Understanding Complex Cloud Patterns

CLOUD COMPUTING CONTINUES TO EMERGE AS THE ENTERPRISE'S ENABLING TECHNOLOGY OF CHOICE. The trend has been to leverage clouds as complex, highly heterogeneous, and distributed architectures, including hybrid and multiclouds. Moreover, as we leverage heterogeneous public clouds, there's a need to provide high-speed data transfer and data integration between public and private clouds. Thus we have the emerging use of the intercloud, or special-purpose networking between public and private clouds.

The use of these approaches and architectures has launched a new set of technologies within the world of the cloud to deal with governance, security, and management of these solutions. Thus, we have to layer technologies to deal with these issues. However, the largest challenge is just understanding these complex cloud architectures, how they work, and the value they're likely to bring. This article provides the information you need to make solid choices around the use of complex, heterogonous, and widely distributed architectures.

DAVID S. LINTHICUM

Cloud Technology Partners

david@davidlinthicum.com

Hybrid Clouds

Back in 2008, the industry first began the hybrid cloud computing model discussion as defined by the National Institute of Standards and Technology (NIST).[1] Cloud computing purists pushed back hard.[2] After all, they already thought private clouds were just another name for the datacenter. To them, the idea of hybrid clouds that used private clouds or traditional computing platforms was just as ridiculous.

However, fast forward to 2016, and the use of the hybrid cloud model is pretty widespread. This complex cloud architecture needs to be understood in terms of value, proper implementation, and its ability to leverage different architectural patterns to best balance the loads between private and public clouds.

NIST defines hybrid clouds as follows:

> The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).[1]

Over time, it became clear that hybrid cloud computing approaches have valid roles within enterprises, as IT tries to mix and match public clouds and local IT assets to get the best bang for the buck. Now it's the cloud computing providers who are pushing back on hybrid cloud computing, as they instead try to promote a pure public cloud computing model. However, these providers are inadvertently hurting the adoption of cloud computing. Although public cloud computing has valid applications, the path to it isn't all that clear to rank-and-file enterprises.

Hybrid clouds provide a clear use case for public cloud computing. Specific aspects of existing IT infrastructure (say, storage and compute) occur in public cloud environments, and the remainder of the IT infrastructure stays on premises. Take the case of business intelligence in the cloud: Although

 2325-6095/16/$33.00 © 2016 IEEE

some people promote the migration of gigabytes of operational data to the cloud, many others find the hybrid approach of keeping the data local and the analytical processing in the cloud to be much more practical.

Using a hybrid model is a valuable approach to architecture since you can mix and match resources between local infrastructures, which is typically a sunk cost but difficult to scale, with infrastructure that's scalable and provisioned on demand. You place the applications and data on the best platforms, then span the processing between them.

The use of hybrid computing acknowledges and validates the fact that not all IT resources should exist in public clouds today—and some might never exist in public clouds. Given the compliance issues, performance requirements, and security restrictions, the need for local computing is a fact of life. The hybrid model helps us all better understand what compute cycles and data have to be kept local, and what can be processed remotely.

Of course, some cloud providers already have their eye on leveraging a hybrid model. These new kids on the block even provide management and operating systems layers specifically built for hybrid clouds. However, most public cloud providers are religious about pushing everything outside of the firewall (after all, that's where they make their money).[3]

Also keep in mind the concept of federated clouds, which are related to hybrid clouds. A federated cloud (also called cloud federation) is a configuration of multiple external and internal cloud computing services that align with the enterprise's business needs. A federation is the union of several smaller parts that perform a common action, and can be fine- or course-grained services, infrastructure resources (such as storage), or even applications.

For more information on specific hybrid cloud architectures, see my "Cloud Tidbits" column in this issue.[4]

Intercloud and Intracloud

Cloud-to-cloud data transfer can be either intracloud (such as within Amazon Web Services [AWS]) or intercloud (such as between AWS and Google). Although these intra- and intercloud services provide wider bandwidth than the open Internet, they typically cost more to leverage, and, in most cases, are add-on options where you're charged by the amount of traffic.

Intracloud data transfer performance is directly related to the size of the pipe within the cloud service, as well as to performance-enhancing services, such as cache services that might be in use. Intracloud moves data from place to place in the same cloud, typically from a set of machine instances to another set of machine instances, usually with storage.

