GRAVITATIONAL

Scaling Managed Services in the Cloud Era

BY GRAVITATIONAL

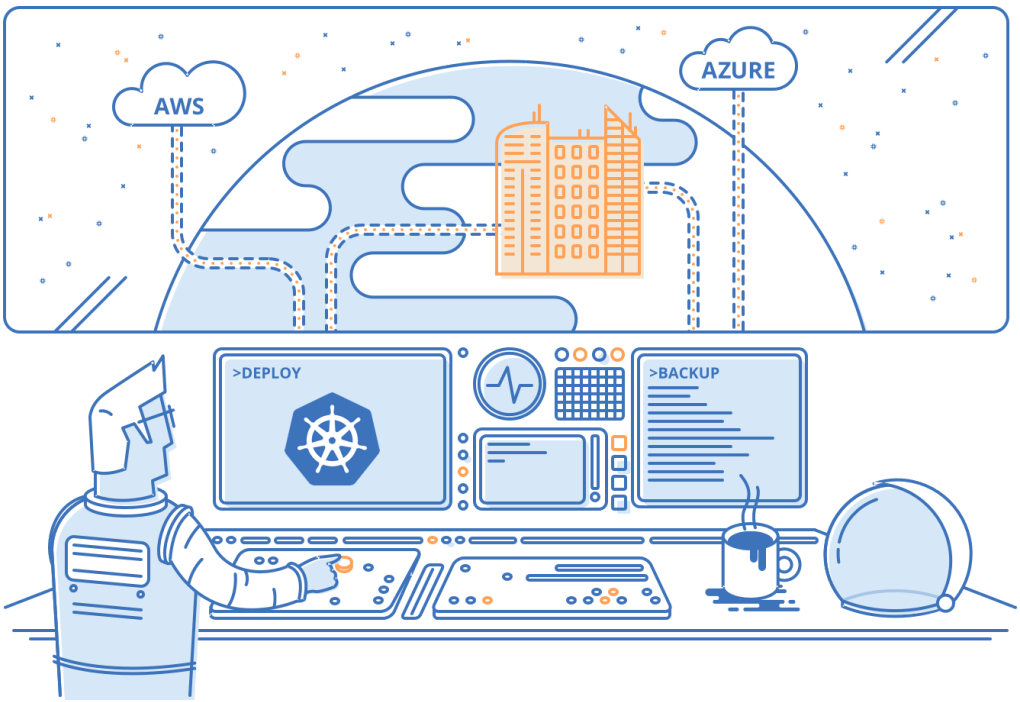

Abstract

The rapid adoption of cloud services is a once in a generation, multi-billion dollar shift for the IT industry. In this paper, we discuss some of the biggest challenges the adoption of Infrastructure as a Service ("IaaS") and Software as a Service ("SaaS") pose to IT departments and application developers; and how Managed Service Providers need to adapt to stay relevant.

info@gravitational.com | 855-867-2538

GRAVITATIONAL

Table of Contents

GRAVITATIONAL

IaaS and SaaS Adoption Create Complexity for MSPs

The rapid adoption of cloud infrastructure or infrastructure as a service ("IaaS") has brought about a revolution in the way companies manage their IT. Everything from how they access third party applications to how they develop their own internal software is rapidly changing. More and more applications and internal workloads are moving from internal data centers to third party public and private cloud providers.

> "Worldwide public cloud services market will grow 18% in 2017 to $246.8B, up from $209.2B in 2016."
>
> *Gartner: http://www.gartner.com/newsroom/id/3616417*

The proliferation of IaaS and software as a service ("SaaS") providers has helped to reduce the difficulty of adoption, significant upfront investment and high cost of maintenance that traditional on-premise software required of IT buyers by shifting the hardware procurement and software operational burden to the vendors. According to Gartner, "The worldwide public cloud services market is projected to grow 18 percent in 2017 to total $246.8 billion, up from $209.2 billion in 2016...The highest growth will come from cloud system infrastructure services..., which is projected to grow 36.8 percent in 2017 to reach $34.6 billion."

These benefits come with some trade-offs to the customer, including less data control, data dispersed across locations and services, inconsistent and unknown security paradigms, usage and procurement tracking issues and many others.

However, it is clear that the movement of IT workloads and services to IaaS and SaaS providers is a once in a generation shift in how IT budgets are spent and the type of managed services companies need from their partners and vendors. In order to stay relevant, Managed Service Providers ("MSPs") need to be nimble and update the types of services they offer and the way they offer them to fit in with this new paradigm.

