White Paper: **February 2023**

Building a Successful Cloud Operations and Governance Practice

Table of contents

Executive summary

Digital transformation and the explosive growth of cloud services are forging an equal transformation in IT and cloud operations. In the cloud era, IT operations has the opportunity to elevate its organizational stature and influence business strategy and outcomes. It no longer should be limited to being solely a simple fulfillment center for end users' IT service requests or a maintenance engine for data center infrastructure. With the opportunity in mind, cloud operations teams must build a framework with best practices that balances the need for innovation with optimizing time, reducing costs and minimizing risk.

This white paper provides tips and best practices for cloud operations leadership and administrators focused on efficient, scalable cloud operations that accelerate their organization's digital transformation in the cloud.

Key learnings include:

- How IT and cloud operations have evolved in the cloud era
- The role of cloud operations teams and the need for an organizational culture of continuous governance
- The different steps in the operational governance maturity model and how to measure success with relevant key performance indicators (KPIs)

The evolution of infrastructure and operations in the cloud era

Enterprises have adopted public cloud to improve agility, gain more flexibility, and innovate rapidly. Agility, flexibility and rapid innovation are possible due to the democratizing effect of easy-to-consume, self-service public cloud services. Unshackled from tightly controlled internal data center infrastructure, application developers and software engineers have an unprecedented ability to spin up new resources and deploy new code in the cloud faster than ever before.

The ease of cloud services and the ubiquity of agile development processes and continuous integration/continuous delivery (CI/CD) frameworks are now, without a doubt, the foundation of an enterprise's business and digital transformation. But it's also driving another transformation: a necessary re-imagining of the role of infrastructure and operations (I&O) teams.

The emergence of decentralized IT

With the exception of companies born and raised in the cloud, nearly every enterprise has an on-premises IT infrastructure in addition to their increasing use of public cloud services. As a result, most enterprises have a strong centralized IT team that responds to end-user requests for internal data center resources and other IT services. In a decentralized model, there can actually be multiple IT teams throughout an organization. As an example, each business unit could have their own IT team to help their users and developers with provisioning cloud services. The ubiquity and accessibility of cloud services lends itself to a more decentralized IT delivery model. While there is overlap in their management tasks, organizations increasingly need to support both centralized and decentralized models.

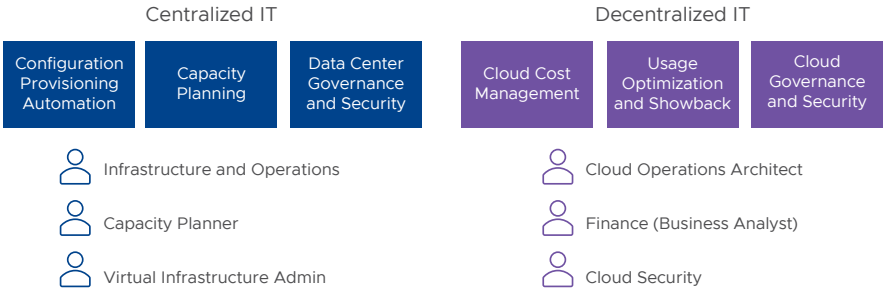

Figure 1: Centralized and decentralized models of IT delivery.

The centralized model of IT delivery requires specific operational tasks, such as configuring and provisioning of infrastructure, planning for capacity, deploying patches and updates, and managing technology refreshes. Most enterprises have been executing these tasks for years and have successfully automated a significant portion of this work. In addition, the roles and responsibilities of I&O teams and leaders is well-defined in a centralized model.

In contrast, operationalizing a cloud-centric, decentralized IT delivery model where users can access self-service resources requires additional skills and tasks that focus more on financial management, optimization and governance.

Kiran Kulkarni, VP for data science and engineering at German media corporation ProSiebenSat1, can attest to this new and unique challenge. “We went from a very controlled IT-run, on-prem server and on-prem development model, which we thought was very complex and was slowing us down, to another extreme where it is a free-for-all, and that is equally complex,” says Kulkarni. “We need to find a middle way—cloud cannot be another word for anarchy.”

Operations for public cloud, therefore, need to transform from a fulfillment function for user requests to one that governs users’ self-service. As it’s still a relatively new discipline, enterprises vary wildly in how mature their cloud operations are.

