IDC
Analyze the Future

I D C A N A L Y S T C O N N E C T I O N

Laura DuBois
Group Vice President

Strategies for Gaining Insight into Your Hybrid Cloud Data Infrastructure

November 2017

As cloud services evolve and adapt to meet the needs of large enterprises, successful IT organizations must also evolve and adapt so they can manage their data and applications across multiple clouds. For many, that begins with a hybrid cloud approach. Others are moving faster and have already encountered a need to manage enterprise data across one or more cloud services. Both hybrid and multicloud environments require visibility into data location, utilization, and performance so that organizations can make the best decisions about where to run applications and workloads.

The following questions were posed by NetApp to Laura DuBois, group vice president at IDC, on behalf of NetApp's customers.

Q. What are the differences between hybrid and multicloud environments?

A. The term "hybrid cloud" has morphed a bit over the past five years, but it is generally accepted as a blending of two different cloud environments. From an implementation standpoint, this often means a blending of on-premises systems and a public cloud service with a goal of leveraging the scale of resources or services in the public cloud ecosystem. Hybrid cloud may also stem from a CIO mandate to use public cloud services while security or compliance concerns may warrant keeping data on-premises, resulting in a hybrid cloud scenario. Examples include doing test/dev in public cloud and running production in private cloud or running analytics jobs in public cloud with results brought back to a private cloud. Another example is running a SaaS application for Office 365 in a cloud service but leveraging another cloud service for disaster recovery.

Multicloud is already a reality today as more than 70% of firms have a cloud-first strategy for SaaS workloads. However, multicloud is more an artifact of many different cloud environments that are in place today but that may not have any interoperation with each other. It often creates many different data silos across different SaaS, IaaS, and PaaS providers. Multicloud is appealing to customers for speed-of-service access and the ability to gain all the advantages of public cloud. For example, you may get best-in-class services for HR, payroll, and other apps with leading SaaS providers; you may get best-of-breed container services with one IaaS provider and a breadth of analytics services from another IaaS provider. Multicloud gives you all the choices.

But most customers quickly become concerned with placing all their data with one provider without having a second source somewhere — either on-premises or maybe with another provider. Hence, multicloud often moves to hybrid cloud. Further, given the rate of innovation and change in public cloud services, the ability to leverage cost, performance, and service-level agreement (SLA) advantages across providers is also a critical motivation.

This ultimately gets to the end state of hybrid cloud, which is seamless and transparent workload, data, microservice, and customer portability between and across cloud environments — both private and public. However, to effectively achieve this end state requires the ecosystem to have a common set of APIs and infrastructure services with the right level of workload and data abstraction to enable true portability.

And while a workload with a lighter footprint might be portable, data itself is far less so. Data is often said to have "mass" and "gravity," which means that it can be difficult to move data to a new location once it reaches a certain size — even over the fastest networks. Because of these barriers to mobility, data is increasingly becoming geodistributed to satisfy an increasingly global customer and stakeholder environment.

Q. What infrastructure management challenges do you see surfacing as IT organizations embrace and deploy hybrid cloud environments?

A. Challenges stem from first identifying which workloads can and should be moved to the cloud and then mapping the performance characteristics and SLAs of those on-premises workloads to the right service class and instance type. This has resulted in managed cloud service offerings to help with the transition of a workload from on-premises to public cloud as well as investments in tools to help profile workload baselines. We also see customers using a number of multicloud management solutions to help identify and triage runaway cloud costs.

Part of understanding cloud costs includes understanding the service and instance types, such as spot versus reserved instances, as well as data volume and egress fees. As data continues to grow, storage consumption also increases. Thus, identifying orphan assets and managing capacity and usage against planned projections becomes critical. All these challenges, from SLA and performance management to cost optimization and planning, span both on-premises and public cloud infrastructures. For firms to get a holistic picture of infrastructure costs, management solutions need to work across the entire hybrid cloud environment.

Q. How do data governance and control factor into an organization's cloud decision-making process?

A. The "right" data placement is paramount to complying with regulatory mandates, ensuring application response time based on customer location, and optimizing infrastructure costs.

Regulations such as GDPR as well as industry-specific mandates are increasingly stipulating where data can reside. I recently met with a company that does oil exploration, characterization, drilling, production, and processing in over 80 countries. Although the company is a large multinational, its data is collected across a vast array of oil rigs and source sites and must remain within each local country. Organizations face the increasingly common challenge of maintaining data in the country in which it originated. This has resulted in the need for prescribed geographic placement and verification of primary and copy data.

The second factor driving data locality is around application and data response time. For example, mobile applications used by customers across geographies need the same level of data access in New York as in China. This requires the geographic dispersal of the application and the data. With these applications, data is dispersed across geographies as well to ensure faster data access.

The last factor driving data locality is the optimization of infrastructure costs. This means placing colder or dark data on higher-density, less expensive tiers, while an active working set is placed locally in memory and closest to the processor.

