

REPORT REPRINT

HPE Flexible Capacity shoots for lower TCO with metered usage model

JEAN ATELSEK, OWEN ROGERS

21 FEB 2017

The company's Flexible Capacity contracts combine a fixed floor with a pay-per-use ceiling, aiming to merge public cloud economics with the control of dedicated on-premises infrastructure.

THIS REPORT, LICENSED EXCLUSIVELY TO HPE, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR REPOSTED, IN WHOLE OR IN PART, BY THE RECIPIENT, WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2017 451 Research, LLC | WWW.451RESEARCH.COM

Enterprises that keep compute workloads in-house still must deal with the time-consuming and financially risky process of planning for and procuring equipment well in advance of when it might be used. The result is often massive overprovisioning: purchasing or leasing a multiple of what's needed to ensure that the system can meet future increases in demand. This overhead comes at great cost, however, as stranded resources remain on the books with no measurable business benefit.

Hewlett Packard Enterprise's Flexible Capacity strives to combine the efficiency and savings of pay-as-you-go computing with the predictability and confidence of guaranteed capacity via a flexible services contract. Flexible Capacity starts with an installation suited for a business's current need, then reassesses requirements at least quarterly, adjusting server, storage and networking provisions with a change order rather than requiring a new procurement cycle. Usage over a pre-set threshold is billed on a metered basis, making it possible to achieve 100% utilization without the risk of shortages.

THE 451 TAKE

For businesses that must maintain on-premises infrastructure, HPE's Flexible Capacity offers a consumption-based model designed to alleviate the pain of procurement for dedicated compute, storage and networking services. By installing and maintaining resources based on near-term demand expectations, HPE shares with customers the financial and technological risk associated with unused infrastructure. Charging above a certain threshold based on actual usage increases utilization and lowers total cost of operation. The question is whether HPE can retain customer loyalty given evolving public and hybrid cloud models. HPE's recent purchase of Cloud Cruiser, which it has long licensed to underpin Flexible Capacity's metering and cost management functions, is a step in this direction.

CONTEXT

The runaway success of cloud computing models has left businesses that must keep certain workloads on-premises (for whatever reason) with what now appear to be antiquated and wasteful routines: capacity planning, demand forecasting, and expensive and painstaking capex outlays for equipment whose value depreciates while it remains unused.

HPE first launched a hardware pay-per-use model in 2002, and in 2008 it started allowing businesses to flexibly scale storage and compute resources, providing a buffer to ensure that customers have capacity provisioned ahead of demand, but charging only for what's used. The concept met with success, particularly in Europe, and in 2012 the program was formalized and rebranded as Flexible Capacity.

HPE purchased Cloud Cruiser in January, a cost management and optimization tool whose technology underlies Flexible Capacity's usage metering, to enhance its support for managing cloud deployments.

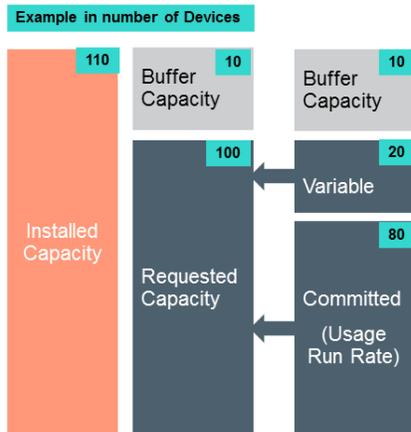
As 451 Research's Cloud Price Index research has quantified, utilization has a huge impact on the total cost of operations for any given infrastructure – the more an available compute or storage resource is used, the lower the unit cost. By shortening the capacity-planning horizon and charging only for what's consumed, Flexible Capacity enables customers to achieve greater utilization – 100% in some circumstances – and lower unit costs.

BUSINESS MODEL

Flexible Capacity is a service and partnership contract designed to share both the risk and benefit of provisioning for growing IT demand. For a customer – be it a midmarket company or a business unit of a large enterprise – acquiring resources with Flexible Capacity preserves capital and enables a better fit between availability and usage.

Here’s how it works: Customers commit to a minimum of resource consumption (typically 80% of requested capacity) based on expected near-term demand, and HPE installs equipment, incorporating a buffer of 10-40% beyond that request. The amount billed varies according to actual usage above the minimum on a daily basis. Compute resources are billed by server (based on whether a server is on), per physical core or per VM (based on allocated memory); storage and backup are charged per GB, and networking is billed per port used. HPE software, including the Cloud Services Automation suite for managing hybrid deployments, can also be included in the package.

Capacity Definitions



Definitions:

- Installed Capacity** – What is actually on site
- Requested Capacity** – Initial estimated usage over the term agreed to with Customer, may increase via change request
- Committed Capacity** – What Customer will pay for at minimum – typically 80% of Requested Capacity
- Variable Capacity** – Difference between Customer’s Requested Capacity and Committed Capacity. This capacity is only billed when used.
- Buffer Capacity** – Additional capacity deployed by HPE above Customer’s Requested Capacity.