In this paper we will discuss how the dispersion of infrastructure and applications from internal data centers to many different locations and IaaS providers changes the landscape for MSPs. We will do this by covering three relatively recent innovations that require a shift in the way MSPs deliver their services:

1. Companies are now using a hybrid of internal data centers, private clouds and third party IaaS providers to deliver IT, internally and to its customers.
2. Repeated breaches behind company firewalls have reduced the significance of perimeter-based network security policies and lead to "zero-trust" network security policies.

info@gravitational.com | 855-867-2538

3. Microservice-based application architectures are being adopted by developers in order to keep applications resilient and efficient when using dynamic cloud infrastructure.

Multi-Region, Hybrid Infrastructure

IT infrastructure spend is rapidly moving from traditional data centers owned or leased by the company to public cloud and private clouds run by third party IaaS providers. According to IDC Worldwide Quarterly Cloud IT Infrastructure Tracker, IT spend on Public and Private Clouds was $36 billion in 2016 and is now approaching that of traditional data centers.

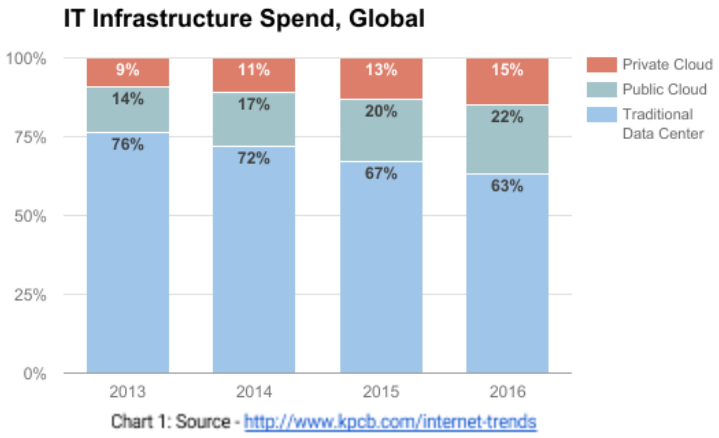

Chart 1: Source - http://www.kpcb.com/internet-trends

The obvious ramification of this is that MSPs that help IT departments manage their internal data centers are seeing their budget allocations getting replaced with IaaS provider spend.

However, these IaaS providers are not as always as turnkey as they would like their customers to believe. In addition, they are notorious for offering little or no support. This has allowed a new generation of MSPs, like Datapipe and 2ndWatch, to create successful businesses that focus on helping companies manage their IaaS accounts, especially on Amazon Web Services ("AWS"). Even companies historically focused on managing infrastructure for their customers, like Rackspace, have successfully expanded their offering to include providing support and services on third party cloud providers.

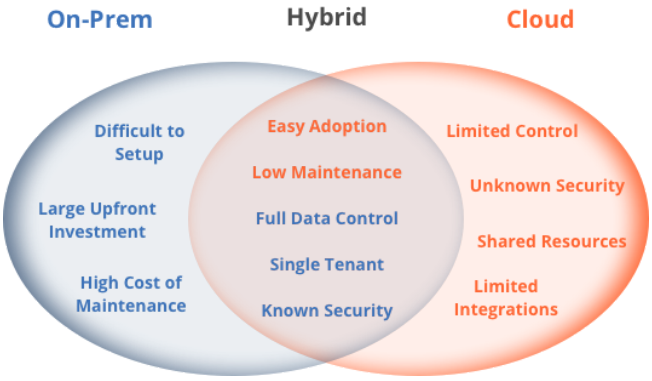

The other ramification is that IT workloads that were traditionally located in a single data center are now located across many different locations and third party vendors. This presents an opportunity for nimble MSPs as IT professionals and software developers need help adopting "cloud native" best practices for their applications so that they can run across a variety of environments. In addition, once deployed across a variety of environments, application owners need help running them in a scalable manner.

This requires tools that can be used across a variety of infrastructure providers that have different network topologies and security regimes. In order for MSPs to efficiently offer services across all these different environments, they need a unified access point that will work across company data centers (behind firewalls) and third party IaaS providers. In addition, this access needs to be compliant with new, modern information security practices. Which brings us to our next topic…

Zero-Trust Network Security

Traditional network security has relied heavily on firewalls and other perimeter security measures to differentiate between an untrusted external network and a trusted internal network. This allowed trusted people and services, behind the firewall, to have relatively unrestricted access to the trusted network. However, frequent and costly breaches of this security model, either by social engineering, disgruntled employees or sophisticated, state-sponsored attacks, have resulted in a new security paradigm that reduces reliance on perimeter network security and assumes that intrusions by bad actors will occur by default.

