


WHY is DevOps so HARD?

The Guide to Overcoming
DevOps Challenges



NUTANIX[™]
YOUR ENTERPRISE CLOUD



Businesses today wishing to run as efficiently as an Amazon or a Netflix need to deliver products and services at the new, faster speed consumers have come to expect. Yet many don't have the strategy or tools in place to compete with these tech giants. Organizations must address barriers where the internal silos between people and technology have become the scourge of productivity.

This is all too prevalent when it comes to the relationship between developers, testers, and operators. The solution companies are beginning to adopt to solve the problem? DevOps. But what is it, exactly, how does it work, and is it as hard as it seems? This eBook will share the top challenges in adopting DevOps and provide solutions so you can avoid pitfalls on your adoption journey.

Let's dive in!





What is DevOps?

DevOps is the combination of cultural philosophies, practices, and tools that increase an organization's ability to deliver applications and services at high velocity; evolving and improving products at a faster pace than organizations using traditional software development and infrastructure management processes. This speed enables organizations to better serve their customers and compete more effectively in the marketplace. However, what these cultural philosophies and tools are and how to use them vary from company to company. It's no wonder so many organizations attempting to adopt DevOps are so confused!

Private Cloud



Distributed Cloud



Public Cloud



01

Which Cloud to Adopt?


Which Cloud to Adopt?

Today, the “public vs. private cloud” debate is less prevalent than it was even a year ago. IT teams want to harness all the benefits of a public cloud (agility, fractional consumption, etc.) while also maintaining security and control from on-prem environments. Data shows that **24%** of IT teams eventually mature past just one cloud. However, it can be difficult for IT teams to choose one over the other because public clouds tend to be expensive and both private and private clouds run the risk of vendor lock-in.

IT teams tend to turn to public clouds first because they are easy to implement and run. The catch is the cloud “sticker shock” — IT teams are often surprised by the cost of consumption when they get their monthly bill, but then it’s too late to switch as teams get stuck locked in to one expensive public cloud because their work is not reproducible in other clouds.

IT and DevOps teams are looking to become more “cloud smart” and re-evaluate their cloud-first strategy. To do this, they need to determine the right provider for each workload they run, starting with rationalizing their use case and application business needs appropriately.

It’s crucial for IT to have a firm understanding on how to consume public cloud resources in a blended, multi-cloud manner, which is the ideal situation according to **91% of businesses**.



A Cloud Smart Plan

Use a Cloud Management Platform

A comprehensive cloud management platform is key to optimizing resource spend in the public cloud and on prem. Using a platform like this can offer a cost savings of 30%; however, cloud costs could start to creep back up. The right cloud management platform can continually remediate, replan, and optimize resources.

Leverage Multiple Clouds

To leverage multiple clouds, it requires proactively managing workloads to plan their capacity and resources efficiently. In order to model workloads for cost and time, implement tools that focus on cost governance and compliance across cloud and on-premise deployments as well as other cloud services and self-service, quotas and orchestration capabilities.

Shifting Workloads

Plan, deploy, and manage workloads in a multi-cloud scenario. Once provisioning workloads with a cloud management platform has been adopted, IT can then decide which workloads to put on-prem versus in a public cloud. The Cloud Smart plan allows teams to get the highest performance and as close to the edge as possible with the widest use of resources between CapEx and OpEx. The ultimate end goal is to steer resources in a blended way over multiple vendors so there is no single point of failure and so teams can dial in their fiscal performance appropriately.

Provisioning analytics tools allows IT teams to deploy workloads on VMs and Docker containers across different hypervisors and public clouds. This inherent capability also allows IT teams to build a practice that is cloud-independent.

Cloud management platforms optimize usage and show IT teams when they need to rent or own their infrastructure, and also when they should spin up or spin down a certain cloud and shift from on-prem to off-prem when needed on demand. An example of this would be using 25% of AWS, 25% of Azure, and 50% of Nutanix—blending resources so there is no single point of failure and teams only pay per active workload use or per transaction. This gives you the best of both worlds—**public and private cloud**—when and where needed to be as flexible and as close to your customers as possible.