The critical role of operations in the cloud center of excellence

To bring some centralized stewardship of decentralized public cloud usage, organizations are establishing a formalized cloud center of excellence (CCoE), sometimes known as a cloud strategy office or a cloud program office. The CCoE is a cross-functional working group that governs the usage of the cloud across an organization, driving best practices across functions.

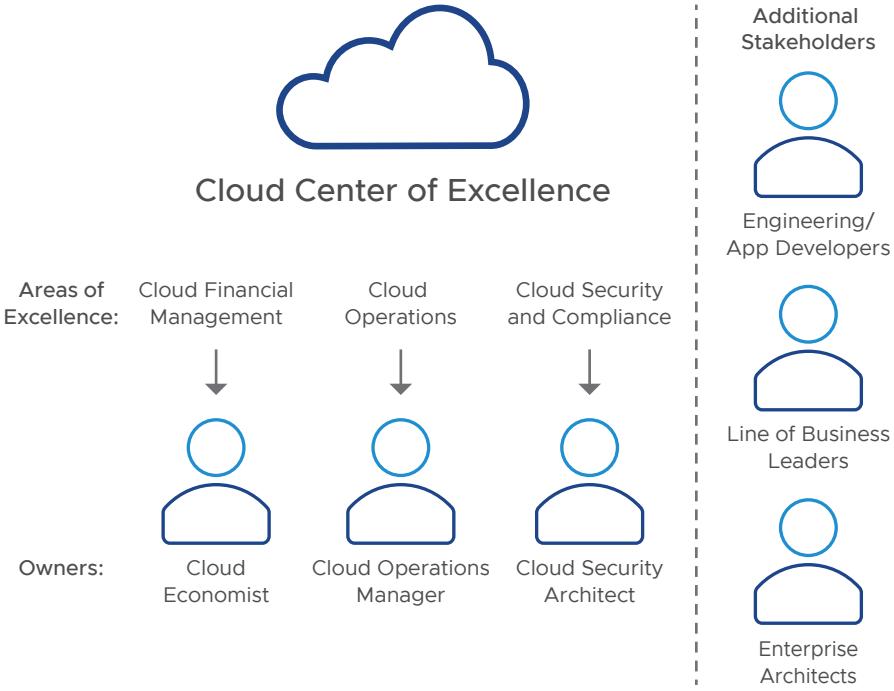

Figure 2: Example of a cloud center of excellence.

The three pillars of a successful CCoE are cloud financial management, cloud operations, and cloud security and compliance. This white paper focuses specifically on cloud operations and its influential role in the maturation of an organization's cloud strategy and digital transformation.

Cloud operations explained

The goal of cloud operations, within the shared responsibility and security model of public cloud providers, is to ensure optimal functionality and performance of cloud platforms and associated workloads for their full lifecycle. The key term is optimal. In the centralized IT model, the metrics for optimal focus almost exclusively on reliability, availability and scalability (RAS). Centralized IT maintains and optimizes RAS by keeping tight control over resources. With public clouds, that approach is simply not possible and likely a fast track to failure.

In the cloud operations world, RAS is almost exclusively a function of the cloud provider and outlined in the provider's terms of service and service-level agreements (SLAs). The role of operations in the cloud era must therefore shift from RAS to governing users' self-service access to public cloud services. This shift to an operational governance model is a new challenge that requires new skills and new thinking. And once an organization adopts multiple clouds, the challenge becomes greater because there are no standards across cloud providers for provisioning, reporting and administration.

Operational governance is more than just protecting the organization and enabling self-service access. The more mature cloud operations becomes, the easier it is to enable users to increase adoption of innovative cloud services, such as serverless computing, container services, managed databases and other new platform services. Mature cloud operations make embracing new technology less stressful for operations staff and facilitates adoption of higher-value services.

