Podcast Episode

Public Sector Embracing Digital Transformation



Benefits of Multi-Cloud Architecture

by Darren W Pulsipher - June 22, 2020

A multi-hybrid cloud architecture allows organizations to take advantage of the benefits of both private and public clouds, optimizing resources and cost efficiency. This model has five main advantages: agility, flexibility, predictive performance, security and compliance, and efficiency.



Video: Youtube Channel Podcast: Podcast Episode Blog: Blog Entry Keywords: multicloud

Agility

Agility is the ability to move quickly to adapt to changing conditions in new technologies, new applications, and new threats such as security and competition. A multi-hybrid model allows the ability to move things among private and public clouds, or even from legacy infrastructure to clouds. It may be tempting for development teams to initially use the public cloud under tight deadlines because they can quickly spin up and spin down infrastructure, but there can be integration problems in the late stages of deploying applications into product. This can result in costly delays. By using the multi-hybrid model, many of these integration points are exposed to the development team early. Application development uses this to fail fast and early in the development process.

Flexibility

Flexibility in infrastructure is a close relative of agility. Where agility is the ability to move quickly, flexibility is the ability to change. For maximum flexibility, developers need the ability to deploy an application anywhere across private or public clouds or legacy infrastructures, and not be locked into any one cloud's service or infrastructure.

By deploying a cloud management platform (CPM) in a multi-hybrid architecture, workloads can be easily redirected to different clouds according to cost, security, and reliability.An example of the consequences of lack of flexibility is Netflix's early decision to use only one cloud service provider. This public cloud had an infrastructure issue that resulted in downtime, and Netflix was unable to stream for several hours on the east coast of the United States. After that disaster, Netflix built in a multi-cloud solution so that they can quickly migrate to another cloud if there are issues. In addition, they are now able to move where it makes the most sense at any time according to cost, security, and reliability.

The applications of a multi-hybrid cloud are portability and operational flexibility. By not being tied to one cloud's way of doing things, developers will write code that can easily be transported between clouds. In addition, you will have the operational flexibility to move workloads between clouds to offer your customers what they want in a secure, cost-conscious, and reliable manner.

Predictive Performance

Public clouds often come with a price: "noisy neighbors," which can impact predictive performance or Quality of Service (QoS).

When using a public cloud, many times you do not know what or who else is running on the same machine, storage array, or network as you. For some workloads, this is not a problem. If you have an application, however, where you need more predictive performance, noisy neighbors, or those who monopolize bandwidth, disk I/O, CPU, and other resources, can interfere with the QoS results you need.

Private clouds can also suffer from noisy neighbors, but since you own the infrastructure and the applications are your own, you have the direct ability to manage them. A noisy neighbor on the public cloud is like living in an apartment building where you have limited options to deal with the partiers next door. On the other hand, noisy neighbors on your private cloud is like living in a house with unruly children that you can immediately manage by restricting resources.

A hybrid cloud strategy gives you the ability to put "sensitive" workloads and applications on private clouds and other workloads and applications on public cloud infrastructure for cost and resource efficiency. Many hybrid tools give you the ability to characterize workloads with QoS requirements to aid in the automatic, optimal placement of workloads on different cloud infrastructures.

Security and compliance

There are some dangers with doing things in an automated way in the public and private clouds. However, if security is built into the multi-hybrid architecture, this automation becomes a benefit because it means a security profile can be imposed across all your cloud assets, whether they are private or public, in addition to legacy infrastructure. This common security profile is applied everywhere, and applications are deployed based on those profiles.

For example, in a private cloud, if you want a specific type of security, such as tying a certain application to a certain machine that only runs in that private cloud, it can easily be a requirement in the profile on a multi-hybrid system.

In a multi-hybrid system, there are also many great tools for auditing and monitoring your infrastructure. You can not only monitor what is happening in your private cloud but also the public cloud, which will alert you to malicious attacks that could potentially infect your private cloud or legacy infrastructure assets.

Efficiency

There are many conflicting ideas about efficiency. If you look at various total cost of ownership (TCO) calculators, you will find different answers about whether private or public clouds are the most cost-efficient. A multi-hybrid cloud solution can solve this dilemma for your organization through heightened visibility and dynamic provisioning.

With a multi-hybrid cloud architecture and an intelligent orchestrator, your orchestrator can use telemetry from your private and public clouds and legacy infrastructure to make optimal decisions about where the workload should land currently and in the future. Another benefit of this visibility is that you can decide whether it's efficient for applications to run continuously. An example of this is how we helped the Canadian government's system to run more efficiently by cutting costs associated with running an application in the public cloud when no one was using it. Instead of running a specific application around the clock, they are now running it 18 hours, 5 days a week, according to actual use.

This visibility will also help you identify and eliminate end of life (EOL) workloads and applications, which will save real money. For the private cloud, this frees up resources that can be utilized for other workloads, in turn, driving up your efficiency. In addition, cloud brokers in the CMP in the multi-hybrid architecture will basically shop around for the lowest price while still maintaining the QoS for the specific workload. This decreases the overall cost of running the workload and also gives you visibility into your actual cost for using a particular public or private cloud.

- Visibility into costs across Clouds and Legacy infrastructure
- Drive workloads and applications to the lowest costs keeping the same Service level agreements.
- Drive higher utilization of private cloud infrastructure.

Call to Action

Multi-hybrid cloud architectures are giving CIOs the ability to get in front of the demands of their customers, but there is still some heavy lifting that has to happen. Building a multi-hybrid cloud strategy includes organizational, behavioral, and technical change that cannot happen overnight. Developing a strong architectural vision and roadmap are key to rolling out a multi-hybrid cloud strategy that can take advantage of multi-hybrid clouds' strengths and prevent the thrash of the technical industries' "shiny object" of the month inefficiency.

Guests

Meet Darren, the Chief Solution Architect for Public Sector at Intel. With over a decade of experience in executive and management positions, Darren has established himself as a trusted advisor to government organizations at the Federal, State, and Local levels, as well as enterprise organizations such as IBM, GE, and Toyota.

Darren's expertise lies in modernizing IT organizations, leveraging his unique ability to bring together technology, people, and processes to deliver transformative change. He is a firm believer in data transformation and focuses on data architecture, workload migration, cloud-native application development, service orchestration, and multi-hybrid cloud data center architectures to help organizations realize the benefits of digital transformation.



Darren's passion for technology and digital transformation is evident in his various contributions to the industry. He has eight patents in Cloud and Grid computing architectures, which have helped companies streamline product development lifecycle times through build, test, deployment optimization, virtualization, and containerization. Darren is also a published author with three books on technology and technology management, and has written over 100 articles in various industry trade publications.

As a thought leader in the industry, Darren shares his insights on his weekly podcast, "Embracing Digital Transformation," where he engages with industry experts and thought leaders to discuss the latest trends, challenges, and opportunities in the world of digital transformation. Through his work, Darren is making a significant impact in shaping the future of technology and driving real change for organizations worldwide.



Intel® technologies may require enabled hardware, software, or service activation. No product or component can be absolutely secure. Your costs and results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. © 2023 Intel Corporation