Enterprise Governance

Once IT teams acquire multiple clouds and toolsets, it becomes much harder to manage them all. Different vendors offer different management services, but they often don't work in tandem. This creates high failure rates; it's impossible to check on all the management tools for all vendors and get the big picture when it comes to metrics. The challenge for DevOps is how to manage all these vendors from one console. How do you know if you should increase or decrease usage of one set of tools or another? How do you measure the bottom line? As you've potentially heard before, "if you can't measure it, you can't manage it."

Entire businesses are also often fragmented and siloed with each team caring about their own area of responsibility for each workload. However, when they don't care about the big picture, that's when issues happen. Utilizing individual tools without a unified management platform doesn't go far enough to represent the whole business.

DevOps needs to break down silos between teams to work together and align for the best business outcomes and metrics.




SOLUTION:

One Pane of Glass

For DevOps to keep track of all their toolsets and clouds, they need to use one management system to rule them all. This single pane of glass experience enables end-user self service and administrative tasks to be conducted from one place across all platform services with role-based access controls and logging. Look for management tools that are vendor-neutral and tied to LDAP and SAML identity management sources.

By managing everything from one console, DevOps teams can eliminate the complexity of infrastructure management. IT teams can manage their entire environment, from storage and compute infrastructure all the way up to virtual machines, seamlessly and easily. With a single management tool, DevOps can streamline maintenance and upgrades, simplify workflows, and consolidate visibility into cluster statistics, making toolsets hassle-free, with no downtime or maintenance windows.

Monitoring clouds and toolsets is just as important as managing them. Using full stack monitoring can provide health assessments and performance analysis that can eliminate bottlenecks and prevent outages. Monitoring-as-a-service can also show DevOps teams minute-by-minute sales and buying trends. Having one monitoring alert on all toolsets eliminates the challenge of managing too much data. Without it the entire business risks being fragmented and siloed. **A monitoring service checks the health of the whole business, not just per team.**



Fiscal Governance

DevOps teams can get so consumed with buying all the tools and clouds they need to run their applications that they might overlook the cost effectiveness of them all. Before they know it, their budgets are blown, and CIOs are brought in to reign in spending. Teams need to manage their toolsets while keeping an eye on their budget but don't know the best way to go about it. Most cloud management platforms only evaluate public clouds, not the rest of the business. Teams choose cloud computing for its cost-saving benefits, but without having a full understanding of where money is being spent, it can quickly get out of control. By only managing the cloud cost, the company-wide budget can get thrown out of whack. The business as a whole needs to be fiscally governed with a DevOps mindset. **So how do IT teams know what cloud or tool is the most cost effective?**



SOLUTION:

Cost Optimization

With a cost optimization tool, IT teams can get **visibility and analytics** detailing cloud consumption patterns and one-click optimizations across both public and private clouds. This type of tool can also identify idle and under-utilized cloud resources, offer specific recommendations to resize infrastructure services, and use machine intelligence-driven algorithms to provide reserved instance purchase recommendations for continued cost savings.

IT teams should also provide advanced app-level orchestration across teams and clouds. To do this they should enable application total cost of ownership to compare with different deployment scenarios, providing responsible consumption choice to deliver simple, repeatable, and automated management of application creation, consumption, and governance. These cost optimization tools also excel at fiscal governance for public clouds and can help to optimize and secure their resource management.



Running Kubernetes

Kubernetes is the highly popular open-source system for automating the deployment, scaling, and management of containerized applications.

In other words, it's a system that makes it easier for IT to deploy and manage modern applications faster, for less money.

While traditional virtualization enables IT to split physical servers into multiple virtual machines (VM) to share resources, containers enable them to split up applications into smaller more portable pieces of code to more efficiently use the underlying infrastructure and move more easily across different clouds and internal environments. However, large applications can be made up of hundreds of different containers. Companies with lots of applications may find themselves with more containers than their development teams can possibly manage and maintain. That's where orchestration tools like Kubernetes come in.

However, because it's so new, most IT teams don't know how to use it. Everyone wants to test it out, but they don't know how. How easy is it to experiment, learn, run, and manage Kubernetes? What is it, how do I test it out? Why would anyone want to use it?