Google has been a pioneer in this movement through their BeyondCorp approach to enterprise security which states as its mission, "to have every Google employee work successfully from untrusted networks without the use of a VPN". Google has also rolled out this security approach as a feature on its Google Cloud Platform.

A core principle of this model is to use an identity-aware proxy or bastion as an entry point to infrastructure, applications and services available to trusted people and services. These proxies can be integrated with other identity management services that manage trusted identities and the permissions allocated to those identities, also referred to as role-based access control ("RBAC").

Another key principle is to use short-lived authentication so access must be affirmatively granted to prevent unauthorized access through stale authentication. In addition to this, multi-factor authentication and end-to-end encryption on internal networks is enforced.

This model presents challenges for MSPs that provide infrastructure or application based services. Copying and pasting SSH keys on servers is no longer an acceptable practice and the use of traditional VPN solutions becomes impractical as a company's infrastructure spans internal data centers and third party cloud providers.

An additional complicating detail is that many modern applications are now architected to have dynamic, ephemeral services. Which brings us to our last topic…

Microservice Application Architectures

Traditionally, accessing servers in order to manage applications was done through “SSHing” into servers based on IP address. This generally worked because servers were more stable and monolithic applications were statically located on a known server or set of servers.

The adoption of cloud infrastructure has required changes to these assumptions due to the use of more volatile but cheaper virtual machines and the practice of “bursting workloads” to additional servers in times of heavy load. This has led to the creation of application architectures that are designed to be resilient for more dynamic infrastructure, like the twelve-factor application approach.

More recently, the adoption of more dynamic building blocks like software containers and software defined networking has led to the decoupling of application services even further with microservice application architectures. Microservices have been made popular by large scale internet companies like Netflix, Amazon and Ebay. This approach achieves more resiliency and flexibility at the expense of additional complexity.

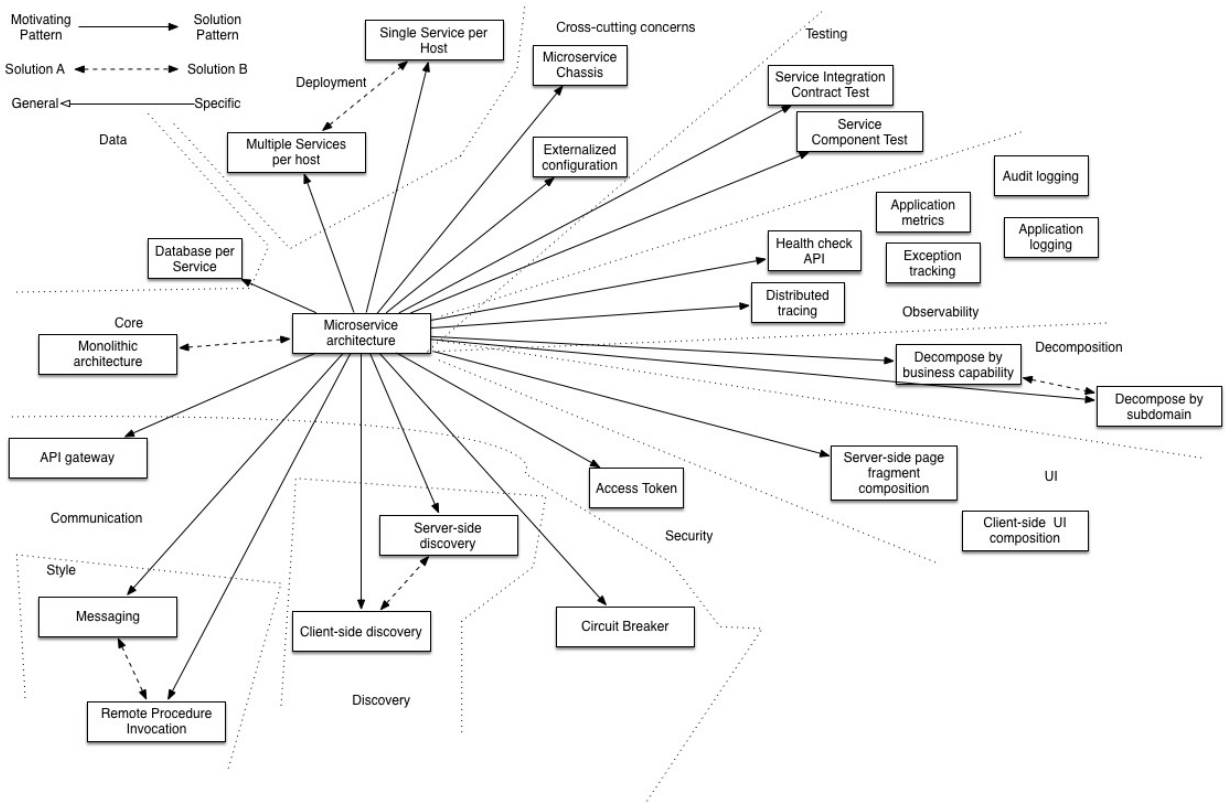

Diagram 1: Microservice Patterns. Source - http://microservices.io/patterns/microservices.html

