Hybrid & Multicloud Management

Five Strategies to Increase Agility & Efficiency in an Evolving Cloud World

Philip Trautman

Hybrid and Multicloud Management

Five Strategies to Increase Agility and Efficiency in an Evolving Cloud World

Philip Trautman

Hybrid and Multicloud Management

by Philip Trautman

This work is part of a collaboration between O'Reilly and Nutanix. See our statement of editorial independence.

Table of Contents

Preface

To succeed in the digital age, businesses of all types need to accelerate the delivery of new applications and services. To help speed innovation and improve service delivery, many enterprises are turning to a hybrid cloud model—with some IT infrastructure and services on-premises and some in one or more public clouds—with the goal of providing interoperability that makes it easy to extend, burst, or migrate operations between environments.

Many companies have progressed to a multicloud approach, using services from multiple public clouds in addition to on-premises IT. However, as enterprises progress on the path to hybrid cloud and multicloud, they struggle to integrate disparate cloud operations.

Who This Report Is For

Designed for IT leaders and business decision makers, this report will help you make smarter cloud choices, gain greater control over private and public cloud operations, and increase integration across clouds.

What You Will Learn

This report describes five key strategies to help you:

- Unify cross-cloud infrastructure management
- Accelerate application deployment and automate life cycle management
- Migrate applications among clouds more easily

- Implement consistent policy-based security and compliance everywhere
- Track and optimize private and public cloud spending

These strategies can eliminate operational silos and make your team more efficient while reducing the risk of planning errors, increasing operational agility, and giving you greater control over your cloud spending.

Hybrid Cloud Is Critical for Enterprise Agility and Efficiency

The 2020 global pandemic underscored the need for enterprises—and enterprise IT—to increase agility and accelerate the pace of private and public cloud adoption. With more and more of the world's economic activity taking place online, achieving cloud success has taken on new urgency as companies rush to fill gaps and take advantage of emerging opportunities.

In a digital economy, IT is a key enabler of success. To react quickly to changing business needs, provision apps more quickly, and support business on a global scale, IT teams need to take full advantage of both public and private clouds.

Private clouds can offer control over consumption, more layers of security, and the ability to run legacy applications that cannot leave your on-premises datacenter, whereas public clouds let you provision and scale resources in minutes, take advantage of innovative services, and eliminate the challenges of low-level infrastructure management. Businesses that can create an effective hybrid cloud are able to leverage the benefits of both public and private clouds (as shown in Figure I-1) with seamless interoperability across environments.

Figure I-1. A properly implemented hybrid cloud combines the advantages of private and public clouds.

Why Hybrid Cloud?

Whether it was intentional or not, most enterprises are already working toward a hybrid cloud model that combines on-premises datacenters and public clouds. And many need the ability to extend, burst, or migrate operations from private to public or vice versa, and to support traditional and cloud native applications. As enterprises seek to rationalize cloud operations, they prefer the hybrid cloud model because it combines the predictability and control of datacenter operations with the elasticity and agility of public cloud. This has a number of advantages:

Flexibility and agility

Hybrid cloud provides ready access to resources to support new applications, accommodate development and testing projects, or to address unanticipated needs quickly. On-premises operations address the majority of predictable resource needs with easy access to public cloud resources as necessary.

Elasticity

The hybrid cloud model gives you the ability to respond elastically to seasonal or other periodic fluctuations in resource requirements. Applications that experience big resource fluctuations may run best in a public cloud, where they can grab resources when they are needed and release them when they are not.

Cost control

With hybrid cloud, you don't have to equip datacenters to accommodate peak loads, which is both expensive and wasteful.

Instead, public cloud resources can be added during peak periods, meeting your needs while reducing expenditures. This is especially helpful for businesses that have seasonal or periodic spikes in demand. Predictable workloads are easily supported and more cost-effective on-premises, whereas elastic workloads can run on public cloud to achieve the best cost profile.

Optimized user experience
Delivering IT services to customers or employees from a distant datacenter introduces latency. Hybrid cloud makes it possible to expand quickly into new geographies, using public cloud resources to bring digital services closer to customers or employees.

Hybrid cloud promises better digital services, faster delivery of new features and applications, and happier customers and employees. However, although the potential benefits are clear, your actual results depend on the maturity of your hybrid operations. Many organizations find they face significant challenges achieving their cloud goals.

Is Your Private Cloud Meeting Your Needs?

A private cloud is a critical element of a successful hybrid cloud. However, many private clouds fail to meet enterprise needs due to:

Inflexible architecture
Your private cloud may need to adapt to a variety of application needs, from traditional enterprise applications to cloud native applications.

Inadequate performance and business continuity
In a poorly architected private cloud, one workload can have a negative impact on the performance of another. Tools may be inadequate to protect applications and data from downtime.

Difficult scaling
It can be difficult to predict when you'll hit resource bottlenecks, and scaling can be disruptive and expensive.

Complex data services
Diverse data storage needs can require a variety of hardware, adding cost and complexity. Separate storage pools decrease overall utilization.

Automation limitations
> An inflexible architecture and complex data services complicate automation and make automation failures more likely.

To get the full benefit from hybrid cloud operations, it's important to ensure that your private cloud is pulling its weight. Hyperconverged infrastructure (HCI) solutions—combining compute, storage networking, and virtualization into scalable building blocks—are proving to be a simpler, more effective infrastructure foundation for private cloud. The right HCI solution delivers the cardinal virtues of cloud, including self-healing, simplified capacity planning, easier automation, and reduced management overhead.

Hybrid Cloud Pain Points

As enterprises have operationalized hybrid cloud over the past several years, a number of pain points have emerged.

Operational Silos

The first and biggest problem that companies face as they move to hybrid cloud is that operations differ significantly from one cloud to the next. A majority of enterprise IT organizations have operational silos to manage private and public clouds. Separate teams are responsible for each cloud environment, leading to resource sprawl, wasted resources, and increased costs.

Complex Application Life Cycle Management

Advanced digital services are becoming far more complex to deploy and manage. A single application may depend on dozens of components. Operational silos, not surprisingly, only add to the complexity of application management. Deploying application components in different locations or different clouds becomes difficult and time-consuming and can require significant infrastructure, networking, and software expertise.

Difficult Migration and Lack of Application Portability

Application portability is a significant enterprise concern. An application that runs in one cloud can't simply be migrated or redeployed in another; it has to be replatformed and possibly rearchitected to operate effectively. Because of these challenges, many enterprises get

locked into particular cloud environments and are unable to make changes without disruption and additional investment.

Lack of Multicloud Security

Each private and public cloud environment has its own security model, making it difficult to ensure that your operations adhere to the same security policies and regulatory requirements everywhere, and increasing the chance for user errors and data breaches. Regulatory compliance has become so difficult that many companies outsource certifications in hopes that the outsourcing vendor will have the expertise to achieve certification across clouds.

One of the most common and persistent causes of data exposure, leaving data exposed in unprotected AWS S3 buckets, is largely the result of simple human error. Such breaches have affected many of the largest enterprises in the world—companies with significant cloud expertise and experience.