SOLUTION:

Kubernetes Infrastructure

Standing up and running Kubernetes by hand is difficult. It can take IT teams months to get it up and running, which can be very expensive. However, running it in the cloud as a service is fast and easy. At first glance the service appears to be an inexpensive add on, but the problem is most IT teams leave it on and running, increasing the cost over time. What IT teams should be doing is buying it on prem and letting it depreciate over time or testing it out as a service before committing to it and turning it off when not in use. This is much like leaving the light on when you're not at home. Every team needs to be using Kubernetes as a rental or utility and only have it on while using it, but many think of it as being "always on." For those IT teams that have yet to treat the public cloud as a utility, deploying Kubernetes on prem is the most cost efficient way for them to store and manage workloads.

With the right vendor, Kubernetes management can be vastly simplified, with easier upgrades and expansion, and persistent storage to enable database and file workloads, including logging, health, and metrics. But beware of Kubernetes vendors who break compatibility with the Kubernetes project by offering advanced features and proprietary tools, as it's a form of lock-in.

IT can leverage Kubernetes to abstract the complexities of the infrastructure at no additional costs. With highly available deployment support, a control plane, and reliable and performant storage for block and file through a Container Storage Interface (CSI), and object storage, this platform makes it simpler for anyone to deploy a Kubernetes cluster. It also offers integrated monitoring, logging, alerting, and zero down-time upgrades for Kubernetes cluster nodes as well as on-demand scaling. Running Kubernetes in the cloud allows developers to **easily operate web-scale workloads.**



Complex Database Management

Databases are a huge component of running a business. Large databases can be expensive and hard to manage and the typical Database Administrator (DBA) must provision, manage, refresh, restore, and perform other database operations for hundreds or thousands of database instances. The complexity and time-consuming nature of managing all these instances is further exacerbated when the databases run on a wide variety of legacy software and hardware technologies.

Provisioning each database requires considerations such as configuring compute as a single VM or multiple VMs in a cluster. Storage provisioning often requires multiple disk groups to handle different kinds of database data such as data files, software, operating systems, logfiles, and temporary workspaces. Once the DBA provisions the compute and storage and has a ready environment, the database server setup process starts with installing the database software. A clustered system requires installation, configuration and testing of additional components. The DBA must also protect the database by configuring backup policies requiring integration with different backup software and hardware technologies.

Once a database instance is provisioned, the complexity extends into cloning and data refresh processes as well. Cloning requires backup set identification along with any log files that are needed for clone creation. The DBA must first locate the backups (tapes or secondary sources) and then perform a complex recovery process that includes setting up the database server, connecting to the database, restoring database backups, and finally replaying the database logs to a specific time. Then the DBA must regularly refresh all these database copies and clones with the source data to be useful. **Now imagine scaling that effort to hundreds of databases to support different groups (test/dev, BI, QA, etc.) within the organization.**

05



Agile Database Development & Test

To help manage these monstrous databases, implement a hyperconverged infrastructure (HCI) solution that can sit on a cluster, which allows developers to get a copy of databases instantly to fix bugs. Without this, developers often get stuck waiting on a copy from the DBAs who are already busy managing the database. Instead of waiting, developers often take matters into their own hands by creating a dummy copy of the database to test for the fix, only to find out later that in reality the actual bug has not been fixed at all, or worse, has caused a performance impact with a production-sized workload.

By utilizing HCI, DBAs can automate and simplify database management, and bring one-click simplicity and invisible operations to database provisioning and lifecycle management (LCM). Copy Data Management (CDM) enables DBAs to provision, clone, refresh and restore their databases to any point in time. Through a rich but simple to use UI and CLI, they can restore to the latest application-consistent transaction. This software automates database provisioning and lifecycle management—slashing both DBA time and the cost of managing databases with traditional technologies.

HCI can make rapid copies of production databases, enabling self-service with easy provisioning, snapshots and restore, expansion and upgrades of databases without needing specialized teams, allowing DBAs to manage databases on top of a cluster. IT needs a production database so they always have extra copies and can survive failure. An HCI platform uses clusters to make sure data stays safe so DBAs can fix an entire node when needed. If all databases are built for failure, then a cluster of databases can be managed with **zero downtime**.



