



BUILDING AN INTERNAL DEVELOPER PLATFORM WITH NUTANIX

By Brian M. Posey

NUTANIX



www.nutanix.com

Since the release cadence of software innovation began accelerating a few years ago, from years to months or even weeks, the demands on organizations to provide developers with a well-architected internal developer platform has also been accelerating. In a recent webcast, Jose Gomez, technical marketing engineer at Nutanix, talked about ways to provide developers with self-service access to the resources they need and how infrastructure as code can assist with automation, helping developers achieve greater efficiency, consistency, scalability, and reliability in their operations.



THE EMERGENCE OF SELF-SERVICE INFRASTRUCTURE

Gomez explained that self-service infrastructure is designed to solve an age-old problem. In the past, the user would request resources and it would take several weeks for those resources to be provisioned because the service desk had to communicate with various other teams in order to fulfil the request.

With time, these processes have improved with automation as a way of accelerating service delivery. For example, whereas IP address management once involved using an Excel spreadsheet, now a team might use an API that is tied to an IP Address Management (IPAM) solution so that they are able to orchestrate the IP provisioning process.

Self-service Infrastructure takes this orchestration one step further. Automations are made available through a marketplace so that end users are able to access those automations and provision resources without having to involve the service desk. According to Gomez, this can reduce the provisioning time from weeks or months to hours, minutes, or potentially even seconds.

Another reason why self-service capabilities are so important is because they help prevent the use of shadow IT. Gomez explained, “The infrastructure teams of the past were not listening to developers when they explained what they needed, and that was how we ended up with shadow IT. Shadow IT came about because developers were trying to get access to the resources that they needed in order to deliver what the business was asking for.”

THE SHADOW IT PROBLEM

Shadow IT is once again becoming a problem as organizations are changing their cloud strategies. According to Gomez, several years ago, most companies embraced a cloud-first strategy and many signed long-term commitments with cloud providers. As time has gone, on however, companies are starting to realize that using the public cloud is not always the best option. Organizations are transitioning to hybrid multicloud and many workloads are being repatriated. The problem with this is that developers have gotten used to the self-service capabilities provided by the large

cloud providers. Gomez said, “If we want to bring workloads back on-premises, then we have to do things differently to how we did it in the past, because if not, we are going to be on the same situation as before with regard to shadow IT.”

Gomez cited a recent [Gartner study](#) that found that by 2027, most workloads are going to be misplaced, meaning that those workloads are not going to be running in the best location based on the workload’s needs. This study clearly indicates that self-service portals are going to need to do more than simply orchestrate deployments in the same way that has been done in the past. Gomez says that what is needed is a single self-service portal that can deploy applications to several places. If, for example, an application is only going to be used for a few hours, then it probably doesn’t make sense to deploy that application on-premises, whereas an on-premises deployment might be best if the application will be used long term.

THE NEED TO STANDARDIZE ON A COMMON SET OF TOOLS

The need for such a portal isn’t just about preventing shadow IT. It’s also about standardization. Gomez said that in many organizations, the infrastructure team will use one set of tools for deploying workloads on-premises, while a DevOps team uses a completely different set of tools to deploy workloads to the cloud. Gomez explained, “Most of time, it’s not just two different tools, but actually three or four or five. So organizations are trying to look how they can [standardize]

around a common set of tools based on both teams’ knowledge and expertise.”


In order for self-service to succeed, self-service capabilities must be bound to a governance framework. According to Gomez, “This governance framework must delight the end users, but it must also take each team’s needs into account and give the various teams the sense that they are in control and that they have visibility into everything that is happening.”

ADOPTING INFRASTRUCTURE AS CODE

One of the most important things an organization can do when deploying self-service capabilities is to adopt infrastructure as code. Infrastructure as code is simply a template that ensures that infrastructure is deployed in the prescribed manner. According to Gomez, history has shown that when developers manually deploy infrastructure, they occasionally make mistakes. They might accidentally open a port that they shouldn’t or forget to click a particular box. Gomez went on to say, “If experienced developers occasionally make mistakes when manually deploying infrastructure, just imagine what is going to happen when you open those self-service capabilities up to your end users.” Adopting infrastructure as code takes the potential for human error out of the deployment process.

Gomez went to say that the best way for companies to begin approaching self-service for their organizations is to engage with developers and ask them to share their requirements. From there,

More than Just VMs



organizations should create roadmaps and milestones. These milestones might initially involve automating individual services such as IP address management. You can then use these milestones to build your marketplace. As Gomez said, "If you try to do everything at once, your milestones become a two-year project and the business will start to complain and wonder what is taking so long. Small, achievable milestones offer visible progress as each milestone is met."

NUTANIX BRINGS CONSISTENCY ACROSS ENVIRONMENTS

Nutanix provides customers with a consistent experience by creating software that can run at the edge, in the datacenter, or in the public cloud. This means that the way that users deploy resources on-premises is the same as how users deploy resources in the cloud or at the edge.

Of course, Nutanix understands that organizations already have resources running natively on public clouds (as opposed to running on the Nutanix operating system). Even so, the Nutanix

Cloud Manager integrates with existing tools such as TerraForm to be able to manage all of those applications through a single pane of glass.

Nutanix Cloud Manager can be also consumed through infrastructure as code. Additionally, Nutanix allows for data portability. Gomez explained, "We talk about application portability with containers, but a container is just an application package and the data is completely decoupled. Nutanix lets you move data from the edge to the datacenter to the public cloud, or the other way around, almost as though that data were containerized." This is the way that Nutanix is able to provide a single platform across any location for any application and for all of your data.

TO LEARN MORE about Nutanix approach to cloud native infrastructure, visit Nutanix at: <https://www.nutanix.com/solutions/cloud-native>

NUTANIX