Cost Control

The challenges of cost control in the public cloud—including shadow IT, wasted resources, incorrect sizing, and reserved instance management—are widely recognized; getting them under control is a high priority. (These challenges are discussed more in Strategy 5.) Many organizations are surprised by the difficulty they have reining in public cloud spending. However, private cloud cost control is an important—and overlooked—savings opportunity. It can be extremely challenging to estimate and assign private cloud costs.

Once you've identified costs for each cloud, you still have to determine where resources are being wasted. And, ideally, you need to figure out where and how to run each application to optimize overall costs.

Summary

Enterprises have adopted hybrid cloud to increase agility, control costs, deliver a better user experience, and accelerate digital transformation. However, as hybrid cloud operations have grown in popularity, a number of significant pain points have emerged. These pain points can reduce the benefits of agility and elasticity. Strategies

to tackle each of these challenges are discussed in the sections that follow.

Key takeaways

- Operational silos reduce the effectiveness of hybrid cloud.
- Deploying applications or application components in different clouds can require significant expertise and manual effort.
- Significant re-architecting or retooling may be necessary to migrate an application between clouds.
- The fact that every cloud has its own security model increases the risk of data breaches and other cybersecurity threats.
- Cost control and cost optimization are among the most critical cloud initiatives.

Strategy 1: Create a Unified Infrastructure Control Plane

A fundamental challenge when it comes to managing infrastructure across a hybrid cloud is that every cloud—public or private—is unique. This strategy digs deeper into the challenges of hybrid cloud operations, suggests a plan of action for addressing those challenges, and looks at the landscape of possible solutions.

Hybrid Cloud Infrastructure Management Challenges

Although things may be conceptually similar between different cloud environments, at a practical level there's no standardization. Knowing what compute instance or what storage type to use for a workload in one cloud doesn't translate easily to another environment.

The fact that each cloud essentially operates as a silo causes challenges to multiply.

### Management Interfaces Differ from One Cloud to the Next

The tools for selecting, deploying, and configuring infrastructure services are different from one cloud to the next, and expertise in one cloud doesn't mean you'll be immediately efficient in another. Given the rapid growth in infrastructure and other services in each of the major public clouds over the past few years, just keeping up to

date with the capabilities and best practices for a single cloud platform has become a challenge.

Siloed Cloud Management Teams

Siloed cloud management is almost certainly why most enterprises have separate teams dedicated to each cloud platform they use, as illustrated in Figure 1-1.

Figure 1-1. A typical hybrid cloud consists of operational silos, with separate teams responsible for each cloud.

Lack of Integration

There's a surprising lack of integration between private and public clouds and across different public clouds. Public cloud providers mostly expect you to use their integration tools—and take widely different approaches to private-cloud integration.

There is still very limited interoperability between clouds or integration across clouds, since it's to a public cloud's benefit to lock you in. Many public clouds make it cheaper to move data in than out, and the hybrid cloud tools they offer are more about colonizing your datacenters with their technology than helping you achieve broader interoperability or integration outside their ecosystem.

These operational silos—although they may seem a necessity given current cloud realities—are an impediment to hybrid cloud success. Siloed teams make collaboration more difficult and ultimately slow progress.

Plan of Action: Unify Infrastructure Management

As a result of these challenges, most enterprises have an approach to hybrid cloud management that has more to do with trying to get the most from each cloud than it does with getting clouds to work together seamlessly. What's needed to fix this problem is a unified control plane that gives you the ability to monitor, manage, and orchestrate across all environments with a single set of tools. This is the only way to operate at the highest level of maturity and achieve the full benefits of hybrid cloud.

The following approach can help you achieve the level of integration among clouds that you've been lacking:

1. Choose a single control plane that you can apply everywhere, abstracting the differences between different environments. (Some available options are discussed shortly.)
2. Modernize your datacenters to use that control plane as broadly as possible.
3. Choose public clouds where you can use that control plane.

This strategy was the main focus of my report *Designing and Building a Hybrid Cloud* (O'Reilly 2018), where I referred to this concept as a cloud operating framework rather than a control plane, but the advice remains the same.

Identify your Goals

Creating a coherent, top-down plan results in a more unified and efficient hybrid cloud environment. Before choosing the best unified management approach, assess your environment and think about what your goals and priorities are. It's likely your goals include some or all of the following:

Eliminate silos
Eliminating operational silos will make your operations more efficient, but it's important to identify the highest priorities across your company's operations (or the operations that are under your control).

Unify management

You may need to unify management across all private and public clouds, or you may have just one or two high-priority cases in mind.

Standardize tools

Standardizing tools across environments increases flexibility, decreases the need for specialized skills, and offers a host of other benefits—even if your organizational structure remains siloed.

Increase pace of innovation

Simplifying management and enabling self-service can help developers and DevOps teams deliver products faster.

Increase cloud security

See Strategy 4 for more on multicloud security.

Reduce cloud costs

Maintaining separate operations teams for each cloud is expensive and results in a lot of redundancy. (See Strategy 5 for more on cost governance.)

What Are Your Must-Have Control Plane Capabilities?

Sometimes it can be helpful to create a list of must-have capabilities in addition to or instead of goals and objectives. Here's a short list of possible must-have capabilities for your unified control plane:

- Deploy infrastructure in every location from a central console
- Manage and monitor infrastructure life cycle from a central console
- Eliminate manual procedures/increase IT automation
- Support self-service for cloud users
- Simplify updates of operating software
- Provide support for virtual machines (VMs) and containers

This list is by no means comprehensive, but it should help your team think about its own list of must-haves. There's nothing worse than spending weeks or months testing—or even deploying—a solution only to discover a crucial capability is missing.

Choose the Right Solution for Your Needs

There are two general classes of solution for unifying the control plane across private and public clouds:

Unified cloud management
> Provides a layer of abstraction on top of existing cloud interfaces.

Unified platform
> Integrates closely with the underlying infrastructure, providing the same interfaces and services everywhere.

The best option for you will depend on your requirements. Here are some possible solutions in both categories:

Unified cloud management

Hybrid cloud management
> A wide variety of hybrid cloud management solutions have emerged in recent years. *The Forrester Wave™: Hybrid Cloud Management, Q4 2020* report compares nine available solutions based on a broad range of criteria, including provisioning, automation, monitoring, and orchestration.

Configuration management and orchestration tools
> There are a number of well-known tools in this class that are popular with DevOps teams. These tools may already be part of your infrastructure stack. Although some of these products operate both on-premises and in public clouds, they can require significant expertise to configure and use in each cloud. As a class, these tools may be facing significant headwinds in the DevOps domain due to the growing popularity of Kubernetes.

Unified platform

Kubernetes
> Kubernetes is a platform for automating and managing the execution of containerized applications, particularly cloud native applications that use a microservices architecture. Kubernetes by itself is not a unified platform, but Kubernetes services are now available in all major public clouds, or you can deploy your preferred Kubernetes software on the cloud(s) of your choice.

Therefore, you can use Kubernetes to provide essentially the same environment everywhere and run containerized applications in the location(s) of your choice without modification.

Although the control plane may be the same or at least similar for each cloud, Kubernetes by itself does not include the ability to manage multiple clusters across different clouds from a single control plane.

If your operations are focused entirely on containers, Kubernetes may be a good choice. However, if you need to support both VMs and containers, or your operations need a control plane with a greater level of abstraction, you will likely want a solution that incorporates Kubernetes but isn't limited to Kubernetes.

