SIMPLIFY ENTERPRISE HYBRID CLOUD COST MANAGEMENT WITH HPE ONESPHERE

CONSOLIDATED VISIBILITY & CONTROL OF COSTS ACROSS CLOUD ENVIRONMENTS

ENTERPRISE HYBRID CLOUD MANAGEMENT CHALLENGES

As enterprise cloud adoption has grown, cost management has emerged as a considerable challenge due to gaps in visibility and accountability across the organization for spend and resource utilization. Common use of multiple public clouds, combined with private infrastructure for hybrid cloud capability, amplifies this challenge. The increased complexity of cost management in this hybrid cloud approach is similarly true in other functional areas, including ease of use and security management.

Growth in cloud adoption is being driven by enterprise digital transformation initiatives to address digital disruption through improved products that enhance customer engagement and experiences. Critical to these initiatives is cloud transformation designed to deliver speed and agility in application development to improve innovation and time to market. The struggle is achieving these benefits in a hybrid cloud approach without sacrificing ease of use and management of cost and security.

The fundamental issue is that each cloud infrastructure has its own toolset based on each vendor's focus on maximizing the value of their own cloud platform, not the hybrid cloud outcome enterprises are seeking. Enterprises face a fragmented set of platforms and tools that inhibits consolidated hybrid cloud management of:

- **Infrastructure costs** through visibility and control
- **Developer enablement** for rapid application development and deployment
- **Self-service infrastructure** for low operational burden on IT

HPE OneSphere offers enterprises a new hybrid cloud management platform provided as a service, enabling IT to address these challenges by empowering application developers and line of business leaders with easy-to-use, self-service tools for cost management and application deployment across all of their public and private cloud infrastructure environments.

This paper focuses on the enterprise challenge of delivering strong visibility and control of hybrid cloud utilization and spend. It is one of a three-part series on these enterprise hybrid cloud management challenges and HPE OneSphere as a solution to consider.

CUSTOMER CHALLENGES WITH HYBRID CLOUD COSTS

Enterprise IT originally perceived cloud computing to be primarily an infrastructure cost reduction lever allowing more tightly fitting capacity to application needs. Common examples were application scenarios benefitting from elastic, on-demand capacity in public cloud for temporary seasonal usage spikes (e.g., online retail) or unpredictable usage loads (e.g., online game launches). Startups adopting cloud to avoid the early expense of purchased infrastructure found and demonstrated to enterprise IT that cloud infrastructure is more valuable for the agility it provides the business to accelerate innovation and growth than its cost efficiency benefits. Speed of application development became the top focus in cloud adoption for improved innovation through easier experimentation and reduced time to market.

A decade into the adoption of cloud computing, many enterprises have now made such strong progress in capturing the benefits of agility that they have turned their focus to optimizing their infrastructure cost efficiency. This stems from both a desire to realize the originally expected cost benefits, as well as to solve widespread challenges with cost overruns in public cloud. These cost overruns occur for a variety of reasons, including:

- Budgeting for ideal capacity without enough allowance for uncertainty,
- Forecasting higher infrastructure utilization than achieved,
- Under-anticipating development and test environment needs beyond production,
- Not accounting for smaller costs beyond compute and storage; e.g., data transfer, load balancing, and application services,
- Resources not being de-provisioned when no longer needed, and
- Use of higher cost platform services instead of infrastructure services; e.g., using proprietary public cloud data store services instead of running open source software on compute and storage services.

While each is understandable as part of maturing cloud use, attacking each would not address the fundamental underlying problem, which is that application developers commonly lack visibility and accountability for cost-efficient infrastructure use. Application development leaders see a bill per public cloud infrastructure provider by cloud account that often spans many application teams and individuals. While some organizations progress to systematic use of metadata via tagging to categorize costs, even that approach is fragmented into per-cloud provider efforts.

With private infrastructure, the challenge is similar to public cloud when operating in a private cloud implementation across a large organization. However, more common in the larger enterprise infrastructure footprint is a virtualized resource fleet that is allocated between organizations and down to application teams but lacks visibility into resource utilization. Therefore, there is a lack of accountability for driving higher utilization for cost efficiency. The struggles with utilization in public cloud bring added attention to this in private infrastructure, versus in the past when "allocated" by IT to an application team was often considered fully "utilized." For example, application team experiences with virtual machine provisioning taking days-to-weeks encourages reporting of resources as fully utilized to avoid releasing them back to IT, which further strands capacity. Without aggregate cost visibility across infrastructures with allocation by organizational unit from business unit down to individual application teams, leadership has a hard time driving accountability.