Due to the dynamic nature of these new architectures, various types of automation tools have evolved to manage “infrastructure-as-code”. The general idea is to treat your infrastructure as immutable - meaning that you make changes by creating and deploying another instance of your application instead of modifying the running application. This practice is designed for the dynamic nature of cloud infrastructure but requires more sophisticated tooling and operational practices.

Configuration automation tools like Chef, Puppet, Ansible and Saltstack can be used to automate the installation of applications on existing servers (bare metal or virtual). Traditional orchestration tools like CloudFormation and Terraform can be used to provision servers and other resources on IaaS providers. These tools are generally designed for when you have a similar infrastructure environment to automate. For example, if your application is running on AWS or on various AWS regions.

However, for larger companies that have investments across internal data centers, colocation facilities and multiple cloud providers, operating system level virtualization through software containers can provide a further abstraction so that applications can run across all of these environments without needing configuration or application code modifications.

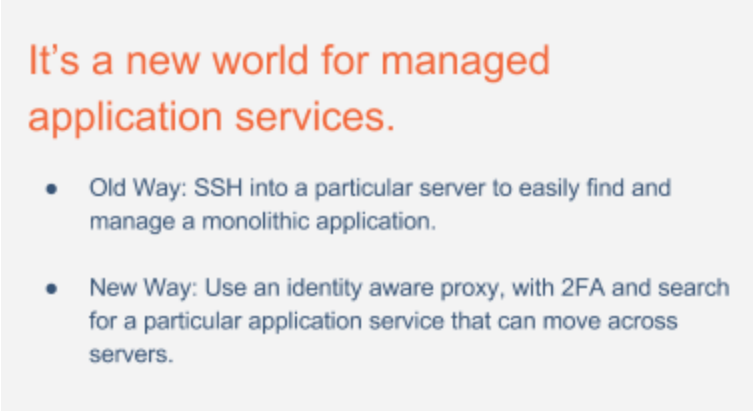

The use of these containers has led to another evolution in how applications are developed and managed. A new crop of automation tools, called container orchestration systems, have evolved, which include Docker Swarm, Mesosphere and Kubernetes. The promise of containers and their orchestration tools is that an application can be deployed and managed across a variety of infrastructure footprints and the application environment remains relatively consistent. This approach also theoretically allows for more efficient server density because application services can be easily moved and allocated where infrastructure resources are available.

This means that applications are broken out into small components located across servers, regions or even infrastructure providers and these application services may be moving across these different locations. This creates new challenges for internal IT and outside parties that help remotely manage these applications.

Conclusion: New Tools and Best Practices Are Needed For MSPs

The rapid adoption of cloud services alters the IT landscape, changing everything from how budgets are allocated to how applications are developed and network security is treated. Managed Service Providers need to modernize their workflows and tools in order to adapt to this new paradigm.

- As applications become more distributed across geographic regions and public/private cloud providers, MSPs need unified access that is designed to work across all environments.
- This unified access needs to be identity aware and compliant with new, modern information security practices, like Google's Beyond Corp, which do not rely on perimeter network security.
- Increasingly complex application architectures require management tools that are designed for microservices and can locate and follow dynamic services as they move across environments.

The MSPs that can adopt workflows and tools designed for these new approaches will be the ones poised to take advantage of the new cloud era.

GRAVITATIONAL

About Gravitational

At Gravitational, we build automation tools designed from the ground up for modern infrastructure environments and application architectures. Below is a brief description of these products with links for additional information.

Multi-Region Automation Tools

POWERED BY GRAVITATIONAL

Gravitational is a Y Combinator® company founded in 2015 by alumni from Rackspace® to create the ultimate platform for managing multi-region applications. The founders previously created Mailgun, which was acquired by Rackspace in 2012. At Rackspace, they launched the On-Metal product line and created open source libraries such as Vulcand.

info@gravitational.com | 855-867-2538