Platform vendor solutions

A number of prominent vendors in the IT space offer public cloud solutions that enable their platforms to run also in public clouds, creating a single unified platform spanning private and public clouds. Prominent examples are VMware and Nutanix.

Look for the ability to monitor and manage private and public cloud environments from a single control plane and to facilitate extending operations, bursting, and application migration among all supported environments.

Hybrid solutions from public cloud vendors

Solutions such as AWS Outposts and Azure Stack create hybrid clouds by extending public cloud capabilities to your datacenters. These solutions may be appropriate for hybrid clouds that connect your on-premises datacenters to a single public cloud in situations where you aren't worried about vendor lock-in.

If what you want is a seamless operating environment that encompasses both private and public clouds, you'll likely be best served by one of the unified platform options. Once you've identified a few candidates based on your requirements, the final decision may depend on which solution offers the most compelling roadmap and vision for your hybrid cloud.

Evaluating Solutions

Many of the strategies in this report, including this one, suggest classes of solutions that satisfy various hybrid cloud requirements. These lists are by no means exhaustive, and new solutions are entering the market all the time. When evaluating any hybrid cloud solution for suitability, keep these points in mind:

- Does the solution work with the private and public cloud technologies you already use or plan to use?
- How well established is the vendor, and how good is their support?
- If the solution is open source, does your team have the necessary skills to implement and support it?
- What skills will your team need to deploy and use the solution?
- Will professional services be necessary to implement the solution?
- What new features are on the vendor's roadmap and how does that match up with your future needs?

There are no perfect solutions yet. However, hybrid and multicloud offerings are evolving quickly in response to the clear needs that exist. Making smart choices now should pay big dividends in the future.

Summary

Current hybrid cloud operations suffer from siloed private and public cloud environments with separate management tools and dedicated management teams for each environment. As a result, management and integration capabilities across clouds are often minimal. A smart strategy is to choose a solution that provides a unified control plane, either through cloud management software or via a single platform available on premises and in multiple public clouds.

Key takeaways
- Unifying the control plane for your hybrid cloud will allow it to operate as a more cohesive entity, eliminating the need for siloed teams.

- Your decision should take into account your specific goals and the set of capabilities you require.
- Unified cloud management provides a control layer on top of existing clouds, abstracting their capabilities so that everything can be treated the same.
- Unified platforms integrate with underlying cloud infrastructure, providing the same set of interfaces and services everywhere for uniform operations with guaranteed compatibility.

Strategy 2: Streamline the Application Life Cycle

As cloud operations mature, there is a natural progression from infrastructure services to application services. Once you've unified your infrastructure control plane, the next step is to do something similar for applications.

This section explores the growing challenges around application deployment and management, suggests a strategy for automating the application life cycle, and looks at possible solutions to implement that strategy.

The Cloud Amplifies Application Challenges

Private and public clouds, digital transformation, and the need to increase the pace of innovation have all conspired to increase the pressure on IT organizations to deliver applications and services more quickly, creating new challenges for already stretched teams, as illustrated in Figure 2-1.
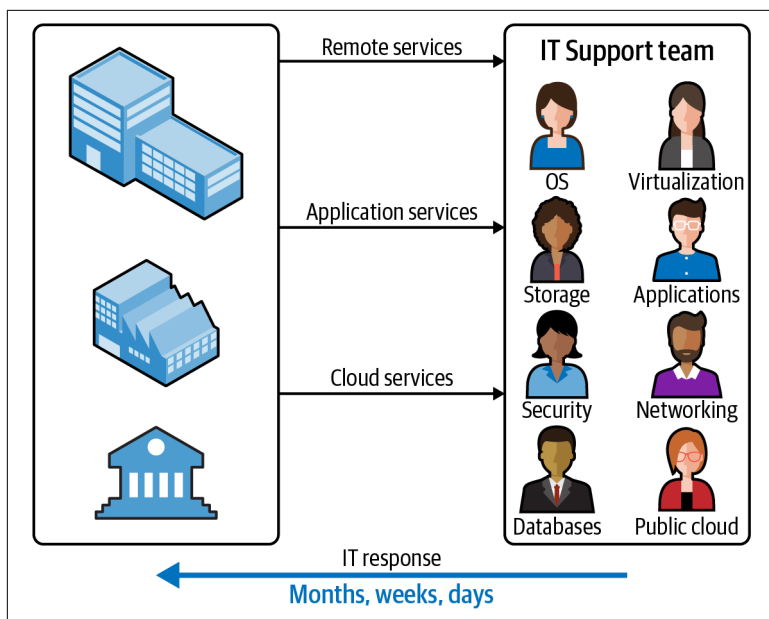
Figure 2-1. Enterprise demands for applications and services are increasing, adding stress on overloaded IT teams. Traditional approaches result in a delayed response to users and dissatisfaction.

Increasing Demands

Enterprise cloud users expect to access infrastructure and other services quickly if not on demand. For example, development teams need access to workspaces, testing systems, continuous integration/continuous delivery (CI/CD) pipelines, and other resources without delay to ensure that their work is unimpeded. If your hybrid cloud can't deliver the desired services, if the interfaces and processes are cumbersome, or service delivery takes too long, users may go directly to the public cloud. This shadow IT circumvents company governance and increases company costs.

Complex Deployments

A single application can have many dependencies and require significant expertise in multiple technical domains to install, configure, test, and deliver to production. Deploying the same application in a different cloud—assuming it's possible—can require different tools, processes, and significant additional effort.

Limited Resources and Low Productivity

At the same time the demands on IT are rising, teams are struggling with staffing shortages, significant skills gaps, and tight budgets that make it hard to meet ongoing needs despite often heroic efforts.

Scaling and Life Cycle Events

A successful application may need to scale during periods of peak load while freeing resources when they aren't needed to save cost. It can also be difficult to predict growth for many applications, making a simple way to scale deployed applications desirable. Application patches and updates can be frequent and need to be deployed with minimum delay or disruption.

Cloud Native and Traditional Enterprise Applications

In addition to the enterprise business applications your team has supported for years, your company may have a growing portfolio of cloud native applications, adding to the demands on your team—and creating a need for new skills that your team lacks and that are in short supply in the labor market.

Given these challenges, continued reliance on manual processes executed by administrators with years of experience and deep domain expertise is not the answer. With your IT team on the critical path for service delivery, you will never be responsive enough—and will fall further and further behind. Automation of application deployment and life cycle management can help your team succeed.

What About SaaS?

Much of the commercial software that enterprises rely on is available via a software-as-a-service (SaaS) model direct from the vendor. Well-known examples include Office 365 and Salesforce.com. Your company probably already uses a large number of SaaS providers, maybe more than you realize. SaaS is an important part of successful cloud operations—so important that there's now an emerging discipline called *SaaSOps*, intended to deal with the complexity that comes with using dozens or hundreds of SaaS services.

This report is mostly concerned with the set of applications and services that—for whatever reason—your company opts to run itself. In many cases, these will be custom applications developed internally, although they may include commercial software components such as databases and so on.

However, you should consider and use SaaS solutions wherever it makes sense; just keep a few common sense guidelines in mind:

- Standardize to the greatest extent possible the set of SaaS providers your company will use.
- Avoid having duplicate services from multiple providers, especially in the same geography. Does your company really need Slack, Jabber, Jive, and Microsoft Teams?
- Pay attention to the corporate data stored by SaaS providers. At a minimum, you need to make sure that your regulatory and data protection requirements are being met.