A mature cloud operations function will:

- Monitor and report on cloud usage and configuration by users and teams across the organization
- Establish standards and policies for proper cloud platform configurations and usage, and take action when configurations and/or usage drift from the desired state
- Maximize operational efficiency and maintain the desired state via automation
- Establish continuous governance that integrates with and enables DevOps and CI/CD processes for rapid application deployment

The risks of failing to establish a mature cloud operations function can include:

- The prevalence of shadow/rogue IT as developers and users provision their own cloud services with no accountability or oversight
- Exploding cloud costs due to lack of visibility into cloud usage by teams/departments and overprovisioning
- Undetected security threats resulting from unsanctioned services and misconfigurations
- Paying down technical debt instead of focusing on innovation and new productivity-enhancing features
- Slower time to market for new and migrated cloud applications

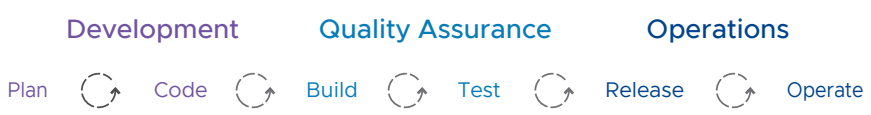

Figure 3: The development lifecycle.

The traditional I&O department is often unfairly treated as a functional silo. They only interact with finance teams when procuring new infrastructure and rarely connect directly with business stakeholders. On top of that, they often aren't involved in the early stages of a CI/CD process and are only brought in at the tail end when they are tasked with getting releases into production.

In contrast, cloud operations cannot function in a vacuum. As cloud grows, so too have DevOps practices and an urgency to shift left. Shifting left means shifting release practices that are typically at the end of a development lifecycle, such as testing or security, earlier in the process. The goal is to make the eventual release into production smoother with few application rollbacks.

Cloud operations and governance need to embrace and facilitate this shift left. A CCoE, by its nature, is a very cross-functional team and the cloud operations aspect is no exception. Cloud operations is tasked with developing and then socializing best practices for cloud usage and working with finance, security, development and architecture teams to embed those best practices in multiple processes to facilitate continuous governance.

In the example of a CI/CD process, policies for proper cloud configuration and usage should be embedded in the process from the beginning (shifted left) and be continuous. Moving to the cloud is one thing. Operating efficiently over time and maintaining alignment with business goals relies on embedding operational best practices (such as tagging hygiene) across the entire organization, and that is a much larger and more complex undertaking.

Meet the head of cloud operations

The head of cloud operations can come from different backgrounds. They may come from a traditional I&O role or a cloud architecture or engineering role. Increasingly, the head of cloud operations may not come from any kind of operations role at all—business acumen and organizational skills are paramount. Again, they are not focused on operating infrastructure; their focus is enabling self-service access by multiple users in a safe and cost-efficient manner.

In a medium to large enterprise, the head of cloud operations typically reports into the VP of I&O and oversees a team of cloud administrators. The team works closely with site reliability engineers (SRE), cloud architects, cloud engineering and cloud security teams to determine what a desired state looks like and create a methodology for maintaining that state.

Building efficient and scalable cloud operations

Building a culture of continuous governance

The head of cloud operations needs to develop a culture of continuous governance that is designed for the self-service consumption of as-a-service solutions—software as a service (SaaS), platform as a service (PaaS), infrastructure as a service (IaaS)—by diverse teams. Continuous governance balances the need to quickly deploy new workloads to the cloud with protecting the organization from harm. This shift is often more about organizational and procedural change than technology change.

As Lukas Lehmann, head of cloud services at Swisscom, explains, "I think this is a crucial part of the cloud journey: to have guidelines that enable the internal development teams and the business to rethink their approach, and enable them to move to the cloud in the right way. Because if you use it wrong, and it's easy to use it wrong, then things go badly very fast."

An important role of cloud operations in a multi-cloud world is to provide line of business (LOB) leaders (depicted in the top row of Figure 4) with visibility into how their users are using cloud resources. LOB leaders also need to collaborate and communicate more effectively so that they, and the organization as a whole, optimize their use of cloud services and prevent bad things from happening.