On Prem Scalability

Traditional IT has slowed down turnaround times for infrastructure requests to satisfy business, development, and test workloads for a number of reasons. Some of these are running out of storage capacity, compute, and memory limitations for the increasing demand of new workloads, which grow in size and complexity. Businesses are starting to experiment with big data analytics, machine learning, and cloud-native applications, all of which require dynamic populations and new application infrastructure on-demand. These applications flourish in the cloud world, because they can automatically scale up and down without a request for more resources via an IT helpdesk or service ticket for a human response.

Compounding the IT resource backlog, software developers and test teams want instant access to new application environments. With the advent of continuous integration and delivery, these environments could be as frequent or short lived as the time to test every new code change, which could be multiple times per hour. This huge potential demand for data and workloads with the ability to turn them on and off when they are being used requires self-service private cloud experience for applications and data. Correspondingly, IT wants scalability on prem without the hassle of installing and expanding different vendors' network, compute, and storage offerings in a coordinated planning and execution phase, which can be complex and require downtime. Finally, the management of any business wants predictability of their resources and desires a capital expenditure depreciation schedule versus the always growing, periodic operational expenditures of the public cloud, which can be very challenging to maintain.

How do you buy a public cloud experience but consume it in an on-prem manner?



SOLUTION:

Instant Private Cloud


The answer is hyperconverged infrastructure (HCI): physical infrastructure acting like a public cloud experience. IT teams don't need to plan, justify, and purchase capacity months or years ahead of time. They can buy HCI nodes, install them in as little as a few hours, and order them a month in advance when they need it. HCI cluster management should be able to clearly chart current capacity use and advanced workload modeling and help IT predict cluster capacity consumption and alert when reaching capacity thresholds. Adding more capacity to an HCI cluster should be as easy as unpacking the node, installing it into a rack in or adjacent to an existing cluster, then adding power and top of rack networking to complete the "rack and stack" experience in minutes.

With advanced network broadcasting enabled, the HCI management console can recognize the new node and ask to either join the existing cluster or create a new one, then work to auto-provision the new node into the cluster, grow the cluster to use the new storage, compute, and memory capacity. By automatically rebalancing workloads to take advantage of the full cluster expansion, the console minimizes and manages potential hot spot and noisy neighbor constraints on workloads. HCI gives IT teams an Apple-like, Lego brick experience where they can **deliver a public cloud experience for the private cloud on prem**. This becomes the required underpinning for on-demand, self-service workloads.

Brownfield and Existing Workloads

IT teams can have a difficult time managing their workloads, whether on-prem or in the public cloud, because they lack the management capabilities to identify and control the lifecycle of applications and their implied infrastructure. With every cloud-native workload, on-demand autoscaling resources are expected, but most IT teams can only provide VM, and they require a manual approval and provisioning process via an IT ticket. The IT service or help desk manual process provides a documented business owner request, potential show-back or charge-back to their business unit, and resource inventory management in a change management database (CMDB).

IT and their business need better visibility into how applications are working and consuming resources across the entire datacenter as well as in the cloud. In addition to requiring a single pane of glass view for all workloads, they should also have visibility into how their servers are working. Oftentimes, IT has no idea how its data is consumed, so they delegate that responsibility to other teams, such as DBAs and storage teams. This siloed effect means each team is focused just on their slice of running the business, and in aggregate, IT merely is working to keep the lights on to satisfy traditional operating models rather than having the opportunity to innovate.



SOLUTION:

Effortless Monitoring

Utilizing a cloud monitoring service can enable an overview of all your workloads with simple monitoring and alerts for changes on the network and application environments, giving developers, IT, and management a **single pane of glass view**. Placing agents on every infrastructure element that powers a workload yields the insight to everyone in the datacenter, and brings visibility not only to network flows, but can be grouped into application demands and monitored for business health. As application performance becomes increasingly essential to the business, organizations are seeking a full-stack monitoring service to provide health assessments and performance analysis to eliminate bottlenecks and avoid outages.



Conclusion

To cope with the fast moving IT world, companies need to increase their pace, shorten work cycles, and improve delivery efficiency. To achieve this, IT teams need to implement a DevOps strategy to give their business a much needed boost. The right DevOps strategy recognizes the interdependence of software development and IT operations and helps companies produce software and IT services more rapidly, with frequent iterations. To learn more about a high-performance infrastructure that will help make DevOps a reality, check out nutanix.com.