Look for opportunities to outsource applications you're running on-premises today to SaaS providers. Outsourcing can free up staff and infrastructure and reduce expensive licenses, allowing your team to address higher priorities. It may also allow you to recover valuable datacenter space and reduce both your capital and operating budgets. Just be sure you're making an apples-to-apples cost comparison. SaaS may not always be more cost-effective than running an application yourself.

Plan of Action: Automate Application Deployment and Management

As you expand from infrastructure to application services, your goal should be to capture the institutional and operational knowledge of your team using intelligent, automated tools so that application deployments and other life cycle processes can be repeated with a minimum of technical understanding or user inputs.

Done correctly, automation and self-service not only shorten delivery times and avoid the risk of human errors (see Figure 2-2), they encapsulate best practices and ensure that policies for governance, security, and data protection are applied correctly every time. Many teams find that as they automate and shift to a self-service model, the rigid team structures of the past no longer make sense, and they naturally evolve toward more fluid organizational structures and less reliance on specialists to keep everything functioning.
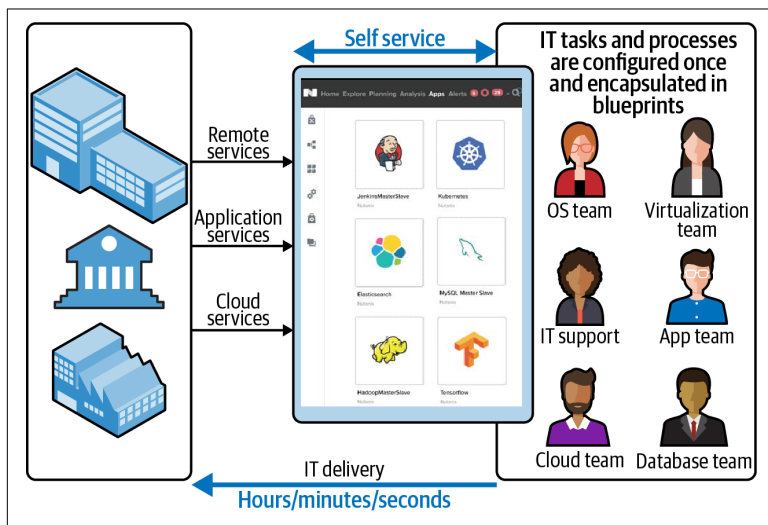
Figure 2-2. Automation and self-service can speed user access to critical services, eliminate risks, and get your IT team out of the critical path for service delivery.

A number of tools are emerging to accomplish this task—often referred to as application orchestration—resulting in an application blueprint, as illustrated in Figure 2-3. For each application, a blueprint must define all the essential elements of the application stack, such as:

- Will the application run in containers or VMs?
- How many instances of each type are needed?
- What storage services does the application require?
- What are the networking requirements?
- Does the application require a firewall or additional firewall settings?
- Is a load balancer needed?
- What backend database(s) and other services are needed?
- Which cloud environment(s) will the application be deployed in?
- Which protection policies should be added?
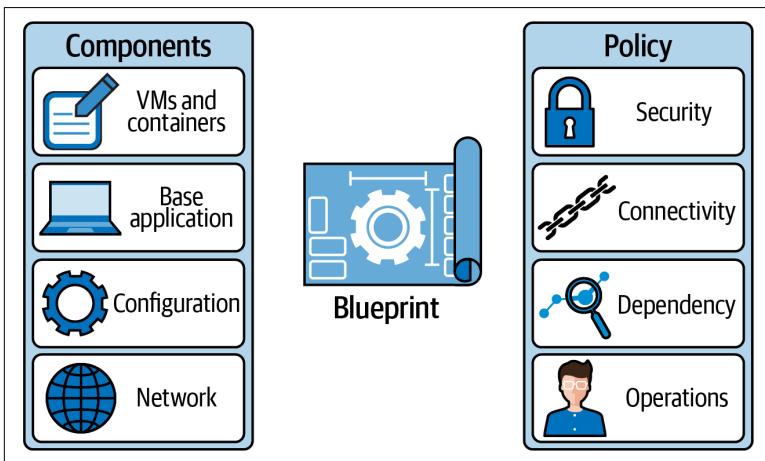
Figure 2-3. An application blueprint encapsulates the full set of components and policies necessary to deploy an application.

There are many potential variables. Different tools approach tasks in different ways, and each has its own strengths and limitations. Some provide visual interfaces, while others require coding skills to use effectively. It's important to understand your requirements and capabilities before choosing a solution.

Identify Your Goals

The first step is to assess the current (and future) private and public clouds your company is using and think about what your goals and priorities are. Your goals probably include some or all of the following:

Accelerate application delivery
> Although this may seem like a given, it may not be your highest priority.

Standardize deployments
> If every application deployment is a snowflake, it complicates your IT environment, increases the risk of human errors, and makes troubleshooting more difficult.

Eliminate manual tasks
> A large percentage of your team's time may be devoted to performing manual infrastructure and application maintenance, management, and monitoring tasks. Breaking out of this pattern will enable you to focus more attention on higher-value projects.

Enable self-service
> One of the biggest impediments to velocity in software development teams is slow access to resources. Enabling developers, business teams, and other IT users to access the infrastructure and application services they need on demand increases the pace of innovation.

What Are Your Must-Have App Orchestration Capabilities?

Creating a list of must-have capabilities is a good way to ensure that you don't come up short after a lengthy selection process. Examples include:

- Flexibly deploy full applications, application components, and/or infrastructure components from a central console
- Enable application life cycle operations (upgrade software, scale up/down, terminate and clean up) from a central console
- Scale up or down as needed (manually or automatically in response to one or more triggers)
- Provide application/component self-service
- Deploy components of the same application instance across multiple clouds for resiliency, scaling, and lock-in avoidance
- Support containers and VMs
- Incorporate security and data protection
- Provide visual design tools or code-based interfaces or both
- Integrate with other tools and services in your environment such as IT service management (ITSM), DevOps tools, CI/CD, and so on

Choose the Right Solution

The solutions available for multicloud application orchestration may overlap with those described for multicloud infrastructure management in the previous strategy, so it may be possible to choose a solution that satisfies your needs in both areas.

Many of the available solutions for application orchestratioion offer predefined blueprints for deploying popular services and applications such as web servers and databases. You may be able to use these as is or modify them for your needs. It may be worth looking for solutions that support the applications your organization uses most:

Hybrid cloud management software
> Some of the hybrid cloud management solutions as described at the end of the previous strategy include a degree of application orchestration.

Configuration management and orchestration tools
> These tools may have a high learning curve, but they can be a smart choice for organizations that already have staff with expertise.

Multicloud Kubernetes orchestration
> Kubernetes is designed to orchestrate containerized applications, particularly cloud native apps. Kubernetes services are now available in all major clouds, or you can deploy your preferred Kubernetes software on the cloud(s) of your choice. Kubernetes by itself lacks a control plane that spans different cloud platforms.

Platform vendor solutions
> Platform vendors that offer hybrid cloud solutions may also have tools for application automation and orchestration.