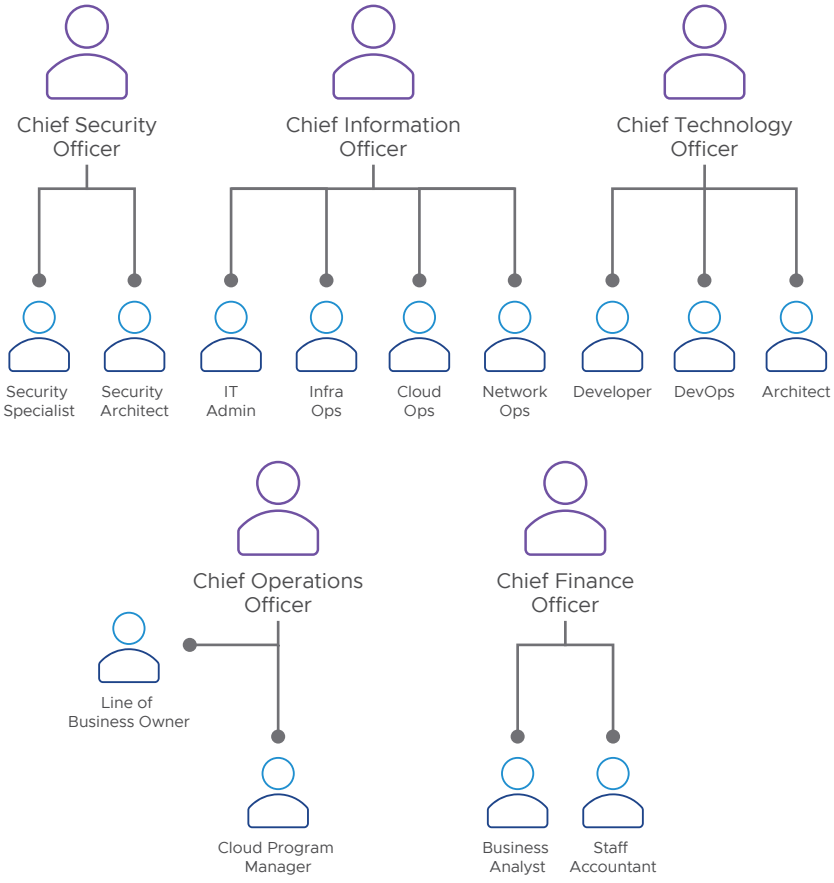

Figure 4: Organizational chart for cloud operations.

Cloud operations needs to strike a delicate balance between giving cloud consumers what they need exactly when they need it, and putting rules in place to govern usage. Continuous governance defines best practices, socializes them, then takes action when a policy or standard is violated. There are several methods for accomplishing continuous governance.

Create guidelines and guardrails for efficient cloud operations

There are two methods that an IT operations team can use for cloud governance. The first method is for IT operations to act as a cloud broker where they curate all sanctioned cloud services and act as a fulfillment engine for user requests. This has the advantage of locking things down, but it limits choice and speed. Additionally, with the number of changes public cloud service providers can roll out in a year (for example, AWS pushes more than 2,000 updates per year), IT simply can't keep up. The second method is referred to as "on the side," where IT allows users to provision cloud services directly from the cloud providers but with controls to protect the organization.

This white paper focuses on best practices for the "on the side" model as this creates the most complex challenge to I&O teams. Continuous governance in this context can be implemented in several ways, each with its own strengths, weaknesses and use cases.

In-band

An in-band policy is evaluated before a user takes an action that would potentially violate best practices. The advantage of this approach is that it prevents users from taking actions that could be dangerous or expensive. The disadvantage is that it can potentially hamper user productivity and cause them to go around corporate IT.

Out-of-band

An out-of-band policy is evaluated after a best practice violation is detected. The advantage of this approach is that it stays out of the way of users and performs cleanup after the fact. Conversely, the disadvantage is that it allows users to violate best practices, and if they aren't fixed in a timely manner, these violations can be damaging.

In addition to in-band and out-of-band detection, there are two primary ways companies can take action on governance policies:

- Guidelines – Guideline policies will communicate a risk boundary that informs the user of the best practice but will not take action to prevent or correct the action.
- Guardrails – Guardrail policies will both communicate and take action to correct a violated best practice.

Figure 5 provides examples of what happens in each scenario when a user violates, or attempts to violate, a best practice.

	In-band	Out-of-band
Guideline	Before deploying, the user is notified they are violating a best practice and provided with instructions on how to properly deploy the resource. The user can ignore the notification and proceed with deployment.	After deploying, the user is notified their recent deployment violated a best practice. The notification includes instructions on how to fix their mistake and conform to the best practice.
Guardrail	Before deploying, the user is notified they will not be able to deploy until they fix the violation, or the violation is automatically corrected before deployment.	After deploying, the user is notified their recent deployment violated a best practice and that it has been automatically fixed.