As mentioned previously, the use of public cloud IaaS starts out with predictable costs, but this soon changes. Over time, it's common for bills to go well beyond what has been planned and budgeted for by organizations. This unexpected increase in costs often relates to egress fees. While there are strategies to mitigate these unexpected costs, it takes discipline and centralized oversight to implement and realize the benefits of the various approaches. Further, the use of IaaS commonly spans accounts and teams, which has made achieving centralized visibility a challenge. This has given rise to many different cloud management tools for cloud service costing and optimization as well as improvements by IaaS providers themselves. Of course, cloud costs are only part of the equation. You also need insight into on-premises costs.

IT organizations are investing in management tools and hybrid cloud offerings to incorporate cloud costs and provide showback or billback to the business units. But an accurate cost picture must also include administrative, facilities, network, and power and cooling costs to get a complete cloud cost accounting picture.

Q. What recommendations do you have for IT leaders who need to maintain control of their data as it becomes spread across hybrid and multicloud environments?

A. The first step needs to be an accurate and holistic picture of all your workloads and data. I refer to these as data maps that include discovery and visualization but should also include understanding of the data, key data owners and custodians, level of sensitivity of the data, data integration sources and targets, API controls, and so forth. Having this holistic picture serves as the foundation for beginning to implement a common control strategy. Then you want to consider what your controls are, and these should span security, privacy, retention, disposition, access, permissions, and so on. It is really the responsibility of IT operations to implement this level of data control. However, it is up to constituents such as lines of business, data owners, legal, security, and HR to assign business rules for how data is secured, retained, disposed of, or otherwise controlled.

With a clear picture of the data maps, IT can then begin to think about managing the workload associated with the data. This workload is likely abstracted as a virtual machine or maybe even a container with a persistent volume. This abstraction gives the workload greater portability, and this is where consideration of running the workload in a public or private cloud begins. The existing workload already has set SLAs defined with the business. If not, users have come to expect a certain level of availability and performance even if an SLA has not been defined. So you want to benchmark throughput, IOPS, latency, and availability metrics in the current environment and ensure that moving the workload to the new cloud environment leads to results that are as good as or better than those associated with the previous environment.

Assuming the workload is easily moved from a traditional three-tier infrastructure architecture to a public cloud IaaS service, operations teams must monitor cost and performance over time and ensure the workload is placed on the correct computing and I/O tier or service

instance. There are tools you should evaluate that provide a dashboard spanning on-premises and clouds. These tools not only collect and report on environments but also use predictive analytics to develop a data-driven approach to managing your infrastructure.

Q. Can new technologies, such as analytics and machine learning, be used to help IT teams manage and control data sprawl?

A. Artificial intelligence (AI) and machine learning, a subset of the larger discipline of AI, are really going to transform enterprise infrastructure and business processes. Machine learning involves the design and creation of systems that can learn based on the data they collect. AI algorithms will be able to review and understand log files throughout the IT infrastructure in ways that are impossible with traditional collection and analysis methods today. AI will be able to predict a system crash or component failure minutes or hours before a human might notice anything was wrong. In this regard, AI holds great promise with its ability to learn patterns in networks, devices, and systems and decode deviations that could reveal problems before an event occurs or detect in-progress cyberattacks.

AI will become integrated into public and private cloud architectures where machine learning is used to transition data seamlessly between infrastructure tiers and cloud environments without human intervention. Automated data movement across heterogeneous infrastructure, effective utilization of current assets in conjunction with public cloud adoption, and automation of provisioning and repetitive tasks will lead to a reduction in infrastructure capex and opex. Data will be analyzed, and logic engines will automate the data movement. Organizations will also be able to more effectively utilize infrastructure resources, from processors to memory and persistence. Further, as IoT progresses, the amount of unstructured machine data created will far exceed our ability to make sense of it with current analytical methods. Organizations will use AI to mine billions of data points for actionable insights, which will be valuable for incremental revenue streams and competitive differentiation.

Ultimately, AI will drive self-configurable and self-healing infrastructure, improving productivity and eliminating processes prone to human errors. An example is intelligent network automation using network virtualization platforms such as Cisco's ACI and VMware's NSX. Often customers do not know what policies they should implement because they do not have enough insight into what is happening on their network. By leveraging machine learning, customers will benefit from the ability of the system to make recommendations on optimal policies to implement. Analytics is the first step in the journey, while automation will be the next step to enhance overall productivity and accuracy.

ABOUT THIS ANALYST

Laura DuBois serves as group vice president for IDC's cloud IaaS, computing platforms, storage, and system infrastructure software research, quarterly trackers, end-user research as well as advisory services and consulting projects. She leads research spanning topics such as IoT, cloud computing, SaaS infrastructure services, software-defined infrastructure, and the impact of machine learning and artificial intelligence on enterprise infrastructure.

ABOUT THIS PUBLICATION

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com