It is possible to address these challenges with tooling available per infrastructure. However, there are two cautions to that approach. First, it results in management complexity with a different tool chain to train on and manage per cloud platform for the more than 90 percent of enterprises expected to use multiple cloud services and platforms by 2020[1]. This is further challenging in private infrastructures where there is often a lack of conversion from resource usage to cost, so the resources appear to be free to the application team. Second, it fails to build towards the strategically desired infrastructure sourcing model of the future that provides flexibility of application deployment across infrastructures. The vision for cloud platforms is that applications are developed and deployed to one of multiple available cloud infrastructures based on best fit cost efficiency, availability, features, security, and compliance. It is reasonable to expect that future platform capabilities will afford the flexibility to automate the migration of many types of applications across cloud infrastructures with zero-downtime purely for cost savings opportunities.

There will be exceptions, for example, when dependencies are taken on specifically valued features like proprietary cloud services (e.g., AWS Redshift for data warehousing is a popular example), but those represent a minority of an enterprise application footprint. An enterprise may also choose to aggregate a dominant portion of its application footprint on a single cloud infrastructure to gain optimal pricing at scale from a vendor – as is naturally the pitch from the cloud providers – but that does not absolve the enterprise from achieving flexibility across infrastructures to maintain future pricing leverage against the vendor. As cloud vendors shift in availability, feature and cost

---

[1] IDC FutureScape: Worldwide IT Industry 2018 Predictions, October 2017

competitiveness, no development or IT leader wants to have to explain to business leadership that they lack flexibility to re-allocate their application portfolio across providers to take advantage.

Development leadership has typically led cloud infrastructure and platform tooling decision-making in taking responsibility for achieving the benefits from agility for the business, but as focus turns to challenges of cost efficiency through visibility and optimization across infrastructures, they want to look to IT to take responsibility for enabling them and their line of business leaders in IT's traditional cost management role. The challenge for IT is delivering these hybrid cloud cost management capabilities in a budgeting environment where they are often expected to shrink their costs year-on-year. IT can pitch investment to do so, but it can be difficult for a big cost efficiency investment to rise up the prioritization list versus growth initiatives. They really want a solution with low operational requirements that allows delivery in a pay-per-use orientation of getting started by proving out value in cloud spend and utilization optimization by a set of application teams, which then justifies the cost of scaling to broader use.

HPE ONESPHERE FOR HYBRID CLOUD COST MANAGEMENT

HPE has built out a portfolio of capabilities supporting customers working towards cloud transformation of their private infrastructure with software-defined offerings. The company has also been working many IT services projects helping customers with cloud management needs across multiple public and private infrastructures, including via the public cloud services expertise added by HPE Pointnext's acquisition of Cloud Technology Partners (CTP). In seeing the patterns of enterprise struggles, HPE identified a clear need for hybrid cloud management beyond the application management focus of most cloud management platforms. HPE invested in and has delivered a new option for enterprise IT to deliver self-service infrastructure with strong developer empowerment with its HPE OneSphere platform for hybrid cloud cost management and application deployment.

Hybrid cloud cost management is a top capability focus for HPE OneSphere. Cost management starts with achieving spend visibility across cloud infrastructures. To deliver this, HPE OneSphere provides a structure for tagging resources upon provisioning that enables reporting capabilities. As an API-driven, services-based SaaS platform, HPE OneSphere allows an administrative user to add existing cloud infrastructures to its set of options for application developer use. Developers set up HPE OneSphere "Projects" to group infrastructure resources when provisioning them on any

available cloud, including across multiple clouds, and the resources are tagged as members when provisioned as part of the project.

This tagging for cost categorization is applied to resources provisioned through the HPE OneSphere user interface, through automation via its API to the cloud infrastructure (e.g., the HPE OneSphere API to EC2 API for VMs on AWS), and for all resources in an existing cloud account when it is mapped directly to a HPE OneSphere Project (e.g., an AWS sub-account). HPE has committed to support AWS, Google Cloud Platform, Microsoft Azure, Kubernetes, and VMware. HPE intends to add support for additional infrastructure providers, including OpenStack, bare metal and application-focused cloud management platforms layered on top of infrastructures like Mesosphere DC/OS.