Figure 5: Examples of violations of best practices.

As a company matures in its cloud adoption, cloud operations teams should implement in-band and out-of-band guidelines and guardrails in a progressive fashion that builds on established best practices and constant learning. There are several published frameworks that organizations can use to guide their development, including the AWS Well-Architected Framework and the Microsoft Azure Well-Architected Framework. Each is specific to their particular cloud platforms, but both offer best practices that apply generally.

A maturity framework for efficient and automated cloud operations

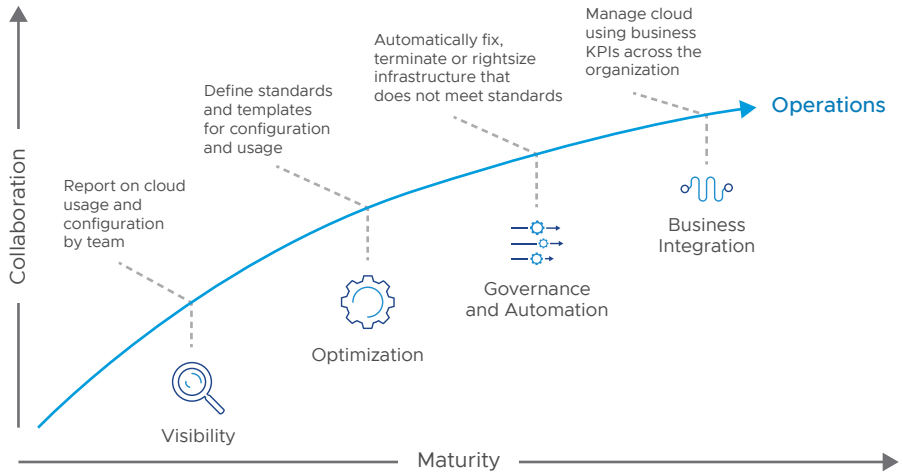

Figure 6: Cloud operations maturity framework.

Phase one: visibility

You can't manage what you can't see, so the first obvious step is to gain visibility into all the cloud resources in use across the organization. From a cloud operations perspective, visibility focuses on usage, configuration, performance, security and tagging hygiene. Some best practices to get started with in this phase include the following.

Discovering and linking cloud accounts

A first step is to find and link disparate and individually owned cloud accounts so you can centrally view who is using what and how much they are paying. You may save considerable money by simply combining purchasing power and taking advantage of discounts offered by the cloud provider. You should also establish policies for new account setup and work with cloud security teams to migrate shadow accounts and comply with identity and access management (IAM).

Establish good tagging hygiene

It's critical to establish standards for tagging all cloud assets. Without proper tagging, cloud operations, cost analysis, and security management are much more difficult and less effective. Each cloud provider's tagging rules vary, so it's important that you either establish individual tagging rules for each or use a third-party cloud management platform that can apply groupings that comply with each cloud provider's rules. Without proper tagging, it's nearly impossible to know which assets belongs to which teams and applications, or to report on costs.

Group all cloud assets by teams, owner, application and business unit

With good tagging hygiene established, you can then group assets in a context that's meaningful to the key constituents in the organization. For example, you can group them by cost center, LOB, owner or application.

Identify misconfigured and non-standard assets

Using a Well-Architected Framework, you can compare how well your cloud resources are adhering to configuration standards and best practices. Examples include non-standard virtual machine (VM)/compute instance families, old or non-standard machine images and operating systems, or deployment to non-approved regions. This visibility shows you cloud assets with non-standard configurations or those provisioned sub-optimally, so you can notify the owner with instructions for how to adhere to compliance standards.

Establish showback/chargeback

With assets properly tagged and viewable by different groupings, you can then report on costs and show each LOB, cost center or owner exactly how much they're spending. You can identify your resource hogs and which applications are using and spending the most. You can begin to establish budgets for different departments and make sure everyone is held accountable for their usage and spending.

Find actively running cloud resources that are unused

One of the first ways in which enterprises save money is by finding and decommissioning actively running cloud resources that are no longer being used. Zombie infrastructure can include unused VMs, storage volumes that are no longer attached to a compute instance, old snapshots and disassociated IP addresses, just to name a few.