Look for solutions that offer appropriate design tools for your needs, whether visual or code-based or both, integration with the other tools you use, and ability to support life cycle operations, and any other requirements you identified as you read this section.

Summary

Cloud users such as developers and architects increasingly need fast and easy access to application services in addition to infrastructure services. The hybrid cloud model and the growing complexity of application stacks make manual deployment methods time consuming, error-prone, and expensive. Application-level automation and orchestration is quickly becoming essential.

Key takeaways
- A variety of tools are available to facilitate application deployments—both in a single cloud and across multiple private and public clouds.
- Evaluate your organization's requirements and skills to ensure that you choose a solution that is appropriate for your needs.

- If you have a significant investment in existing automation using scripts or other tools, you should ensure that you can continue to use them in conjunction with any new tools you select.

Strategy 3: Migrate Applications More Easily Among Clouds

The difficulty of migrating applications between different cloud environments is a significant concern for many enterprises. In a recent survey of cloud users, 73% said application portability was an issue, noting:

> Organizations are also facing significant challenges when moving applications....What is preventing them from moving more applications to public cloud is the need to re-architect or re-platform applications when moving them across clouds (75%) or the complexity of the migration (71%).

This strategy explores the challenges around application migration and mobility, and suggests a plan of action that provides the ability to move applications more freely.

Migration and Portability Challenges

Many of the migration and portability challenges that enterprises encounter involve traditional enterprise applications, but challenges can arise with cloud native applications as well.

Traditional Enterprise Applications

Most enterprises have hundreds of applications—both commercial and in-house—that are based on traditional development methods. In many cases, it doesn't make sense to go through the significant effort to modernize these applications, but you may still want or need the ability to run them in a public cloud rather than—or in addition to—on-premises. The associated challenges can include:

Multiple dependencies

Many applications have dependencies on services and capabilities that are available in your datacenter but may not be available in the cloud. For example, many enterprises use data management and data protection capabilities built into enterprise storage such as snapshots, replication, and data reduction. Some applications may be written to use these and other capabilities of enterprise infrastructure. Identifying these dependencies late in a migration will increase the migration timeline and cost.

Licensing

For commercial application components, licenses may not be portable from your datacenter to the cloud.

Post-migration management

You may use a variety of management and monitoring tools and other services on-premises. When an application moves to the cloud, you'll have to identify equivalents. For example, datacenter storage often includes accelerated snapshots, cloning, and replication that may not exist in your chosen cloud. If the data protection strategy for an application depends on these services, you'll need alternatives.

Note that solving these challenges for migration to one public cloud does not ensure that your application will be portable to other public clouds. Bi-directional mobility between private and public clouds is needed for the greatest flexibility. Solutions should automate the transfer of networking settings and provide the ability to test a migration prior to cutover.

Cloud Native Applications

Cloud native is the term for applications that take advantage of modern tools and methods, including containers, a scale-out microservices architecture, and agile development practices. Although we tend to assume that these applications run in the public cloud, they can also run on-premises, given the right infrastructure services.

Cloud native application design can result in applications that are more portable, but there are still impediments to easy application movement:

Dependencies

One of the things that's attractive about the public cloud is the availability of a variety of services that can be incorporated as part of an application to simplify and accelerate the development process. However, this can cause an application to get locked into a particular cloud. To move the application on-premises or to a different public cloud, you'll have to find alternatives to those services and re-architect.

Your cloud native applications may also need to access and update data from backend systems of record, creating additional dependencies that complicate migration.

Orchestration

If your cloud native applications run on Kubernetes, you may need the identical version of Kubernetes running everywhere. Some teams deploy Kubernetes in public clouds for portability instead of using services such as Amazon EKS, Azure AKS, or Google GKE.


Public Cloud Services: Pros and Cons

There's no question that the public cloud can deliver immense value to enterprises, from elastic and nearly inexhaustible infrastructure services to popular higher-level services such as AWS Lambda and Azure Cosmos DB. However, as with everything, it's important to think carefully about what you're getting, what you may be giving up, and what the future implications are. This is especially true with higher-level services.

For example, AWS Lambda provides serverless operations. You hand Lambda a function, and it executes that function whenever it is called, scaling automatically to address load. This approach is fast, simple, and free from infrastructure management headaches. However, although similar services are available in other public clouds, if you use Lambda—or any of the other specialized services across the public cloud—that necessarily increases your degree of lock-in. Does that mean you shouldn't use these services? Certainly not. Just make sure that you weigh the risk versus the reward and think about your exit strategy if you need one.

Plan of Action: Enable Easy Application Mobility

Your goal when it comes to application migration is obviously to make it easy—even trivial—to move applications where you want them, as illustrated in Figure 3-1. In this case, the results you achieve will depend on the choices you made for Strategy 1 and Strategy 2, described earlier. With the right unified infrastructure management and unified application management, migration across clouds becomes much less challenging. Note that this is the only strategy in this report that's wholly contingent on earlier recommendations.

Figure 3-1. Enterprises want easy, bi-directional application mobility among all the clouds they use.

Migrating Traditional Enterprise Apps

The choice of *unified platform* versus *unified cloud management* in Strategy 1 has direct implications for application migration:

Unified platform

If you choose a unified platform approach—the same infrastructure platform everywhere—to provide a universal control plane, you'll be able to move applications between clouds with far less concern about dependencies, management, and so on. Essentially, the process becomes the same as moving an application from one datacenter to another, a task that your IT team already knows how to accomplish.

Once you have the application VMs and data in the desired cloud location, you should be able to shut down the application in the original location and start it up in the new location. The data movement can be accomplished using asynchronous replication, or you can employ disaster recovery tools to failover the application, avoiding a service outage.

Unified cloud management

> If you choose a unified management approach—a layer of abstraction on top of existing private and public cloud interfaces—you will have more work to do with regard to dependencies across environments. Some cloud management solutions include migration planning tools to simplify the migration process.

> Alternatively, the right application orchestration solution (see Strategy 2) can abstract the differences between one cloud environment and another, enabling your application to be redeployed quickly in a new environment.

Networking

An important consideration for either migration scenario is networking. When you migrate applications, your solution should include the necessary network integration. If you're using a unified platform, here are three important things to consider:

Datacenter connectivity

> In a hybrid cloud, applications will likely need to connect efficiently between your datacenters and cloud locations. IP address management (IPAM), network portability and connectivity, and routing are important considerations.

Virtual networking

> All the public cloud providers implement some type of software-defined virtual networking (referred to as a virtual private cloud [VPC] in AWS, or VNet in Azure). Virtual networks allow for application connectivity, firewall security segmentation, and other networking features to be defined and managed independently from the underlying physical infrastructure or location. Unified solutions should provide similar constructs and management across on-premises and cloud locations to ensure portability and ease of connectivity between locations and applications. Some platform solutions run in a dedicated virtual network and may require specialized networking overlays that can't take advantage of native cloud networking features. Connecting to applications running in other virtual networks may be slower.

Public cloud services

> Often, an important reason for migrating an application is to take advantage of a particular public cloud service (see "Public Cloud Services: Pros and Cons" on page 21). Your migrated application should be able to access cloud services at native speeds or as close to them as possible.

If you're using a unified cloud management approach, your application blueprint should encapsulate the application's networking requirements, which ensures that the environment is configured appropriately.