Typically, intracloud data transfer is between tenants, virtual machines, or applications and datastores. The approaches and technology vary greatly, so you should run your own benchmarks to duplicate the scenario you plan to implement. Moreover, do a total cost of ownership (TCO) analysis to be sure that you won't be surprised by the data transfer bill your public cloud provider sends you at the end of the month.

Intercloud data transfer is even more complex, having to exchange data with cloud services that might not like each other. Moreover, the open Internet is the typical mode of transport, so the same issues arise here as with cloud to enterprise. I suspect that cloud providers will get better at this architecture in the future, for the sake of their mutual clients.[5] However, as we progress in time, more intercloud connects are becoming dedicated circuits running between providers just for the purpose of high-speed data transfer. This trend will continue as public cloud providers understand that it's in their best interest to work with other public cloud providers for the best outcome of their cloud-to-cloud solutions.

Multicloud

The use of multiple clouds has evolved from architectural patterns required to solve business problems. Many business departments want to use cloud computing within public clouds, outside of the company firewall. Today we have many types of public and private clouds that provide security, governance, and management tools to support other clouds, and which can be combined to create a composite solution.

Enterprises move to multicloud for a variety of reasons:

- Single cloud solutions typically don't provide the breadth and depth of functionality that enterprises require for all of their cloud computing solutions.
- The rise of cloud management platforms (CMPs) gives enterprises a single interface to help provision, manage, and scale complex environments.
- Usage-based pricing makes it easier for enterprise IT to evaluate the cost of cloud computing, including show-back and charge-back services.
- Companies that want to move applications into public clouds need a range of services, including different database, middleware, development, and compute services, and this drives the use of multiple cloud computing platforms.
- The growing use of platforms, infrastructure, and software in the cloud results in multiple forms of clouds. Consequently, IT must often support two or more public clouds for development and operations teams that use the cloud to create new business applications and services.

The *RightScale 2015 State of the Cloud Report* proved that the movement to multicloud is real, with 82 percent of those surveyed identifying multicloud as their current strategic direction.[6]

Enterprises moving to multicloud face some critical choices, including the types and brands of cloud to leverage and the approaches and technology to use in managing a multicloud solution. These approaches seem to have some common patterns of failure, as well as patterns of success.

Multicloud architectures have their own sets of pros and cons. The core question that many enterprises ask is: How much cloud heterogeneity is a good thing? Also, when does heterogeneity bring too much complexity and risk? Indeed, you can think of multicloud as a complex hybrid cloud that has more than two brands of clouds within the architecture.

Enterprises that tried to maintain homogenous on-premises IT infrastructures lost the battle a long time ago. Typical enterprise architectures have been built through years of solving tactical problems with whatever technology seemed to be right at the time. Over the years, these on-premises technology solutions became very heterogeneous, and thus very complex.

Multicloud shouldn't evolve the same way. Those charged with picking the right cloud technology should consider the solution patterns to fit the problem patterns, and use that as a guide to select and deploy the right public and private cloud technologies. Although some might find that a single cloud provider is the best solution,[6] most are driving cloud solutions using best-of-breed technology, and thus end up with a number of cloud types and brands.

Although this isn't the first time enterprises moving to new technology have faced the "homogeneous versus heterogeneous" question, the move to cloud computing brings some new challenges and confusion. Clouds are, indeed, platforms, but they also provide common resources that other cloud brands may share.

Clouds take more of a service-oriented approach to architecture, and the enterprise typically ends up with a common services catalog that might link back to many types and brands of clouds. Clouds represent a collection of services that can be mixed and matched to form applications, more so than monolithic applications themselves (see Figure 1).

For example, an enterprise could take storage services from one public cloud provider and mash them up with compute services from another provider, and perhaps introduce database services that are running on premises. The ability to build solutions out of best-of-breed cloud services provides a solid foundation to justify a multicloud approach and is the primary driver of multicloud use. The more clouds you leverage, the more services you have in your catalog, and thus application development becomes more of an assembly process. This allows enterprises to quickly build or change applications, providing the value of agility and speed-to-market for the business. Agility is the fundamental way that cloud computing provides value. The more clouds (such as with multicloud), the more the business is able to solve problems or adjust to changes in the market, which means it can make more money.

Finally, consider off- and on-premises cloud models, such as virtual private clouds (VPCs). Each model is sometimes preferred over private and public because you can have a dedicated zone in a public cloud provider for one customer. These models provide cloud computing's cost effectiveness, but also provide the data privacy that business might prefer.

