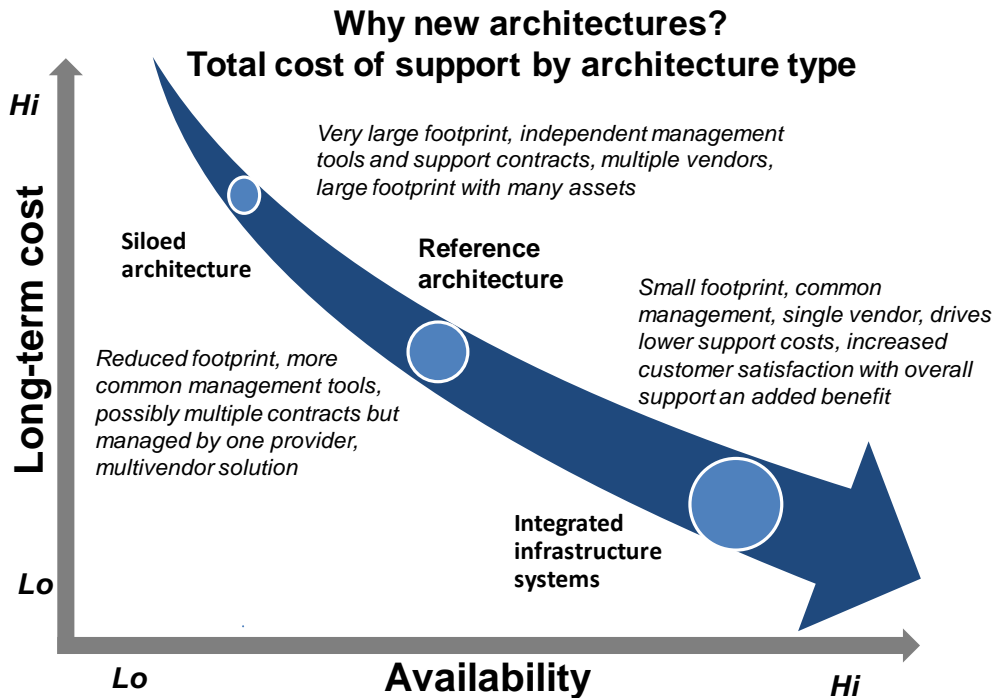
Flexible capacity solutions allow IT organizations to focus on driving the business forward instead of working on mundane IT tasks such as patching or system deployment. Datacenter managers can rest assured that their workloads are being served properly by these integrated infrastructures, freeing up valuable IT resources for other projects.

Market Trends

IDC has seen a dramatic uptake in the as-a-service model over the past few years. One of the most obvious user benefits of this service model is that it is presented as an all-in, "turnkey" solution: The customer can access the offering without the need to own, manage, or understand any underlying resources required to support the offering (see Figure 1). The service provider bears that burden, offloading day-to-day operational tasks of maintaining and provisioning capacity from the customer, making management much simpler while allowing for faster customer adoption. An important implication of these services being technologically "easy" is that many LOB organizations without IT skills have the ability to leverage capacity directly. Private "cloud" services provide customers with self-service capabilities for service provisioning and administration. In the IT cloud services world, the range of self-service capabilities varies widely up and down the stack. Customer self-service is a key tool for providing greater operating efficiency, deployment speed, and customer satisfaction.

F i g u r e 1

Support Costs and Architecture Availability

Why new architectures?
Total cost of support by architecture type

Very large footprint, independent management tools and support contracts, multiple vendors, large footprint with many assets

Siloed architecture

Reference architecture

Small footprint, common management, single vendor, drives lower support costs, increased customer satisfaction with overall support an added benefit

Reduced footprint, more common management tools, possibly multiple contracts but managed by one provider, multivendor solution