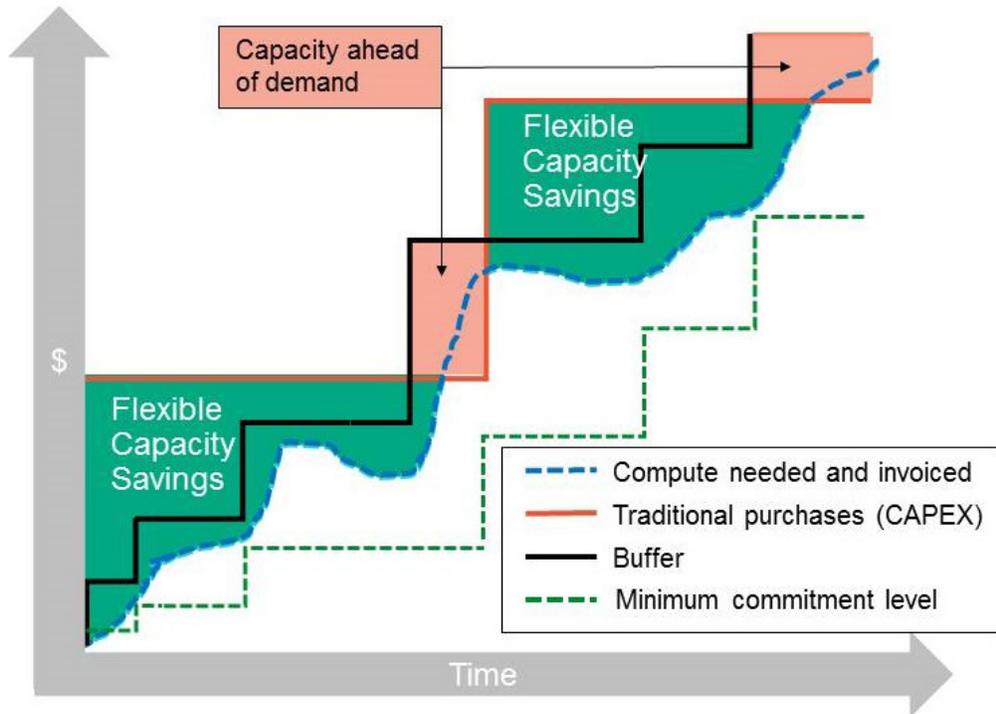


Source: Hewlett Packard Enterprise

HPE meets with customers at least quarterly – more often if a trigger point is reached (e.g., if utilization exceeds 90% of the buffer limit) – to review budgeted versus actual usage and agree on necessary adjustments. When new physical capacity is required, it can be set up quickly with a change order rather than requiring a new procurement cycle. The program also avoids technology risk – as a customer’s IT needs grow and resources are added, the caliber of the underlying hardware can be upgraded.

In effect, Flexible Capacity is exchanging the customer’s commitment to continue adding HPE equipment for HPE’s assumption of the risk that the capacity above the minimum will not be used.

How Flexible Capacity Works



Source: Hewlett Packard Enterprise

Using the Cloud Price Index plus data on private cloud management efficiency from an HPE-sponsored survey of 500 enterprise IT decision-makers, 451 Research found Flexible Capacity to be priced on-par with public cloud for a typical use case. Compared with a private cloud planned for peak capacity, Flexible Capacity was found to be 29% less expensive. Individual benefits depend on specific circumstances – most critically variability of workloads – but these savings are common-sense findings: Better capacity planning leads to less wasted resources, reducing sunk costs.

What's in it for HPE? Sustained revenue for one thing: The minimum contract term is three years, and four years is typical. But HPE says that the typical Flexible Capacity deal grows at 20% per year. According to the company, many customers start with just storage, then add compute and networking resources when they see how the model works. And HPE touts a Net Promoter Score of over 90% for Flexible Capacity, indicating that customers like what they're getting.

The program has about 300 customers, many of them service providers that need on-premises infrastructure but want to scale capacity in keeping with revenue growth. Other businesses are in demanding verticals such as healthcare and financial services, which require the low latency of a private network and the security and control desired for regulatory compliance.

In December 2015, HPE entered into an agreement with Microsoft, enabling Azure public cloud resources to be incorporated into a Flexible Capacity contract with one invoice and one portal for usage reporting. This option is designed for customers that anticipate the need for resources beyond what's provisioned in the buffer – for example, those with spiky test/development workloads.

COMPETITION

Flexible Capacity's value proposition is up against some tantalizing alternatives, including public and hybrid cloud choices, and traditional IT hardware purchases. Another option is leasing. Other hardware companies, including Cisco and Dell EMC, have financial services arms that offer equipment leasing for businesses of all sizes (as does HPE). Dell EMC launched a competing set of services at its Dell EMC World event in October 2016, including Flex on Demand, Datacenter Utility and Cloud Flex.

SWOT ANALYSIS

STRENGTHS

With pay-per-use experience going back 15 years, 300 customers and a high Net Promoter Score, HPE has a track record and credibility in capacity planning and managing on-premises deployments under this model.

WEAKNESSES

The target market for Flexible Capacity is narrow - customers that anticipate growing (but not bursty) IT demand, require on-premises infrastructure and are willing to commit to a minimum level of usage for three years or longer.

OPPORTUNITIES

HPE's partnership with Microsoft Azure to offer hybrid private and public capacity billed and managed through a single interface could be extended to other hyperscale providers.

THREATS

Competing vendors are jumping on the pay-per-use hardware bandwagon, most notably Dell EMC.