Build a baseline for utilization and performance

Now that you have visibility into the environment and all the applications and services that cloud assets support, cloud operations teams need to work with other constituents of the CCoE to establish baselines for utilization that are balanced with performance requirements.

The cloud operations team works with developers and cloud engineering teams to view performance data of the cloud assets and establish minimum/maximum thresholds for CPU, memory and disk utilization to establish how well-sized the environment is for the workloads it supports.

Organizations should track the following KPIs to measure success in phase one:

- The percentage of assets that meet tagging standards
- Resource utilization/performance (CPU, disk, memory, etc.)
- Misconfigurations identified per month per team
- The number of unsanctioned cloud resources in use by team/owner
- The number of deployed, yet unused cloud resources (zombie infrastructure) by application/team/LOB
- Adherence to budgets by application/team/LOB

Phase two: optimization

Once you have this visibility and a baseline for good tagging hygiene, proper configuration and required performance metrics, you can see where you have wasted spend and take steps to optimize the environment. Some best practices to get started with in this phase include the following.

Identify unsanctioned services, untagged assets and misconfigured assets

With input from software development teams and finance, identify any deployed assets that are not sanctioned or do not comply with configuration standards. For example, you can let users know what compute instances and instance families are not acceptable to deploy. You can also alert users to misconfigured or untagged assets, and give them instructions on how to fix.

Rightsize infrastructure

Based on the benchmarks established collectively with developer and finance teams, identify underutilized assets and move workloads to lower-cost resources. This starts with an out-of-band policy guideline that looks for assets with CPU, disk and memory metrics that are well below thresholds.

For example, a workload running on two AWS i3.8xlarge costs about $240 per day. If the instances are running well below established thresholds, the workload could potentially be moved to a lower cost instance, such as one i3.4xlarge that would only cost about $30 per day without affecting workload performance.

Identify and decommission zombie infrastructure

As mentioned earlier, one of the first ways companies reclaim cloud expenses is deleting deployed but unused infrastructure. A simple example is when a user deletes an Azure VM but neglects to delete the disk storage. When a VM is launched, disk storage is usually attached to act as the local block storage for the application. However, when you terminate the VM, the disk storage remains active, and Microsoft will continue to charge the full price of the disk, even though the data is not in use. Other common types of infrastructure that are at risk of becoming zombies are compute, snapshots, databases, IP addresses and load balancers.

Optimize committed discounts

Cloud operations works very closely with the cloud economist to put policies in place to optimize utilization of committed discounts. Every organization is different, so your use of committed discounts will vary. Some companies with a lot of consistent workloads running 24x7 will cover 90–100 percent of their compute environment with reservations. Others with very dynamic environments will use more on-demand and spot instances. A very general rule of thumb is to run only 25–30 percent of your compute environment on demand, with the rest covered by committed discounts or running on spot. The savings can be 60–80 percent depending on the commitment type.

In the visibility phase, you should have gained a view of how much of your infrastructure is covered by committed discounts, such as reservations or Savings Plans. A first step is to analyze your compute environment and try to increase the percentage of your environment covered by reservations. Secondly, you want to make sure you are utilizing your commitments properly.

It's a good practice for cloud operations to include a list of services covered by an existing reservation in their in-band guidelines for sanctioned services to encourage users to prioritize their usage. It's also important to have at least an out-of-band policy so you can continuously report on commitment utilization and make recommendations.

Organizations should track the following KPIs to measure success in phase two:

- The number of assets that do not meet configuration standards (wrong VM type, location, image, OS, tagging)
- The percentage of assets meeting or exceeding performance and utilization metrics (rightsizing)
- The percentage of infrastructure running on demand (in comparison to the percentage of infrastructure covered by reservations, Savings Plans, committed use discounts, spot, and the like)
- Reservation utilization
- Mean time to repair (MTTR) and mean time between failures (MTBF)

Phase three: governance and automation

Once you've gained visibility and established areas in need of optimization, you can begin to set policies using in-band and out-of-band guidelines and guardrails to help every user be an active participant in continuously optimizing the environment. Once policies are created, you can begin

to automate remediation actions. Governance and automation use both in-band and out-of-band guardrails to automate routine tasks and maintain the desired state with minimal human intervention.