Moving Cloud Native Apps

An advantage of the cloud native approach is that containerized apps (should) have fewer external dependencies, so it becomes more a question of starting the application in the new location rather than migrating it. The process is trivial in a unified platform environment—assuming your app isn't locked in due to use of cloud-specific services—since the tools are the same everywhere.

Summary

The traditional approach to application migration between clouds is based on modifying an application or the management of the application to be compatible with the new environment. However, if you've chosen a unified platform approach as your solution in Strategy 1, this becomes unnecessary. An application that will run in your private cloud will also run in your public cloud environment and vice versa.

If you've chosen unified cloud management as your solution in Strategy 1 and a unified application management tool in Strategy 2, this will also streamline application migration, although it won't be as simple as having the same platform everywhere.

Key takeaways

- If you think carefully about your needs as you choose solutions for unified infrastructure management and unified application management as described in the previous two strategies, application migration between disparate clouds becomes much simpler. These strategies help enable application migration.

- A unified platform approach provides the same environment everywhere, making migration trivial.
- A unified cloud management approach in conjunction with unified application management simplifies the process of application migration and can reduce or eliminate the need for re-platforming and refactoring.

Strategy 4: Enable Consistent Security Policies Everywhere

As your cloud operations expand to encompass multiple private and public clouds, security challenges can increase exponentially. When the Flexera 2020 State of the Cloud survey asked about top cloud challenges, security ranked number one (chosen by 83% of responders). Compliance also ranked high (76%).

This strategy explores why the challenges around security and compliance are growing, and it suggests an approach for simplifying security management for hybrid cloud.

Hybrid Cloud Amplifies Security Challenges

As with infrastructure management in general, the default for hybrid cloud security is a siloed approach that increases complexity and risk.

Security Silos

Each cloud has its own security model and tools. To ensure security, you therefore need people on your team skilled in the security tools for each cloud. This gets complicated quickly:

- With different tools in each cloud, there's no standard approach to security monitoring or remediation.

- With no global view of security across your hybrid cloud environment, applying the same security policies everywhere is a manual process.

- Changing security policies globally becomes complicated and error-prone.

Increased Compliance Demands

The lines between security and compliance can sometimes seem blurry, so it's useful to think of security as internally driven, based on your organization's own assessment of its digital protection needs. Compliance, on the other hand, is usually driven externally by government regulations or contractual obligations.

No matter what industry you are in, it's likely that your business is subject to regulatory compliance. For example, all companies that operate in Europe have to comply with the European GDPR regulations, which took effect in 2018.

Public Cloud Security Is a Shared Responsibility

You probably already know that private cloud security models are different than those in the public cloud, but there is another important difference that often goes underappreciated. Public cloud security is based on a *shared* responsibility model.

If you're using resources in a public cloud, certain aspects of security are provided and managed by the cloud provider. But how you configure and use security within the provided security framework remains up to you.

For example, if you provision an S3 bucket in AWS, it's up to AWS to secure the infrastructure underlying that storage bucket, but it's up to you to ensure that the bucket has the appropriate access control and policy configured.

Organizations that don't fully understand where shared security responsibilities begin and end will inevitably have poorer security outcomes.

If you're in a highly regulated industry such as finance or healthcare, compliance is probably something you factor into all IT decisions. You may need to ensure compliance with well-known regulations such as HIPAA in healthcare or PCI-DSS and GLBA in financial services.

Proving that your business complies with applicable regulations has become a complicated task and requires significant attention. And it becomes more challenging with every new cloud environment you support.

Growing Risks of Human Error

As your hybrid cloud environment becomes more complex and security needs more stringent, continued reliance on manual security configuration only adds to the risks. As already noted, one of the biggest causes of data breaches is human error.

Every vendor and cloud has unique security controls—with no established guidelines for secure integration among them. This makes it difficult or impossible to apply security policy in a consistent way everywhere and contributes to configuration errors.

Increased Threats

A final consideration where security and compliance are concerned is that threats continue to mount. With a large percentage of employees now working from home—and accessing resources across your hybrid cloud—the number of possible attack vectors has increased dramatically. How do you protect at-home systems from phishing attacks or ensure endpoint security, version control, or anti-malware protections? Research firm Cybersecurity Ventures estimates that the global costs of ransomware alone will reach $20 billion by 2021.

And yet according to the CyberArk 2020 Remote Work survey, "40% of organizations have not increased their security protocols despite the significant change in the way employees connect to corporate systems and the addition of new productivity applications."

Plan of Action: Implement Global, Policy-Based Security

Security silos are just as bad for hybrid cloud operations as management silos are—and potentially even riskier for your company. Your goal should be to find solutions that abstract the differences between various cloud security models to provide global visibility of your company's security posture, as shown in Figure 4-1.
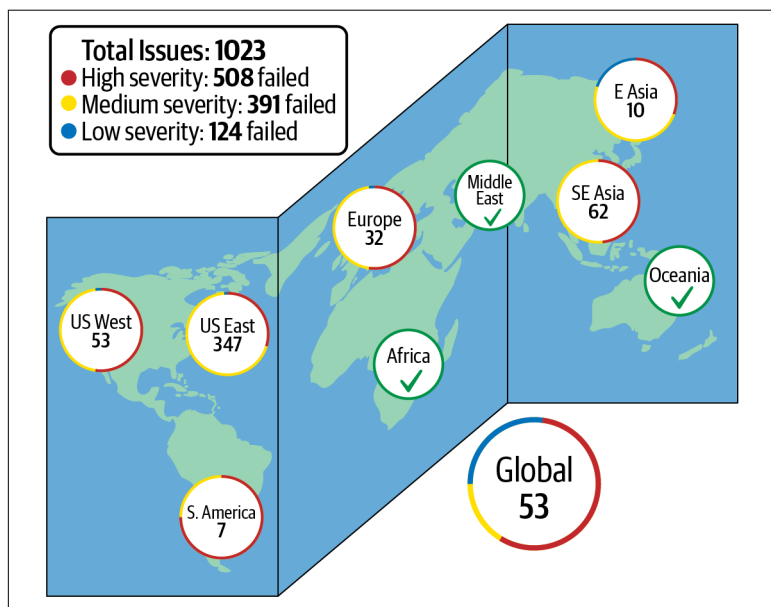
Figure 4-1. The right multicloud security solution should provide global visibility of vulnerabilities.

The capabilities of multicloud security management tools vary widely, and each has its own strengths and limitations, so it's important to understand your requirements before picking a solution.

Identify Your Goals

When choosing security and/or compliance tools, it's important to assess your environment and think about what your goals and priorities are. Your goals may include some or all of the following:

Unified view
> Visibility of your company's security posture across all private and public clouds, including the ability to drill down into particular locations and issues.

Policy and controls
> Ability to manage and enforce consistent security controls across all environments.

Real-time detection
> Ability to identify and flag vulnerabilities as they arise versus hours or days after the fact.

Alerting
> Notification to appropriate personnel when vulnerabilities arise.

Compliance auditing
> Ability to audit private and public cloud environments for compliance with regulations such as HIPAA, NIST, PCI-DSS, etc.

Reporting
> Ability to generate regularly scheduled security and audit reports to demonstrate compliance.