Integrated infrastructure systems

Hi / Lo — **Long-term cost**

Lo — **Availability** — Hi

Source: IDC

In traditional IT outsourcing, it is commonplace for a provider to offer customers a standardized portal for viewing of performance and other usage metrics as part of the contract. But this portal is more for ongoing high-level management than for true self-service provisioning as the capacity requirements in non-cloud outsourcing tend to be fairly predictable on a day-to-day basis. In dedicated private cloud, customers will expect more granular visibility over the workload as a function of "management" and true self-service via a portal when it comes to provisioning new workloads or making changes within the confines of the contract. For example, an application development team that wants a test server for a month could configure and provision the new instance within the broader outlines of an

existing contract (i.e., XYZ Corp. buys the capacity to set up to 10 computing instances, within these parameters, over the course of the quarter) rather than as a net-new contract requiring a new PO and service line to be opened. Customers are truly accepting this pay-as-you-go model.

IT departments can develop a pay-as-you-go model for their internal users and then take the integrated systems and IT a step further by becoming a profit center within their organization. Enterprises that can excel at running IT profitably in their industry may be able to provision resources for other organizations in their industry and charge a monthly fee.

Considering HP

HP Flexible Capacity is a scalable computing offering that is designed to provide customers with the cost benefits of a public cloud experience and the security and control benefits associated with on-premise IT (see Figure 2). The service supports servers, storage, networking, software, and converged systems. HP Flexible Capacity delivers the ability to pay as you go using opex rather than capex.

---

F i g u r e 2

HP Flexible Capacity

Elastic capacity, enterprise support, on premises

- Enterprise-grade support
- Pay only for what you use[1]
- Aligns costs with usage monthly via advanced metering
- Infrastructure capacity that never runs out
- Scalable – add capacity in minutes, not months
- Apply to servers, storage, networks, and software
- HP and multivendor
- Likely qualifies as OPEX

A Datacenter Care building block

[1] Subject to minimum capacity commitment

Source: HP
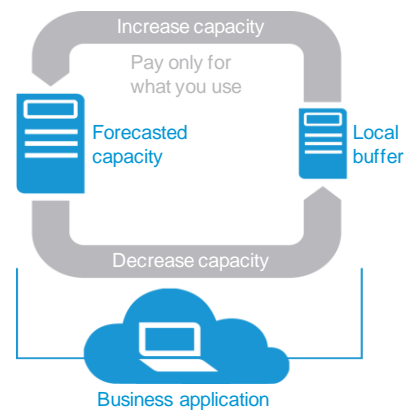
According to the company, HP Flexible Capacity is designed to deliver the following benefits:

- Elastic capacity that enables provisioning in minutes, not weeks

- Zero up-front capital investment with payment through opex

- Optimized cash flows

- Pay per use, match consumption

- Enterprise-grade service for operational efficiency

- Unlimited capacity with capacity refreshed as it is used

As a component of HP Datacenter Care, HP Flexible Capacity service delivers enterprise-grade support to the datacenter. This support includes a personalized onsite account team and a call experience delivered globally and backed by HP Centers of Expertise. In addition, HP Datacenter Care includes a mix of proactive and reactive services.

Service providers need flexible business solutions that can improve revenue streams and time to market. According to HP, HP Flexible Capacity's opex approach aligns costs with usage monthly because the bill reflects actual metered usage. This gives service providers the ability to expand or shrink costs with customers' usage, without using capital.

Challenges

Organizations are becoming comfortable with the as-a-service model, whether on- or off-premise. HP Flexible Capacity truly is a pay-as-you-go model as capacity needs increase or decrease over time. One of the major challenges that HP faces will be how to overcome the skepticism among potential customers who may perceive this solution as a modified lease. HP will need to educate the market that this is a pay-as-you-go model, with the true benefit of capacity on demand.

Conclusion

Today's datacenters need to be agile as well as cost effective. To accomplish these goals, many are turning to virtualized infrastructures that speed provisioning while eliminating the need to purchase excess capacity. To support these environments, resource-strapped IT departments are turning to third-party service providers. Ideally, these providers enable organizations to streamline troubleshooting, pay for what they use, and expand or reduce capacity — and costs — as needed. If HP can address the challenges highlighted in this paper, IDC believes the company has a significant opportunity for success in the important market for datacenter support services.