Organizations should start slowly and focus on repetitive tasks first before deploying automation more broadly. It's important to segment automated actions into things that can be fully automated and those that require human intervention. Some best practices to get started with in this phase include the following.

Set policies and notifications for when assets drift from the desired state

Users can follow guidelines exactly and provision resources in compliance with standards but, over time, as the workloads evolve, assets can drift from their desired state. A simple example is an application that isn't used as much or some of its functionality is replicated with a new application. Consequently, its performance profile would change, and the policy would trigger a notification to the application owner with a recommended action.

Delete zombie infrastructure

One of the simplest, fastest ways to save time and costs is to automate the deletion of zombie infrastructure. For example, you can add an automated action to the policy that looks for unattached Amazon Elastic Block Store (EBS) volumes automatically notifies the owner the volume should be and/or will be deleted after five days. You can apply the same type of policy to snapshots, unused VMs, disassociated IPs, and the like.

One note of caution is to provide a mechanism for teams to opt out of the policy. Cathal Cleary, director of cloud services for VMware Engineering Services, set up policies to be opt-out instead of opt-in. This means that policies will apply universally to all assets unless they are tagged specifically to be skipped. For example, a policy might automatically delete all snapshots once they are 90 days old unless they have the tag value of nodelete. Of course, only authorized users can decide which assets qualify for the nodelete tag.

Terminate assets out of compliance with tagging policies

Cloud operations teams are often ruthless about tagging compliance and rightly so. Tagging is foundational to cloud management. Not tagging assets or not following tagging rules makes it nearly impossible to associate an asset with an owner, BU, cost center, or the like. Put simply, poor tagging compliance nearly renders the asset unmanageable. This is one of the instances where an out-of-band guardrail is required to automatically notify owners of improperly tagged assets and inform them they'll automatically be deleted after a set time if not remediated.

Automate commitment purchase recommendations

It's a very manual and tedious process for cloud operations to analyze their environment across multiple clouds and make informed purchase decisions. Using a cloud management platform, the cloud operations team should set a policy that evaluates the environment and automatically makes purchase recommendations. Other than setting up an approval workflow and making the actual purchase, the entire process of evaluating reservation usage and building a quote for purchasing a new reservation can be automated.

Automate security risk remediation

Security and compliance are their own dedicated functions, but cloud operations teams should work closely with security and compliance teams to share the burden. At a minimum, cloud operations teams can build default policies for their cloud providers' security best practices and the Center for Internet Security (CIS) best practices, and take automated actions should something violate those standards.

Integrate continuous governance to DevOps workflows

Cloud operations, in conjunction with development team leads, should embed operational governance policies early in the development lifecycle (shift left).

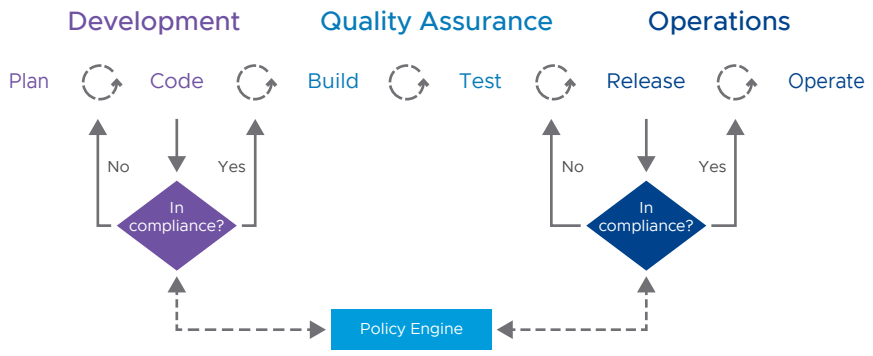

Figure 7: Example of embedding governance policies in the development lifecycle.

In this scenario, developers can make an API call to a policy engine at any point in the CI/CD process and know, well in advance of deployment, if their deployment will violate a policy. Depending on whether the policy is in-band, out-of-band, or a guideline or guardrail, they'll know what prescriptive action to take before they invest more time in their release process. Embedding continuous governance in the release process accelerates deployment times; makes the developer, quality assurance (QA), and release management teams more productive; and makes cloud operations staff more efficient.