This straightforward approach in HPE OneSphere, combined with its use of the billing and monitoring APIs from the cloud providers, enables its "Insights" dashboard for development, IT, and business leadership visibility into cloud spend and utilization across all tracked infrastructures. (For a private infrastructure, an administrator provides HPE OneSphere with its per-resource cost parameters.) More specifically, this allows drill-down by cloud, project (or team, application, etc.), and user. The enterprise gains a common language for its cost categorization with immediate visibility. Business leadership avoids the all too common cost surprises at the end of the month and gains the ability to use showback (i.e., directing team/user attention to spend) to drive increased empowerment to track and own accountability for spend throughout the organization. For example, month-to-date spend is shown on each Project's page.

This developer-empowering approach is very important for IT when introducing new cloud tooling. Its importance increases as the development organization adopts a more cloud-native structure, including a microservices architecture, empowering developers to manage iterative development and deployment of their application. They are very well set up to manage the costs of their application and be held accountable when provided the tools for visibility and control. Given how much development has led more than IT in enterprise cloud transformation, IT leaders know a cloud management tooling initiative is far more likely to succeed with widespread internal adoption if their approach is one of developer empowerment rather than IT attempting to take further control.

From an operational perspective, HPE OneSphere is provided as a SaaS offering to further aid IT in ease of setup and use by keeping its operational burden low. IT can deliver a scalable capability to support successful internal adoption growth without additional operational burden, while the service cost scales based on the average size of fleet being managed.

As a service provided by HPE, HPE OneSphere is developed in a cloud architecture using container-based services operated with Kubernetes container orchestration. This behind-the-scenes context from HPE is part of their setting an expectation of an iterative approach to feature development. Their vision for cost management includes features that build on cost visibility to better enable budgeting, control, and optimization by the project resource owners. It is easy to see that they could build towards surfacing recommendations for optimizing application deployment for cost, enabling capacity and/or spend limits configurable per project by teams to avoid surprises from abuse or mistakes by individuals, and towards the longer-term vision of automated application deployment to the most cost-efficient infrastructure meeting availability, feature, security, and compliance requirements.

CALL TO ACTION

Enterprise focus on hybrid cloud cost management is growing as cloud use matures from capturing business value benefits to optimizing their cost at increasing scale and complexity across infrastructures. There is common recognition of the importance of achieving and internally opening up visibility into spend and utilization split by relevant organizational unit of project, team, or application in order to empower ownership and accountability. Unfortunately, this need is often addressed through independent resource tagging efforts per cloud infrastructure that help with the tactical need but are likely to fall short in the long-term. That approach will constrain consolidated hybrid cloud visibility and therefore inhibit delivery of more strategic capability, for example, in recommended – eventually even dynamic automated – optimization of application deployment across cloud infrastructures.

Moor Insights & Strategy recommends that IT leaders – as the traditional drivers of infrastructure cost efficiency and the tooling to support it – surface this need based on current hybrid cloud cost management challenges and drive an approach to solving it from a long-term strategic infrastructure sourcing perspective, rather than a tactical approach likely to constrain future flexibility. Successfully delivering this type of cloud management capability could help IT take a more strategic role serving business needs in cloud transformation with broader management capabilities.

Moor Insights & Strategy recommends that IT leaders consider HPE OneSphere as a hybrid cloud management platform for addressing these needs by empowering application developers with easy-to-use, self-service tools for cost management. HPE OneSphere similarly delivers faster, simplified application deployment for developers and a self-service interface for private infrastructure to be considered as well in valuing

the offering. These capabilities are covered in the other two parts of this three-part series.

Few alternative cloud platform vendors have invested deeply in this area, particularly of the top vendors with existing enterprise credibility. They have focused more on investing in the feature depth of their own public clouds, private cloud platforms, or application management platforms. Application management platforms may further step into cost management, so IT should weigh the benefits of any included cost management capabilities of those platforms already adopted versus those from HPE OneSphere and its additional capabilities. If considering a standalone cost management service provider, IT should weigh the depth of cost management feature capabilities versus HPE OneSphere (e.g., consolidated cost management of private infrastructure along with public clouds), the benefits of the additional HPE OneSphere capabilities versus any from an alternative (e.g., helping cloud-enable private infrastructure), and any business risk assessed in adopting a newer / smaller vendor's platform.

More information on HPE OneSphere is available at: www.hpe.com/info/onesphere

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR
Rhett Dillingham, Vice President & Senior Analyst at Moor Insights & Strategy

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

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

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

LICENSING
This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

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

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

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

©2018 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.