

Cloud Expert Q&A: Optimizing the Cloud Mix, Powering High-Performance Applications



We sat down with ViON's EVP of Operations Rob Davies, and Dan Fallon, Senior Director of Public Sector Systems Engineering at Nutanix, to discuss how to best manage high-performance applications that are part of a multi-cloud strategy, and the key considerations every agency leader should be evaluating in that equation.

MeriTalk: What is the biggest issue affecting IT leaders using cloud in enterprise environments?

Rob Davies: As more organizations adopt resource-intensive applications like Artificial Intelligence (AI) and Machine Learning (ML), simulation software, financial modeling, business intelligence, and advanced analytics and visualization, the need for high availability, reliability, and high-performance infrastructure to support these applications becomes abundantly clear. In the midst of this drive for performance, many organizations are turning to the public cloud as part of a broader IT strategy. But they are finding that while some applications are well suited to the public cloud, it generally lacks the speed, power and capacity to support high performance applications.

Customers may think public cloud is the only option. It's seen as a mandate or silver bullet. But people start moving to the public cloud, and have buyer's remorse when they quickly spend their budget in less than a fiscal year. And it becomes evident that there are management and security issues even when using FedRAMP-compliant providers.

Today, there is a constant battle between public and private cloud, and what's best to pursue. Organizations are finding that while public cloud may initially appear to be the least expensive option, the costs really add up over time for some workloads. IT leaders

“A recent IDC survey reported that 86 percent of enterprises are considering “repatriation” – moving applications from public clouds back to the data center for one or more workloads.”

-Rob Davies, EVP of Operations, ViON

often think they can't have a technology as-a-service on-premise – so they move to public cloud and then begin to experience a lack of flexibility and rising costs because their workloads don't align with what public cloud can offer. This issue has put many organizations in a difficult financial position that isn't sustainable long-term. In fact, a recent IDC survey reported that 86 percent of enterprises are considering “repatriation” – moving applications from public clouds back to the data center for one or more workloads.

MeriTalk: What limitations do you see with a public cloud environment for running high- performance applications?

Dan Fallon: High-performance apps are typically the crux of an operation and, depending on where these mission-critical applications need to plug into various systems, organizations may have latency issues. The interdependence between applications requires a multi-cloud strategy to truly optimize operations and minimize costs.

“High-performance apps are typically the crux of an operation... The interdependence between applications requires a multi-cloud strategy to truly optimize operations and minimize costs.”

-Dan Fallon, Senior Director of Public Sector Systems Engineering, Nutanix

This is where Hyperconverged Infrastructure (HCI) can greatly reduce costs. There are private and hybrid cloud models available today, similar to public cloud, that give organizations the opportunity to leverage HCI for greater agility than a public cloud environment. For example, ERP or payroll system applications are typically high performance with a database backend – these are critical operations that are interdependent. Government often has homegrown applications that run systems for citizen services, tax records, etc. These are the applications that are critical to the business of government work, and that run across platforms. The plurality of the applications is forcing organizations to move from Cloud First to Cloud Smart – considering different cloud models for different applications.

RD: In general, higher performance applications equate to higher CPU/memory costs. You have to evaluate that in the cloud the same way; the whole cost model will **absolutely** change when the conversation shifts to high performance applications. Some organizations make the mistake of doing the same cost modeling without awareness that these types of applications dramatically change the model.

Not all clouds are created equally and aligning workloads to the right cloud environment is fundamental. Without that alignment, you run the risk of decreased performance, security exposure, compliance and governance uncertainty, latency issues and increased costs for dedicated servers and bandwidth consumption.

MeriTalk: How can a multi-cloud environment impact a cloud strategy?

RD: A hybrid multi-cloud environment allows IT organizations to leverage the best attributes of the public cloud, private cloud, and on-prem infrastructure. Hyperconverged infrastructure delivers on-premise IT services with the speed and operational efficiency of the public cloud. Users have the flexibility to leverage different environments based upon performance requirements. There are three primary factors that are driving the multi-cloud strategy: 1) Application performance; 2) Data security; and 3) Cost.

Ensuring optimal **Application Performance** requires identifying the right environment for the right workload. The right architecture can simplify management for business-critical applications.

IT leaders should look to HCI to maximize uptime with native **Data Protection and Security**. With HCI, these capabilities are inherent in the system, providing greater control and visibility while streamlining management of the infrastructure.

With a hybrid multi-cloud environment, leaders can move beyond the cost and complexity associated with legacy infrastructure and the need for specialists to maintain it. Additionally, for high performance applications, it's important to evaluate the Total Cost of Ownership. Public cloud may seem less costly up front, but when considering long-term costs and desired performance, it may not be the best option.

MeriTalk: How do you determine what applications should be moved to the cloud and what should remain on-prem?

DF: The first things that go to the cloud are more public-facing applications. The other consideration is dynamic versus static workloads. You can think of public cloud versus on-prem like a hotel rental versus owning a home. During certain times of the year where you need more resources, you can rent public cloud for those times. When you have static steady-state applications – you are building them on-prem and keeping them on-prem. It's important to bridge the gap with steady-state on-prem and scale up and burst to cloud when it is required. On the other hand, DevOps is more dynamic and might need resources in the public cloud and then bring the apps back on-prem when the building phase is done and you are ready to move to production.

For example, AI is so resource-intensive that it should stay on-prem. Similarly, Internet of Things (IOT) applications like self-driving cars need to be done with real-time analytics with real-time processing at the edge. However, these same applications will have data sets that are sent back to a central hub to understand and examine patterns and use different environments to advance different components of the overall strategy. This is where you can really stretch the idea of multi-cloud.

MeriTalk: How does Hyperconverged Infrastructure (HCI) fit into a multi-cloud operation and what value does it bring to the equation?

DF: HCI has been a solid approach to infrastructure for a while and can easily be the backbone of a multi-cloud operation. However, it's important to consider the software with which you run your HCI. This is often overlooked, but where the data is stored has a big impact on minimizing latency, which is a base requirement for running high-performance applications.

The space has definitely matured. HCI originated when Google indexed the internet at enterprise scale (with a distributed file system). More recently, Nutanix took that same idea and applied it for enterprise use. They wanted to solve the latency issue of HCI. In some HCI operations, data can end up spread out everywhere. Now, Nutanix does HCI with data locality – keeping the data as close to the compute as possible, making it the fastest way to run applications. The Nutanix solution moves hot data into the RAM, using various tiers of storage – memory-class storage like flash, SSD – making sure the hottest data lives in the highest tier of storage. One platform for both transactional and analytical workloads creates faster application performance.

MeriTalk: How can organizations avoid common procurement challenges associated with changing cloud operations?

DF: Oftentimes, organizations think public cloud is the only way because they haven't seen the possibilities that other options provide. When considering a move, the procurement journey seems arduous and full of unknowns. The as-a-service financial model enables organizations to get HCI with the right software to create a private cloud in their own data center with no upfront investment and greater agility to scale up and down.

RD: Today we have many Federal agencies using these exact models, with examples across all different types of workloads throughout government. Agencies have generally tested this model and shift to on-prem in phases. In every case, as they prove out a concept by trying a new workload/workload application, they continue to expand out to additional phases as they see the performance and cost improve. It's really changed the way agencies look at procurement and cloud workloads.