Organizations should track the following KPIs to measure success in phase three:

- The percentage of policies in a compliant state
- The admin hours saved as a result of automated policies
- The number of reservations automated
- The number and cost of zombie infrastructure resources identified and remediated
- The time to remediate a security violation
- Service availability
- The time to deploy

Phase four: business integration

For many years, IT operations has been regrettably characterized as a cost center. This has often meant an overly narrow focus on utilization rates and expenses. Today, however, a company's entire technology stack is intertwined with an organization's day-to-day functions and has a direct impact on an

organization's success. The rapid availability of new cloud capabilities presents an even greater opportunity to influence business outcomes. To measure that impact, cloud operations teams need to align their cloud and operational KPIs with business KPIs. There are a couple ways to bolster that alignment. Some best practices to get started with in this phase include the following.

Integrate continuous governance into business applications and systems

The metrics and KPIs established in the CCoE and put into practice by cloud operations should also be embedded in other systems that run the business. In the earlier example of the CI/CD process, cloud operations and governance are embedded in the process from the start. This operational best practice can now be measured in terms directly related to successful application deployments. Operational metrics and KPIs can also be embedded within the DevOps toolchain in tools such as Jira, Jenkins and IT service management platforms. The objective is to know exactly how cloud operations, automation and continuous governance accelerate application deployments.

Align operational expense to profitability

As automation is put in place and begins to take hold, cloud operational metrics (such as MTTR, MTBF and tracked hours) can be measured in how they impact profitability. As an organization deploys more services to the cloud, it should be apparent how cloud operations contributes to profitability.

For example, if a team of 20 cloud administrators was able to manage 100 applications and $5 million of annual cloud spend in 2018, but one year later can manage 250 applications and $10 million of cloud spend, the team has made a huge contribution to the organization's bottom line and adoption of new cloud applications and functionality.

Charmaine Honeycutt, senior director of operations for Ziff Davis, said that effective cloud operations had clear benefits to the company's bottom line: "Not only are we able to benefit from considerable cost and resource savings, but we are more efficient and effective as an organization."

Measure effect of new cloud services on employee and cloud administrator productivity

Organizations often have a laser focus on cost metrics, but the prime reason for cloud migration is to take advantage of new services and focus on innovation. With a mature operational governance framework, cloud operations can safely adopt new cloud services quickly after cloud providers release them to market. Increasingly popular services, such as serverless, containers and managed platform services, help the operations team become nimbler. As a result, they can support new higher-value cloud services and new applications that save themselves and employees time in their day.

Streamline compliance and audits

Nearly every large company, especially those in regulated industries, has a governance and compliance office with a chief compliance officer that oversees adherence to industry regulations, as well as enforces internal company policies. It's a critical but often overlooked marker of organizational success. Cloud operations plays a significant role in creating a safe, secure and highly available cloud infrastructure, as well as in documenting everything and creating an audit trail that shortens the time and expense of audits.

Learn more

Learn more about streamlining cloud operations with VMware Tanzu CloudHealth®.

Organizations should track the following KPIs to measure success in phase four:

- Cloud operations expenses as a percentage of the cost of goods sold (COGS)
- The contribution of cloud operations expenses to profit margins over time
- The time to market for new services
- Developer and engineering productivity improvements
- Employee productivity enhancements from new cloud services
- The adoption of cloud native applications
- Customer satisfaction

Conclusion

The head of cloud operations has arguably the most cross-functional role within the CCoE. In addition to maintaining efficient operations, they act as a liaison between developers, engineers, LOB leaders, security and compliance teams, and the cloud economist. The head of cloud operations brings together input from all these teams and builds workable standards and policies that everyone can adhere to and iterate on continuously.

Cloud operations starts with gaining visibility into all cloud resources and bringing shadow IT into the light. Successful cloud operation functions build on that visibility with optimization and automation that continuously monitors and improves the efficiency and performance of the cloud environment, and aligns that success to the company's business goals. Cloud operations and continuous governance are a value-adding foundation to a company's digital transformation in the cloud and should be prioritized in every cloud journey.

Item No: 1703106aq-wp-bld-succ-cld-ops-gvnce-prac-a4 2/23