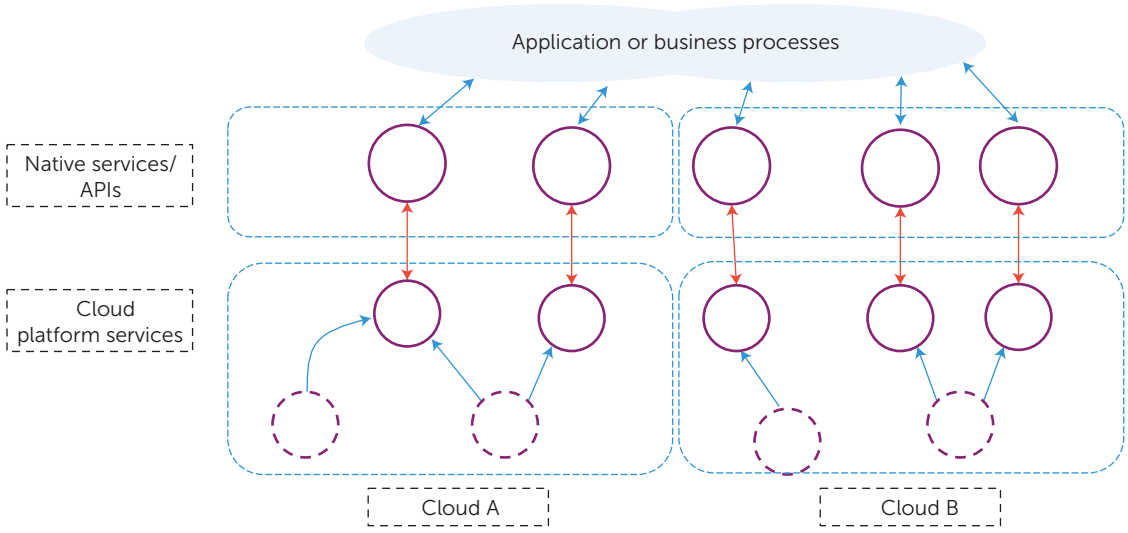

FIGURE 1. When leveraging cloud-based platforms, you can mix and match services to form business processes or applications. This drives the value of multicloud.

THE USE OF COMPLEX CLOUD ARCHITECTURES IS STILL AN EVOLVING SCIENCE. It's helpful to understand the core patterns, and perhaps a match to your ultimate IT solution as you evolve your own cloud computing strategy.

As a path forward, you should consider your "as is" state, including data, processes, and compute, and how those components map to cloud, or should map to cloud computing solutions. Although the task is laborious, it's well worth the time invested due to the efficiencies that can be gained from the use of these complex cloud architectures. Of course, the trade-offs are the cost of dealing with the complexities versus the value that can be had. Some total cost of ownership analysis should be done before moving down this path.

References

1. P. Mell and T. Grance, *The NIST Definition of Cloud Computing, Nat'l Inst. of Standards and Technology*, NIST Special Publication 800-145, 2011; http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf.
2. E. Brown, "Final Version of NIST Cloud Computing Definition Published," *NIST Tech Beat*, 25 Oct. 2011; http://www.nist.gov/itl/csd/cloud-102511.cfm.
3. D. Linthicum, "Why the Hybrid Cloud Model Is the Best Approach," *InfoWorld*, 27 Jan. 2011; www.infoworld.com/article/2625289/cloud-computing/why-the-hybrid-cloud-model-is-the-best-approach.html.
4. D.S. Linthicum, "Emerging Hybrid Cloud Patterns," *IEEE Cloud Computing*, vol. 3, no. 1, pp. 88–91.
5. D. Linthicum, "All You Need to Know about Cloud Data Transfers," *InfoWorld*, 23 Mar. 2012; www.infoworld.com/article/2619818/cloud-computing/all-you-need-to-know-about-cloud-data-transfers.html
6. *RightScale 2015 State of the Cloud Report*, RightScale, 2015; www.rightscale.com/lp/2015-state-of-the-cloud-report.

DAVID S. LINTHICUM *is senior vice president of Cloud Technology Partners. He's also Gigaom's research analyst, and he frequently writes for* InfoWorld *on deep technology subjects. His research interests include complex distributed systems, including cloud computing, data integration, service-oriented architecture, and big data systems. Contact him at david@davidlinthicum.com.*