

Consolidated eBook of three separate articles published by Dave Young.


Three Part Series: Multicloud Strategy

Part I:

Should You Consider a Multi-Cloud Strategy?

Many enterprises adopt a multicloud strategy out of a desire to avoid Cloud supplier lock-in or to take advantage of innovative solutions with a certain Cloud provider. There are many rationales for using multiple clouds, as they each have different strengths and weaknesses for performance and functionality, different pricing schemes, and some workloads simply run better on one cloud platform versus another.

Although neither provide a financial breakdown of their cloud business, AWS continues to dominate the cloud infrastructure market, but its lead is starting to wane as Microsoft Azure and Google Compute Engine make significant inroads. This trend will undoubtedly continue as gains from Microsoft Azure points to a desire of many enterprises to use public Clouds other than Amazon. Further, this multicloud strategy also leads to more public Cloud plus private Cloud, as well as Public and Hybrid Cloud (definition: a mix of on-premises, private Cloud and third-party, public Cloud services with orchestration between the two platforms).



When adopting a multicloud strategy, measuring “cost to performance” becomes important to ensure cost is optimally managed.

Multicloud deployments also reduce the risk of service disruptions with a dependency on a single Cloud supplier, improves negotiating leverage with competition and allows you to choose the best Cloud solution or tool for the job.

Multicloud Strategies:

- A practical application of a multicloud strategy can take a variety of forms, starting simply by placing development and testing in Cloud A and Production Cloud B, and this is probably the most common version of multicloud. Or simply keep Production on-premise and development in the Cloud, which is the most common early adopter approach to Cloud services.
- Another multicloud approach is bursting transactions to the Cloud. In this case, where an online transaction processing application or retail operation during peak holiday season or video streaming service typically runs in an on-premises environment, but during certain peak times, the organization will spin up additional capacity in the cloud. The objective of cloud-based bursting in this manner is to avoid the cost of buying enough hardware for peak capacity when that hardware will sit underutilized for most of the time.
- A third, and arguably more complex, approach to multicloud is Cloud-based disaster recovery, involving running a production application in one cloud or on-premises environment but maintaining a separate, up-to-date copy of the application in another Cloud environment. This is a wise strategy, as any Disaster Recovery Planner knows, that geographic and supplier diversity is vital to ensure a dependency or proximity to a single host for your applications does not put your entire IT environment, and your business, at risk.

Elevating the maturity of a multicloud strategy leads to multicloud operations of a single application, where the same instance of an application runs across multiple Cloud providers in an active to active manner. This provides for a more scalable architecture, above and beyond the Cloud bursting examples, but also provides diversity to potentially survive a DDoS attack or a failure in one Cloud supplier's availability zone.

Part II:

Pricing and Product Transparency

Is a Cloud Broker right for you? A Cloud Broker by definition is an entity that manages the use, performance and delivery of cloud services, and negotiates relationships between cloud providers and cloud consumers. As cloud computing evolves, the integration of cloud services may be too complex for Cloud consumers to manage alone. A Cloud Broker can also be in a unique position to collect, analyze and conduct opportunity assessments on pricing that will be key to a multicloud or Cloud Arbitrage strategy.

Generally, cloud arbitrage is the practice of taking advantage of a price difference between two or more Cloud providers for the same product: striking a combination of matching Cloud deals that capitalize upon the imbalance, the cost savings being the difference between the market prices at which the Cloud product unit is traded.

The challenge is that this imbalance can also disappear quickly as Cloud providers are continually adjusting pricing and changing product generations and feature sets over thousands of product types, so arbitrage opportunities are likely to be transient and short term. It is a sophisticated undertaking to keep track of and typically requires the service of a *Cloud broker* and related *Cloud services partner* to assist as a proxy to help the enterprise navigate opportunities for Cloud service arbitrage.

A Cloud Broker is not a requirement, however. An AWS or Azure partner can be used to help you navigate the myriad product and pricing considerations. There are literally thousands of product line items with AWS, add to this Azure and GCE, as well as others, and you have a daunting task to run the analysis to determine opportunities for cost savings or new technologies or better performance between providers at any given moment.

For instance, AWS has an option to subscribe to Amazon Simple Notification Service (Amazon SNS) notifications to get alerts when prices for the services change. AWS prices change periodically, such as when AWS reduces prices, when new instance types are released, or when new services are introduced. The AWS Price List Service API (AKA the Query API) and AWS Price List API (AKA the Bulk API) enable you to query for the prices of AWS services using either JSON (with the Price List Service API) or HTML (with the AWS Price List API), so you have an option to programmatically extract and analyze products and pricing dynamically as needed.