Remediation
> The ability to correct common vulnerabilities without operator assistance—or at least identify necessary corrective actions.

Extensibility
> No tool or set of tools is likely to address all your requirements out of the box. Therefore, it may be important to choose a tool that can be extended to address unique needs now and as they arise in the future.

Choose the Right Solution

Some of the solutions available for multicloud security and compliance overlap with those described earlier for multicloud infrastructure management. Although it may make sense in terms of your team's learning curve and the total cost to choose a single tool that satisfies multiple needs, be careful not to sacrifice key capabilities in the process:

Hybrid cloud management software
There are a number of hybrid cloud management vendors that offer integrated and add-on multicloud security capabilities.

Platform vendor solutions
If you opted for a platform vendor that provides a unified control plane in Strategy 1, by definition you'll have the same (or a very similar) security model across environments, simplifying security management. Security is an important consideration when evaluating unified platform solutions.

Other solutions

There are too many solutions in the multicloud security space to enumerate them. Do your homework and make sure you pick a solution that will grow with your needs.

Is It Time for Zero Trust?

Traditional perimeter-based security is becoming inadequate. Many organizations are adopting a zero-trust approach in which an application, service, or user only has access to the resources that are required to perform a specific task. Zero trust grants least-privilege access to resources based on application requirements or the identity of the user requesting access.

Zero trust at the network level relies on knowledge of application network communications, authentication, security analytics, and microsegmentation to create fine-grained access control policies for individual workloads. Adopting zero-trust methods minimizes the attack surface, prevents malware spread, and improves audit and compliance visibility, reducing the risks and impacts of security breaches.

Summary

With new malware attacks and data breaches occurring with ever greater frequency, security and compliance remain among the top hybrid cloud concerns. However, only about one-third of enterprises have adopted hybrid or multicloud security management tools. Deploying one (or more) of the available security tools is the best way to gain a global view of security and compliance and improve your company's overall security posture.

Key takeaways

- A variety of solutions are available for hybrid and multicloud security management.
- Multicloud vendor platforms, as described in Strategy 1 provide the same security model everywhere.
- Evaluate your organization's requirements and skills to ensure that you choose a solution that is appropriate for your needs.

- Advanced capabilities such as automated compliance audits and automated or one-click remediation can enhance security and compliance.

- Extensibility to accommodate unique organizational requirements may be important for growing hybrid and multicloud deployments.

Strategy 5: Track and Optimize Private and Public Cloud Spending

Failure to meet cost targets is a consistent cloud complaint. A multicloud IT environment is complex with a wide variety of services, as illustrated in Figure 5-1. Cost overruns can result from excessive upfront costs, unexpected scaling costs, and wasted or oversized resources and services. Cloud users are often better at provisioning resources than they are at releasing them. This can be due to lack of visibility, mistakes, or neglect. To control usage, reduce cloud spend, and avoid wasting resources, you need to be able to monitor and meter actual costs, identify and eliminate unused resources, and identify and rightsize underused resources across your public and private clouds. You also need to be able to optimize your choices based on actual consumption.

This strategy explores some of the challenges around hybrid cloud costs—including the significant differences between private and public cloud—and it suggests an approach for implementing effective cost governance.
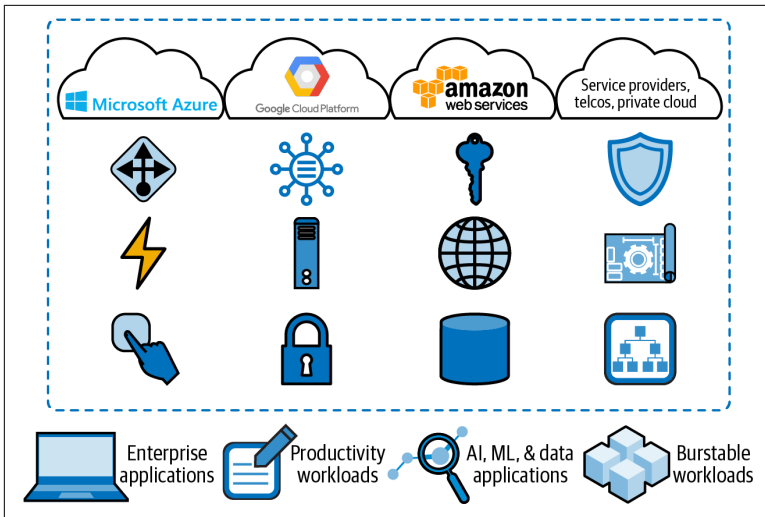
Figure 5-1. Today's multicloud environments include a diverse set of services supporting both traditional and modern workloads. Failure to manage this complexity leads to increased costs.

Hybrid Cloud Cost Governance Challenges

To understand the cost governance challenges created by hybrid cloud, it's important to look first at unique public and private cloud challenges.

Public Cloud Challenges

Most enterprises are already well aware of how quickly public cloud spending can increase. There are several significant aspects to public cloud spending:

Many deployments
> Many IT organizations have dozens of cloud accounts, using thousands of cloud resources—across multiple zones and regions—with new ones added all the time. Because it is so easy to initiate new deployments, it can be difficult to track and control consumption.

Shadow IT
> Developers and line-of-business teams are often tempted to use a corporate credit card and create their own cloud deployments

to avoid constraints and save time. These are completely outside of corporate governance and control, creating potential risk.

Resource waste
As noted previously, people are better at grabbing resources than they are at releasing them. In addition, many cloud resources are overprovisioned, using VMs that are larger than the workload requires or a higher tier of storage than necessary, creating significant opportunities for cost savings through right-sizing.

Reserved instance management
Most public clouds have provisions for reserved instances—VM or database instances that you contract for up front at a significant discount for a period of one to three years. It's important to be able to forecast reserved instance needs accurately. If you're too high, you end up leaving money on the table, and if you're too low, you could end up paying significantly more for on-demand instances to make up the difference. (I've used AWS terminology here. Azure offers Reservations, and Google Cloud Platform has sustained use discounts [SUD] and committed use discounts [CUD]. The concepts are similar. AWS is itself in the process of phasing out AWS Reserved Instances in favor of *Savings Plans* that offer greater flexibility.)

Private Cloud Challenges

Although it may not be that simple to parse, with a public cloud your monthly bill gives you almost everything you need to know about your cloud costs. However, it is surprisingly difficult for IT teams to estimate true costs accurately for a private cloud or to identify resource waste (see Figure 5-2).

Private cloud cost challenges include:

Accurately modeling TCO
To understand private cloud costs, you need a TCO model that accurately accounts for a whole range of overhead costs, including hardware, software licensing, facilities, networking costs, one-time and recurring service costs, and staffing costs. Once those costs are understood, you also need a way to allocate those costs accurately to the various resources (VMs, storage capacity, application services, etc.) being consumed by development

teams, departments, lines of business, and other cloud users. For example, to charge a department for the VMs it uses, you need to have some way to allocate a fraction of your total costs to each VM.