In addition, AWS and Azure have different prices on a region-by-region basis. You may pay 50% or more based simply on what part of the world you're in. To make things more interesting, Azure will match AWS pricing for comparable services.

Further compounding the pricing complexity is the fact that deploying workloads to the Cloud requires more than a single server, plus storage (often bundled, sometimes not depending on the type) and network and other necessary services to meet the total application needs. While the enterprise can discount for reserved instances with a multi-year commitment, as the case with AWS and Azure, the trend remains that Cloud prices continue to drop, and you must exercise caution when making financial commitments.

Evaluating the various Cloud services offered by Cloud providers is also not as straightforward as one might assume. Consider the following when comparing providers:

- Actual performance, as measured by an objective benchmarking process.
- Class or size of a Cloud service
- CPU technology and chip generation
- Service specifications
- Storage technology and configurations
- Underlying technology and data center and network architecture
- Value-added services, e.g. support, managed services, add-on value-added services.

Understanding Cloud service pricing is a science in itself and this is only going to grow in complexity. Cloud providers continuously roll out new versions of existing services, add completely new services, and constantly employ new pricing strategies targeted toward achieving a competitive advantage against other providers.

Part III:

Application Portability

The enterprise's application processing workloads must be compatible to portability between Cloud providers, including technology stacks that are common across all Cloud provider platforms, such as the use of application microservices and open stack services such as containers and Kubernetes. It is also important to consider how you'll integrate data and applications residing in different Clouds.

Containers are ready-to-run applications comprised of the entire stack of services required to run. OpenStack is an ideal platform for containers because it provides all the resources and services for containers to run in a distributed, massively scalable cloud infrastructure. OpenStack is a set of software tools for building and managing cloud computing platforms for public and private clouds. OpenStack is managed by the OpenStack Foundation, a non-profit that oversees both development and community-building around the project.

The concept of containers itself is beneficial to hybrid and multicloud users. The trend in Cloud computing, public or private, is moving away from the traditional virtual machine method. VMs are slower and use more computing power, don't scale as well and are difficult to move. Containers by contrast improve on VMs by running atop the OS kernel. This characteristic makes containers lightweight, portable, and faster on startup. This kind of flexibility is important for DevOps teams utilizing multiple clouds and key to any effort at Cloud Arbitrage.

It is not enough to have a container-based platform, however. Enterprises seeking Cloud portability must have a solution for moving data between container environments and managing it within those environments. Without a data management solution, enterprises will be challenged with being able to execute on multicloud workflows because data acts as a tether on mobility between Cloud environments.

Cost Optimization

A multicloud strategy should not be a distraction from a thorough and constant effort at Cloud Cost Optimization with all Cloud providers. This is the low hanging fruit of cost savings when taking some basic steps, and real material savings with a meticulous, expert approach to uncover the myriad ways to avoid overpaying for Cloud services.

In surveyed organizations, the number one priority in 2019 is cloud cost optimization, as the challenge of managing cloud spend grows as cloud use increases.

In short, at its most basic level, Cloud cost optimization is the process of reducing your overall cloud spend by identifying mismanaged resources, eliminating waste, reserving capacity for higher discounts, and right sizing computing services to scale.

Conclusion

Realizing the promises to efficiencies and cost savings using public Cloud and, in particular, multicloud requires significant organizational maturity and the use of partnerships, like NET(net) and others than can bring the outside expertise necessary to elevate any company's cost savings capacity.

About NET(net)

Founded in 2002, NET(net) is the world's leading IT Investment Optimization firm, helping clients find, get and keep more economic and strategic value. With over 2,500 clients around the world in nearly all industries and geographies, and with the experience of over 25,000 field engagements with over 250 technology suppliers in XaaS, Cloud, Hardware, Software, Services, Healthcare, Outsourcing, Infrastructure, Telecommunications, and other areas of IT spend, resulting in incremental client captured value in excess of \$250 billion since 2002. NET(net) has the expertise you need, the experience you want, and the performance you demand. Contact us today at info@netnetweb.com, visit us online at www.netnetweb.com, or call us at +1-866-2-NET-net to see if we can help you capture more value in your IT investments, agreements, and relationships.

NET(net)'s Website/Blogs/Articles/eBooks and other content is subject to NET(net)'s legal terms offered for general information purposes only, and while NET(net) may offer views and opinions regarding the subject matter, such views and opinions are not intended to malign or disparage any other company or other individual or group.