Resource waste

Although the cost impact of inefficient resource use is less obvious in a private cloud than in a public cloud, you still have to be able to identify underutilized or abandoned resources and right-size them or return them to the pool. Otherwise, you exhaust your resource pool prematurely and end up increasing datacenter spending.
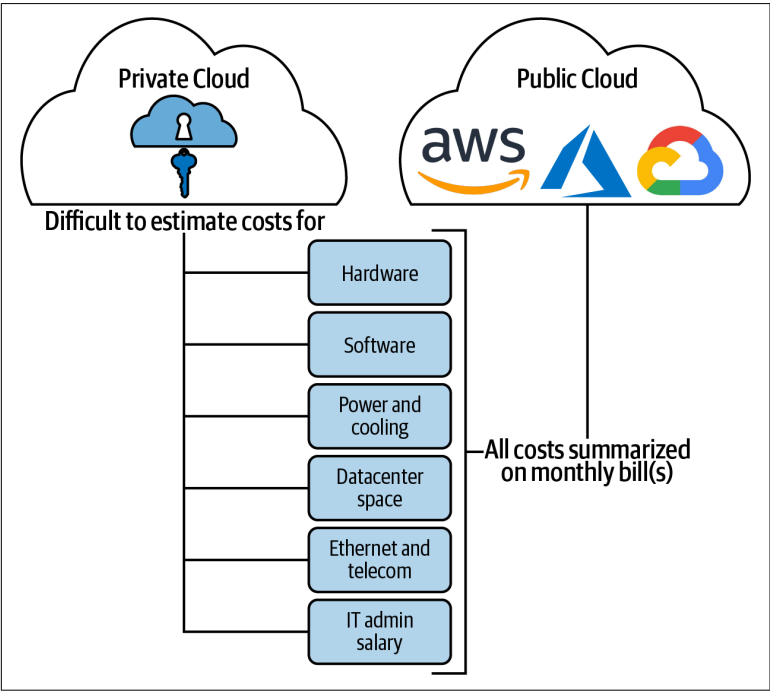
Figure 5-2. Determining your total costs for a private cloud can be challenging because there are many separate cost elements that must be accounted for to enable chargeback and provide a valid comparison against public cloud costs.

Cost Comparisons Between Private and Public Clouds

Even if you're able to understand and control private and public cloud costs, you'd also like to be able to compare costs across different cloud environments to optimize deployments. This sounds simple enough, but in practice these comparisons are often difficult. For example, it's impossible to know at a glance what on-premises VM configurations are good equivalents for specific AWS instances such as *t3.large* or *a1.4xlarge*. Or what AWS storage is the equivalent of your on-premises flash array? Without answers to these types of questions, it can be difficult to compare costs or size the infrastructure needed for workloads moving from one cloud to another. You need some way to *normalize* cross-cloud comparisons to ensure that you're not comparing apples to oranges and making inaccurate decisions. It's the only way to operate a hybrid cloud environment efficiently at scale.

Given these complexities, it's perhaps surprising that the spreadsheet remains the most common tool for managing cloud costs.

Plan of Action: Unify Hybrid Cloud Cost Governance

At the risk of sounding like a broken record, it should be clear from the challenges described that you need more than diligent application of Excel for successful hybrid cloud cost governance. There are a number of tools that provide hybrid and multicloud cost management. Your goal should be to implement tools that provide cost governance across your private and public cloud deployments. This provides a single point of management and optimization.

The capabilities of the available tools vary widely, and each has its own strengths and limitations.

Identify Your Goals

When choosing cost governance tools, it's important to assess your environment and think about what your goals and priorities are. This may include some or all of the following:

Unified view of costs
    Visibility of costs across all private and public clouds, including ability to identify anomalies and drill down as necessary.

The right granularity and aggregation
> You may need to view costs down to the level of business units, teams, or even individuals—across or within clouds. You may also want to be able to aggregate and roll up costs to see the bigger picture.

Identify cost overruns
> Identify and alert regarding excessive or anomalous resource consumption and cost overruns so corrective action can be taken before things get out of control.

Identify and remediate waste
> Policy-based identification of resource waste and automated right-sizing and resource reclamation can prevent unnecessary overruns. For example, you may want a policy that reclaims VMs that have been idle for a week or more.

Implement and track budgets
> The ability to establish budgets for a team, department, or a project; track adherence; and receive alerts can help keep costs in line.

Reporting
> Chargeback reports are useful for controlling cloud resource consumption and encouraging cloud users to be more diligent about resource usage.

Planning
> You may need a cost governance tool that can help you make informed purchasing decisions based on team needs and business objectives.

Tools that provide cost visibility across both private and public clouds simplify cost tracking and make it possible to cost optimize private *and* public cloud operations.

Choose the Right Solution

There are a number of solutions for cloud cost management. Many
of these offer multicloud support, but relatively few are designed for
both private and public cloud cost governance.

Hybrid cloud management software
> Several popular cloud management solutions include cost man-
> agement functions or have an option to add the capability.

Dedicated solutions
> There's a lot of activity in this area, so a definitive list would be
> impossible. Notable options from well-known infrastructure/
> platform vendors include VMware CloudHealth, NetApp Cloud
> Insights, and Xi Beam by Nutanix.

Summary

Hybrid cloud deployments make cost governance much more complicated. However, most IT teams still rely on spreadsheets. Choosing carefully and deploying one (or more) of the available cost management tools is a smart way to get spending under control and optimize your private and public cloud spend.

Key takeaways

- A variety of tools are available for hybrid and multicloud cost management.

- Evaluate your organization's requirements and skills to ensure that you choose a solution that is appropriate for your needs.

- Advanced capabilities such as reserved instance analysis and the ability to identify resources that are over-provisioned or aren't being used can substantially reduce your cloud spend.

What's Your Hybrid Cloud Strategy?

Enterprises are accelerating public cloud adoption and working to take full advantage of private and public cloud resources as they seek to speed up delivery of new and improved digital services. To succeed, you need a hybrid and multicloud strategy that works *for* your company, not against it. This means eliminating operational silos and gaining global control and visibility.

This report has outlined five strategies that can help you achieve greater hybrid cloud success in less time and with less effort:

Strategy 1: Create a Unified Infrastructure Control Plane
> Create a single cloud control plane, either via unified cloud management or a unified platform that doesn't lock you into a single public cloud.

Strategy 2: Streamline the Application Life Cycle
> Streamline application deployment and management with application orchestration.

Strategy 3: Migrate Applications More Easily Among Clouds
> Simplify the process of moving applications across private and public clouds.

Strategy 4: Enable Consistent Security Policies Everywhere
> Enable a comprehensive view of security with consistent security policies everywhere.

Strategy 5: Track and Optimize Private and Public Cloud Spending
   Take control of private and public cloud spending and optimize costs.

It's up to you to identify your organization's pain points and implement the strategies that will deliver the most immediate benefits. The more of these strategies you implement today, the quicker your hybrid cloud will reach maturity, giving your company an advantage in the race to make the most of digital.

About the Author

Philip Trautman has more than 25 years in the IT industry. Philip's career has focused on understanding and writing about enterprise IT infrastructure and cloud. He was senior manager of technical support for Auspex Systems before becoming an industry consultant and writer in 1997. His areas of expertise include storage, data protection and disaster recovery, computer architecture including converged and hyperconverged infrastructure, server and desktop virtualization, and cloud. Philip has done extensive work for current and past industry leaders including Data Domain, LSI, Legato Systems, Microsoft, NetApp, Nutanix, SGI, and SUN Microsystems. He has authored hundreds of white papers, ebooks, success stories, and other material for these and other clients.